Context-Aware Itinerary Planning in Smart Tourism: IoT-Enabled Design and Tourist Profiling

Claudio Marche, *Member, IEEE*, Vlad Popescu, *Member, IEEE*, Luigi Atzori, *Senior Member, IEEE*, and Michele Nitti, *Senior Member, IEEE*

***Abstract*—The Internet of Things (IoT) is increasingly supporting the tourism sector by enabling adaptive services that enhance the visitor experience. Among emerging applications, itinerary planning has gained significant attention, leveraging IoT data to deliver personalized and adaptive routes. However, current systems show key limitations. Personalization often depends on static forms rather than adaptive learning; validation is usually restricted to simulations, and the few real implementations mostly rely on mobile apps, tools that tourists are reluctant to download and often abandon after limited use. To overcome these challenges, this article proposes an itinerary planning system that integrates reinforcement learning for tourist profiling with a genetic algorithm (GA) for multiobjective optimization, implemented in a real-world scenario through a cloud infrastructure and delivered via an interactive totem that serves as the access point for tourists. Results confirm both the efficiency and scalability of the approach, showing that the system can be seamlessly extended to diverse urban contexts.**

***Index Terms*—Internet of Things (IoT), itinerary planning, real-world implementation, user profile.**

I. Introduction

The Internet of Things (IoT) has become a widely adopted paradigm, embedded in our daily lives and transforming the way individuals, businesses, and public administrations interact with their environments [1]. The integration of sensors, actuators, and cloud infrastructures has turned the IoT into a pillar of digital transformation, enhancing experiences in several domains and facilitating decisions. The expansion of IoT has been particularly visible across a wide set of application areas, where industry and health represent only a few examples [2].

Among these domains, tourism has emerged as another important sector where IoT has played an important role [3]. Tourism, one of the most dynamic sectors of the global economy, involves the continuous interaction between visitors, services, and the urban environment. By leveraging IoT systems, cities can collect data, for example, real-time mobility information or environmental conditions, and combine it with citizen content to improve the quality of the tourist experience [4]. These types of solutions have already been used to monitor crowd flows at popular attractions, optimize transport services, and provide information through guides or mobile platforms [5].

Recently, increasing attention has been devoted to intelligent itinerary planning systems, which leverage IoT data and advanced algorithms to deliver personalized travel experiences. Itinerary planning represents a notable application of IoT in tourism, where traditional travel planning approaches, often limited to guidebooks or static websites, fail not only to capture the dynamic nature of urban areas and their factors, such as weather or traffic, but also to adapt to the specific characteristics and preferences of individual users. In this sense, IoT systems facilitate the generation of adaptive itineraries, continuously updating suggestions based on real-time data [6]. The objective is not merely to recommend Points of Interest (PoIs), but to generate optimized routes that respect temporal constraints, reduce waiting times, and consider available modes of transportation [7].

Despite these advances, the state of the art still presents several deficiencies. First, personalization is often limited, relying mainly on static user inputs or explicit questionnaires, rather than on adaptive mechanisms that evolve through interactions. Second, many existing solutions have been validated only in simulated environments or with respect to datasets, while real implementations remain scarce. Finally, the reliance on mobile applications requiring installation poses barriers, since many tourists are reluctant to download additional software during their trips, preferring immediate interaction modalities [8].

To address the limitations identified in existing approaches, this article introduces an itinerary planning system designed to enhance the tourist experience in smart cities. The primary contribution of this work is the design of a cloud-based system that integrates IoT technologies to offer personalized and adaptive itineraries. This design is complemented by the adoption of reinforcement learning for tourist profiling and a genetic algorithm (GA) for multiobjective optimization. The system is implemented in a real-world scenario via a cloud infrastructure and an interactive totem, which serves as the point of interaction for tourists. The main contributions of this article are as follows.

Received 3 February 2026; revised 15 April 2026; accepted 18 May 2026. Date of publication 25 May 2026; date of current version 27 July 2026. This work was supported in part by Italian Ministry of Enterprises and Made in Italy (MIMIT) within the 5G Technology Support Program, on axis 1 "House of Emerging Technologies" (CTE), Project Name "Cagliari Digital Lab" under Grant G27F22000040008. *(Corresponding author: Claudio Marche.)*

Claudio Marche, Luigi Atzori, and Michele Nitti are with the Department of Electrical and Electronic Engineering (DIEE), University of Cagliari, 09123 Cagliari, Italy, and also with the National Telecommunication Inter-University Consortium, Research Unit of Cagliari, 09123 Cagliari, Italy (e-mail: claudio.marche@unica.it; l.atzori@unica.it; michele.nitti@unica.it).

Vlad Popescu is with Transilvania University of Braşov, 500036 Braşov, Romania (e-mail: vlad.popescu@unitbv.ro).

Digital Object Identifier 10.1109/JIOT.2026.3696714

1) Design of a context-aware itinerary planning system that collects and integrates data from various IoT sources across a smart city. These data, including real-time mobility, environmental conditions, and PoIs, are processed on cloud infrastructure to generate personalized itineraries.
2) A reinforcement learning model is employed to create and adapt user profiles, starting from demographic attributes and progressively refining using feedback from tourists' interactions. Additionally, a GA is used for multiobjective optimization, balancing factors such as category preferences, travel time, waiting time, and temporal constraints to generate optimized itineraries.
3) Unlike many existing systems that are limited to simulations, the proposed system is implemented in a real-world smart tourism scenario. It leverages IoT infrastructures and open municipal data for continuous updates, and it is accessible via an interactive totem.

The remainder of this article is organized as follows. Section II reviews the state of the art in IoT-based itinerary planning. Section III introduces the reference scenario considered in this work, while Section IV presents the design of the proposed system, including its main modules. Section V describes the simulation setup and discusses the obtained results. Finally, Section VI concludes this article and outlines directions for future research.

II. IoT Architectures for Tourism Applications

In recent years, the diffusion of smart devices has enabled a new class of tourism services where itinerary planning and recommendation systems are connected with the IoT. These systems collect and process heterogeneous data, from fixed and mobile sensors to social media and open municipal datasets, to generate travel suggestions and assist the tourist's journey [9]. Unlike traditional recommender systems based solely on static information, these IoT tourism applications can integrate dynamic variables, mainly represented by geolocation, weather conditions, and transport status, just to cite a few [10]. These data are continuously transmitted, aggregated on cloud platforms, and processed through advanced filtering and optimization techniques. These mechanisms ensure that recommendations remain current and context-aware, dynamically adapting itineraries to evolving urban conditions [11]. Moreover, beyond real-time adaptation, the use of cloud infrastructures also enables the integration of additional information, such as user content and social interactions, which can provide enhanced personalization [12]. In the following, an overview of recent works and references concerning recommendation systems in the tourism domain is provided, progressing to the most recent. The objective is not to produce a complete survey, but rather to outline the main approaches and trends that have emerged in this area. This targeted review focuses on identifying the characteristics of existing solutions and highlighting their most significant limitations.

Among the recent works proposed in the state of the art, two important proposals are illustrated in [13] and [14]. In the first work, the authors propose a travel route recommendation system that collects tourists' data via smartphones and IoT devices (BLE beacons, accelerometer, and camera). Behavior sequences, including visit duration and photos taken, are processed into pattern sequences and analyzed with an algorithm to discover frequent routes. The system then recommends itineraries tailored to user profiles and route constraints, tested in an exhibition hall. While in the second work, the proposed system gathers PoI information from online sources, clusters PoIs via a k-means algorithm, and optimizes the itinerary using a GA with a parameterized fitness function accounting for PoI constraints, travel times, and user preferences.

Furthermore, two other approaches based on machine-learning techniques are illustrated in [15] and [16]. The first proposal introduces a deep learning tourist recommendation system that integrates information such as location, time, and season with historical user preferences. It leverages a convolutional neural network (CNN) to extract features from social media images and combines them with metadata for PoI ranking. Recommendations are updated as new user data becomes available. The second study designs a tourism platform that integrates service composition methods to support itinerary planning. It applies a passenger flow prediction model based on regression and an itinerary recommendation model using the ant colony algorithm, tailored to time and budget constraints. The platform also incorporates RFID monitoring for real-time management of scenic areas. Moreover, in the same year, another important approach was illustrated in [17], in which the authors introduced an itinerary recommendation engine for tourism that leverages Industrial IoT data and geo-tagged photos. It integrates user interests, POI popularity, weather conditions, and travel expenses, and applies a multiobjective optimization algorithm to generate itineraries. Experiments are conducted on photo datasets, showing improvements compared to benchmark methods.

In addition, two novel works that mainly aim to adapt the systems to tourists' preferences are depicted in [18] and [19]. The first article presents a tourism recommendation system that integrates user search queries, browsing history, and data (such as location and weather) with a fuzzy clustering approach to group similar preferences. Recommendations are generated through a deep learning model. The second article proposes a system that generates tour itineraries by combining user constraints (duration, destination, and accommodation level) and preferences (PoI categories and activities) with a recommended itinerary database. It employs a filtering process and allows user modifications before applying a branch-and-bound algorithm for route planning. Furthermore, an important algorithm is proposed in [20], in which the authors present an itinerary planning framework aimed at reducing congestion and improving tourist experience. The system uses Raspberry Pi devices connected to footfall and vehicle count sensors to monitor real-time crowd and traffic conditions at popular tourist spots. This data is transmitted to a central server and displayed to visitors, enabling them to choose visiting sequences that minimize waiting times and travel delays. Finally, two recent works, which explicitly exploit reinforcement learning to adapt itinerary recommendations based on user feedback, are illustrated in [21] and [22]. In the first article, the authors formulate the itinerary planning

TABLE I

COMPARISON OF RECENT IOT ITINERARY PLANNING APPROACHES

Ref.	Inputs	Algorithm	Personalization	Implementation
[13]	Smartphone data and visit duration	Based on PrefixSpan	User profile form and onsite behavior	Prototype tested in exhibition hall
[14]	POI data and user survey	k-means and Genetic Algorithm	User survey	Mobile application
[15]	User form, weather data and social media images	Deep Neural Network	User profile form	Tested on datasets
[16]	PoI data and flow prediction	Regression and Ant Colony Optimization	User form constraints (time, budget)	Prototype simulation
[17]	Sensor data, photos and expenses	Multi-objective optimization	User interests from photos	Simulated with datasets
[18]	Search queries and weather	Collaborative Filtering and Apriori	Browsing history	Tested on datasets
[19]	User inputs and itinerary database	Filtering and Branch-and-Bound	User form	Simulated
[20]	Footfall and traffic counts	Threshold rules	No personalization	Raspberry Pi prototype
[21]	User preferences and PoI data	Q-learning	Reward-based adaptation from user feedback	Tested on datasets
[22]	Historical visits and contextual data	Deep Reinforcement Learning	Implicit feedback from past itineraries	Tested on datasets

problem as a Markov process and propose a Q-learning-based approach to generate multiday tourist routes. User preferences are incorporated through a reward function that accounts for factors such as travel time, rating, and cost, enabling the system to adapt the recommended routes to the user's desired criteria. Similarly, in the second article, the authors propose a deep reinforcement learning system for generating tourist itineraries, with a particular focus on overcrowding mitigation. The system leverages information and historical visit data to learn policies that guide tourists through sequences of PoIs, balancing user experience with destination constraints. The learning process is driven by implicit feedback derived from observed tourist behavior and past itineraries. However, despite their focus on personalization, both works exhibit a limited integration with the real-time information of the urban environment. In particular, city information is either derived from historical datasets or treated as static input, without being connected to IoT infrastructures capable of providing continuous updates. As a result, the adaptation mechanisms primarily address user preferences, while the ability to react to evolving urban factors remains limited.

The analyzed works are summarized in Table I, which reports for each reference the types of inputs considered, the algorithmic approach adopted for itinerary generation, the personalization strategy, and the kind of implementation provided. As summarized in the table, IoT technologies have an important role in the tourism domain, supporting a variety of approaches through the integration of real-time data. Despite promising results, these approaches still show significant deficiencies, leaving ample room for further research.

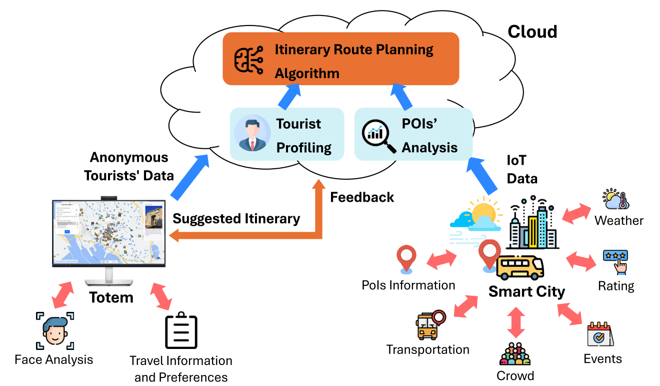

Fig. 1. Proposed system architecture.

First, in many studies, the personalization of tourism services based on the individual tourist is almost absent. When personalization is considered, it is typically limited to explicit information collected through forms filled out by the tourist or, alternatively, through more invasive mechanisms such as the analysis of personal profiles. Another limitation is the lack of real implementations, since the majority of models are tested through simulations, while public deployments are missing. Furthermore, the type of applications offered often relies on the tourist actively installing mobile apps or registering on dedicated platforms that require access to personal data. This is problematic, since tourists are increasingly reluctant to download additional applications on their devices, preferring instead systems that minimize installation [8].

To address these limitations, this article proposes an itinerary planning system that combines user profiling with optimization techniques and that has been implemented in a real-world setting. Unlike previous approaches, personalization is achieved through reinforcement learning, which progressively adapts the tourist profile based on his/her characteristics and interaction feedback. Rather than relying on mobile applications that require installation, the system is delivered through an accessible interactive totem, easily extendable to multiple locations within the same city or to different cities.

III. SCENARIO

The primary goal of this work is not only to enhance smart city tourism by generating itineraries that are not only personalized but also adaptive to contextual factors such as mobility, events, and environmental conditions. To this end, we propose and evaluate in a real-world setting a context-aware itinerary planning system that leverages reinforcement learning for user profiling and a GA for multiobjective optimization. From the tourist's perspective, not only does the system recommend the optimal POIs to be visited for the user, but also their number and the type. It further improves the overall experience, accounting for waiting times, temporal constraints, and preferred modes of transportation.

As illustrated in Fig. 1, the proposed system architecture is structured around three main entities: tourist, IoT, and cloud.

The tourist component is realized through an interactive touchscreen totem, which represents the interface between tourists and the system. The totem collects explicit user inputs via a graphical interface (web form), including travel information, such as departure date, number of days, and preferred transportation mode. In addition, the totem integrates a face analysis module, executed locally on the device, which estimates the tourist's age range and cultural geographical origin using a pretrained deep learning model [23]. This module operates exclusively on-device for privacy reasons: no facial images are stored or transmitted, and only anonymized demographic attributes are forwarded to the cloud. These attributes are exploited as an initial approximation of user preferences, leveraging empirical evidence showing that tourists sharing similar demographic characteristics often exhibit comparable interests [24], [25]. In addition, the totem represents the point where tourists' feedback is collected. By allowing tourists to accept or modify the proposed itinerary, the totem enables a continuous feedback mechanism in communication with the cloud.

The IoT component provides information about the city's state. The collected data include both static and dynamic information. Static data mainly refers to PoI-related attributes, such as category, geographic location, opening hours, average visit duration, and accessibility features. Dynamic data capture time-varying aspects of the urban environment, including crowd levels at PoIs, real-time transportation conditions, weather updates, user ratings, and scheduled events. These data are obtained from a combination of IoT infrastructures and open municipal data services, and are continuously updated to reflect real-time urban conditions. Whenever available, public municipal open data and city APIs are prioritized as primary sources, rather than relying on proprietary platforms. To avoid overloading these services, data acquisition is performed through a limited number of queries: less variable PoI attributes (e.g., opening hours, accessibility, and general descriptions) are cached in the cloud and refreshed periodically, whereas dynamic information (e.g., mobility and crowd data) is queried more frequently. When municipal data are unavailable for specific fields, queries to online services (e.g., travel-time queries via Google map services) are supported. In general, the IoT component allows the system to adapt itinerary generation not only to individual tourist preferences but also to evolving environmental and city factors. These heterogeneous data directly affect itinerary feasibility by constraining visit schedules, modifying travel times, and updating PoI relevance. By continuously integrating IoT information, the system maintains an up-to-date representation of the city, enabling itinerary generation to adapt to both user preferences and evolving urban conditions.

The cloud infrastructure hosts the system logic and is responsible for integrating tourist data with IoT information. Within the cloud, tourist profiling, PoIs analysis, and route planning are performed as distinct but interconnected processes, as will be explained in Section IV. The collected data are aggregated and processed to generate a personalized itinerary that accounts for several constraints, such as user interests, city dynamics, and temporal limits. To speed up itinerary generation and to avoid repeated calls to the public API, while candidate solutions are explored, a cache of the city's state is maintained for PoI attributes and dynamic context and is consumed by the planning module. The cache is refreshed by background jobs scheduled according to data variability (periodic refresh for static attributes and more frequent updates for dynamic data), so that optimization is executed on available data. The resulting itinerary is sent back to the totem and visualized through an interactive interface on the totem, allowing the tourist to accept or modify the proposed route. User interactions are interpreted as feedback and exploited to iteratively refine the tourist profile through a reinforcement learning mechanism, improving the accuracy of future recommendations. Optionally, the itinerary can also be delivered via email, without storing any personal information.

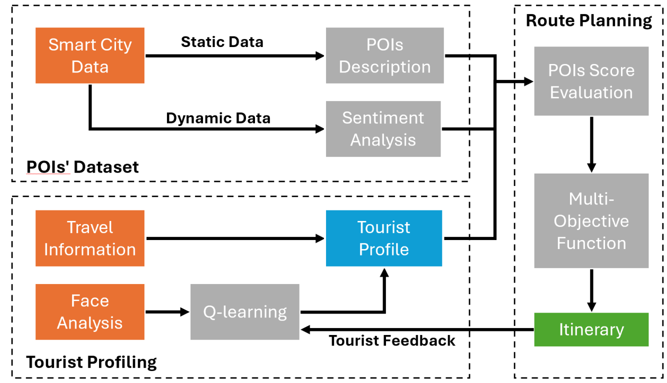

Fig. 2. Overview of the system workflow.

In the remainder of this article, we adopt the following notation. Each PoI is denoted by p_i, where i indexes the PoIs in the dataset; each PoI belongs to a category j, as detailed in Section IV-A, where all the stored and utilized information is described. A tourist profile for each user u is represented by the tuple $\langle a_u, o_u, \boldsymbol{\sigma}_u \rangle$, where a_u denotes the estimated age range, o_u the estimated cultural-geographical origin and $\boldsymbol{\sigma}_u = [\sigma_{u,j}]$ is a vector of interest weights. Each $\sigma_{u,j} \in [0, 1]$ represents the relative interest of user u in category j, with $\sum_j \sigma_{u,j} = 1$. The generated itinerary is defined as an ordered sequence of daily itineraries, $\mathcal{I} = [\mathcal{I}_1, \ldots, \mathcal{I}_d, \ldots, \mathcal{I}_D]$, where D is the number of days selected by the tourist and each $\mathcal{I}_d$ represents an ordered list of PoIs scheduled for day d.

IV. System Design

This section presents the logical architecture of the proposed system, which comprises three main modules: PoI dataset construction, tourist profiling, and route planning. Fig. 2 illustrates the overall workflow, from data collection and tourist characterization to the generation of a personalized travel itinerary.

All system components are hosted in the cloud, granting access to a wide range of smart city open data. The only exception is the face analysis module, which is executed directly on the totem for privacy reasons. This design ensures that personal data (facial images) are processed locally, while only anonymized outputs, estimated age range, and origin are transmitted to the cloud.

The three functional modules depicted in the figure are described in Sections IV-A–IV-C. The first module deals with the construction and scoring of the PoIs dataset; the second focuses on tourist profiling through reinforcement learning; and the third formulates itinerary generation as a multiobjective optimization problem, solved using a GA.

A. PoIs' Dataset

The first module is responsible for constructing and maintaining the dataset of PoIs available in the city. This is achieved through a hybrid data collection strategy that combines APIs from popular review platforms, such as Google Maps, with open data services provided by the local municipality. The retrieved information is categorized into two types.

1) *Static Data:* Attributes that remain largely unchanged over time, such as the PoI category (e.g., historical sites and monuments), geographic coordinates (GPS), average visit duration, accessibility details (e.g., wheelchair access), and opening hours.
2) *Dynamic Data:* Continuously updated information, including ratings and user comments. These are frequently retrieved to ensure the dataset remains aligned with current conditions.

Each PoI is described by a feature vector comprising category, opening hours, visit duration, location, rating, and accessibility indicators. To complement the quantitative data with qualitative insights, a sentiment analysis step is applied to user reviews. The system employs the widely adopted VADER model [26], which classifies textual feedback as positive, negative, or neutral. The outcome of this analysis contributes to a dynamic scoring mechanism of each PoI.

The dataset is updated regularly and stored in a format suitable for the optimization tasks performed during the itinerary planning phase. All the information is obtained by integrating municipal open data with online platforms frequently consulted by tourists, such as Google Maps and TripAdvisor.

B. Tourist Profiling

The second module is responsible for generating a personalized profile for each tourist. As described in Section III, the profiling process begins by estimating the user's demographic attributes, that is, the age a_u and origin o_u. When the tourist provides consent, a facial image is processed locally on the totem device using the DeepFace library [23], an open-source deep learning framework based on CNNs. DeepFace leverages state-of-the-art models such as Google FaceNet and Facebook DeepFace, trained on large datasets. The predicted age range and cultural background are computed locally and transmitted to the cloud as anonymized inputs for the reinforcement learning model.

In our formulation, tourist profiling is modeled with Q-learning, enabling profiles to evolve iteratively based on user interactions and feedback. The main components are defined as follows. The state S denotes the tourist profile, represented by the tuple $\langle a_u, o_u, \boldsymbol{\sigma}_u \rangle$, where $\boldsymbol{\sigma}_u$ is the distribution of interests across the PoIs, represented as a probability vector over the categories. The action A corresponds to a distribution of categories proposed by the system, which is then used to generate an itinerary by the route planning module (described in Section IV-C). The feedback F represents the updated distribution of categories derived from the tourists' feedback, which is provided by the user accepting or modifying the proposed itinerary. The reward R measures the similarity between A and F, quantifying how well the system's proposal matches the tourist's actual preferences. Finally, the Q value $Q(S, A)$ represents the expected utility of taking action A in state S, balancing the immediate reward with the value estimated for the next state.

The process begins by checking whether a state S already exists for the demographic attributes (a_u, o_u). If so, the corresponding state with the maximum Q value is retrieved; otherwise, S is initialized with a uniform interest distribution across categories, and in this specific case, the first action A is selected randomly. Once the initial state has been set, an action A is selected according to an ε-greedy policy.

1) *With Probability ε:* The system explores by selecting a random action.
2) *With Probability $1 - \varepsilon$:* The system exploits current knowledge by choosing the action that maximizes $Q(S, A)$.

Where for each tourist group defined by demographic attributes, the exploration rate ε is gradually decreased over time, following a power-law decay $\varepsilon_k = (1 + k)^{-\eta}$, with k denoting the generic interaction within that group and $\eta > 0$ the decay exponent controlling the rate.

Based on the selected action A, the system generates a distribution of PoIs across categories and produces a tourist itinerary. This itinerary is returned to the totem and presented to the tourist through the interactive interface.

At this point, the tourist provides feedback F by either accepting the proposed itinerary or modifying it according to his or her preferences. This feedback directly translates into a new distribution of categories, which serves as the basis for updating the profile. The reward R is then computed as the similarity between the proposed action A and the feedback F, both represented as probability vectors over the categories. Similarity is measured through the Jensen–Shannon distance $d(A, F)$, and the reward is defined as

$$R = 1 - d(A, F), \quad \text{with:}$$
$$d(A, F) = \sqrt{\frac{1}{2} \sum_{j}^{n} \left(A_j \log_2 \frac{A_j}{A_j + F_j} + F_j \log_2 \frac{F_j}{A_j + F_j} \right)} \tag{1}$$

where n is the total number of categories.

To retain temporal consistency across interactions, the new state S_{new} is computed by blending the previous state S and the updated distribution F through exponential smoothing

$$S_{\text{new}}(j) = \lambda \cdot S(j) + (1 - \lambda) \cdot F(j) \tag{2}$$

where $\lambda \in [0, 1]$ is a decay parameter that regulates the influence of past information. The updated state S_{new} is then adopted as the new state of the system.

Q-learning then proceeds by updating the Q-values associated with each state–action pair. The Q value $Q(S, A)$ denotes

Algorithm 1 Tourist Profiling

```
Estimate age a_u and origin o_u using DeepFace
if state S for (a_u, o_u) exists then
    Load state S with max max_A Q(S, A)
else
    Initialize S with uniform interest distribution
end if
Select action A using ε-greedy policy
Generate PoI distribution based on A
Present itinerary to user and collect feedback → F
Compute reward R
Update state: S_new
Update Q value
Set S ← S_new
Reduce exploration rate ε
```

the expected utility of taking action A in state S, considering both the immediate reward and future rewards. The update rule is given by

$$Q(S,A) \leftarrow Q(S,A) + \phi\left[R + \gamma \cdot \max_{A'} Q\left(S_{\text{new}}, A'\right) - Q(S,A)\right] \tag{3}$$

where ϕ is the learning rate, controlling the weight assigned to new information, and γ is the discount factor, balancing immediate versus long-term rewards. The term $\max_{A'} Q(S_{\text{new}}, A')$ represents the estimated best-possible reward achievable from the updated state S_{new}.

The full sequence of steps involved in the tourist profiling process is summarized in Algorithm 1.

C. Route Planning

The third module is responsible for generating the tourist itinerary, formulated as a multiobjective optimization problem. A GA is employed due to its ability to handle multiple, often conflicting objectives simultaneously [27].

Classical optimization techniques, such as linear and nonlinear programming, are often unsuitable for complex multiobjective scenarios involving heterogeneous constraints. In contrast, GAs are global search heuristics inspired by biological evolution [28]. They evolve a population of candidate solutions (chromosomes) over successive generations, applying genetic operators, such as selection, crossover, and mutation, to progressively improve the population's fitness. This makes them particularly well-suited to itinerary planning under temporal, spatial, and preference-based constraints.

In this context, each individual in the population represents a complete itinerary for a tourist, distributed across a fixed number of days. Each chromosome encodes the assignment of PoIs to daily slots, including the order and time of visits. The initial population is generated by selecting PoIs based on the tourist profile and associated priority scores.

Each PoI i is assigned a score P_i, computed as

$$P_i = \alpha V_i + \beta C_i + \gamma T_i, \quad \text{with } \alpha + \beta + \gamma = 1 \tag{4}$$

where V_i is the normalized rating of PoI i, C_i is the ratio of positive user comments obtained via sentiment analysis, and T_i is a binary indicator of accessibility. All components are normalized in $[0, 1]$. These scores are used both for prioritizing PoIs and as input to the GA.

Within each category, PoIs are ranked by score P_i and the top entries are retained. The GA then operates on this filtered subset. Input parameters include the following.

1) *Metadata for Each PoI:* Category, coordinates, visit duration, and opening hours.
2) The priority scores P_i.
3) Number of days in the itinerary.
4) Tourist interest distribution across PoIs categories.
5) Optional temporal constraints such as lunch breaks (e.g., 12:30–14:30).

The initial population is generated by randomly sampling feasible combinations of PoIs for each day. Individuals are evolved using the following genetic operators: *Selection*, which identifies promising solutions based on their fitness scores; *Crossover*, which combines two different itineraries by exchanging portions of their daily schedules while preserving feasibility; *Mutation*, which introduces diversity by reassigning PoIs to different days or altering their order; and *Repair*, which ensures feasibility by adjusting visits that violate opening hours or by rebalancing the category distribution to better align with the user profile.

Each solution is evaluated through a fitness function F, which quantifies the overall quality of the proposed itinerary. In the context of this problem, the fitness function is designed to balance multiple goals. Higher fitness values correspond to itineraries that better satisfy user preferences while remaining feasible within the constraints of the visit. The evolutionary process is repeated for a fixed number of generations or until convergence. At each generation, individuals are evaluated via the fitness function.

The fitness function is defined as a weighted combination of four distinct objectives, each reflecting an important aspect of the itinerary.

1) Maximize PoI Relevance: The itinerary should prioritize highly rated and positively reviewed PoIs

$$f_1 = \sum_{i \in \mathcal{I}} P_i \tag{5}$$

where P_i is the score assigned to PoI i according to (4).

2) Minimize Travel Distance: The total distance between consecutive PoIs should be reduced

$$f_2 = \sum_{d=1}^{D} \sum_{i=1}^{n_d - 1} dist\left(p_i^d, p_{i+1}^d\right) \tag{6}$$

where D is the number of days, n_d the number of PoIs scheduled on day d, and p_i^d the ith PoI scheduled on day d. Distances $dist(p_i^d, p_{i+1}^d)$ between PoIs are retrieved from the local open data and Google Maps API, based on the selected transportation mode by the tourist, that is, walking, public transport, taxi, or a combination of them.

a) Minimization of time penalty for visiting PoIs outside their opening times: To ensure the itinerary is executable, each PoI should be visited during its valid time window. Let v_i be the scheduled visit time for PoI i, and $[O_i, C_i]$ its opening interval. The temporal penalty is defined as

$$f_3 = \sum_{i \in \mathcal{I}} \delta_i, \quad \text{where} \quad \delta_i = \begin{cases} 0, & \text{if } O_i \le v_i \le C_i \\ 1, & \text{otherwise.} \end{cases} \tag{7}$$

This term directly enforces the temporal constraints of the itinerary by ensuring that each visit is scheduled within the PoI opening-time window; jointly with the average visit durations (available from city data), it enables a scheduling of all visits while respecting the overall daily timing constraints.

3) Match User Profile Preferences: The itinerary's category distribution should align with the tourist profile. Let z_j be the proportion of PoIs of category j in the itinerary, and σ_j the target interest weight. A tolerance τ accounts for small deviations

$$f_4 = \sum_{j=1}^{K} \max\left(0, |z_j - \sigma_j| - \tau\right) \tag{8}$$

where K is the number of PoI categories.

The overall fitness function F is a weighted combination of the four objectives

$$F = w_1 f_1 - w_2 f_2 - w_3 f_3 - w_4 f_4 \tag{9}$$

where w_1, w_2, w_3, w_4 represent scalar weights used to control the tradeoff between competing objectives, used to prioritize the different aspects.

V. Experimental Evaluation

In this section, we first describe the considered scenario for performance evaluation, and then present the obtained results.

A. System Setup

To validate the proposed system in a real-world scenario, a working prototype has been developed. The architecture includes two core components: the totem and the cloud server, both containerized using Docker and implemented with Django as the web framework.

The totem, which should be installed at the cruise terminal (the main tourist interaction PoI in Cagliari), is an all-in-one touchscreen device equipped with a quad-core processor at 2.40 GHz, 16 GB DDR4 RAM, and 512 GB SSD storage. This machine runs the local GUI and the face analysis module, which relies on the DeepFace library to extract facial features. All biometric processing occurs locally on the totem; no facial images are stored or transmitted. Only nonidentifying attributes, that is, the estimated demographic attributes, are sent to the cloud, not linked to any identifier or contact information, and used solely to parameterize the itinerary generation. Geospatial data used for visualization is retrieved in real time via the Google Maps Platform, using the Directions API.

The cloud backend, deployed on Microsoft Azure, hosts the itinerary planning services, including the Q-learning and genetic optimization modules. These services are executed within a container instance provisioned with 4 vCPUs and 8 GiB of RAM. The backend receives inputs from the totem, performs tourist profiling, and generates the personalized itinerary accordingly.

All components communicate via RESTful APIs. Tourist sessions are stored in a central database to enable Q-learning updates based on implicit user feedback. An example of the totem interface displaying a personalized itinerary generated is shown in Fig. 3.

Fig. 3. Totem interface displaying an example of personalized itinerary.

TABLE II
Overview of PoIs per Category

Category	# PoIs	Avg Rating	Avg. Visit Time (min)	Accessibility (%)
Beaches	4	0.87	95	25.0
Churches and Places of Worship	28	0.86	30	78.6
Historical Sites and Monuments	34	0.85	35	70.6
Markets and Shopping Streets	13	0.88	25	100.0
Museums and Art Galleries	14	0.85	50	100.0
Parks and Gardens	8	0.86	45	100.0
Scenic Spots	10	0.92	30	50.0

B. Analysis of Performance

The performance assessment begins by analyzing the characteristics of the PoI dataset used for itinerary generation. The dataset comprises 111 PoIs, grouped into seven categories that reflect common urban tourism experiences in the city. Table II provides an overview of the dataset, reporting for each category: the number of PoIs, the average user rating V_i (normalized in the [0, 1] range), the average expected visit time (in minutes), and the percentage of accessible locations. The dataset is constructed by combining open data made available by the municipality of Cagliari with information retrieved from the most used tourist platforms, such as Google Maps. Overall, the dataset presents a balanced category distribution, with *historical sites and monuments* and *churches and places of worship* being the most represented. *Scenic spots* exhibit the highest average rating (0.92), while *beaches* show the longest expected visit time (95 min). Accessibility levels vary significantly across categories: while *museums and art galleries*, *parks and gardens*, and *markets* are fully accessible, only 25% of the *beaches* meet accessibility criteria. To support deployment, PoI information is persisted in the cloud (to avoid repeated calls to the smart city API) and periodically refreshed to capture updates such as changes in availability or opening hours. Each entry includes the PoI identifier and

TABLE III
IMPACT OF Q-LEARNING HYPERPARAMETERS ON THE FINAL REWARD

α	γ	λ	**Reward**	α	γ	λ	**Reward**
0.1	0.1	0.1	0.8153	0.5	0.9	0.1	0.8140
0.1	0.1	0.9	0.8473	0.5	0.9	0.9	0.8499
0.1	0.9	0.1	0.8121	0.9	0.1	0.1	0.8328
0.1	0.9	0.9	0.8479	0.9	0.1	0.9	0.8640
0.5	0.1	0.1	0.8246	0.9	0.9	0.1	0.8108
0.5	0.1	0.9	0.8574	0.9	0.9	0.9	0.8430

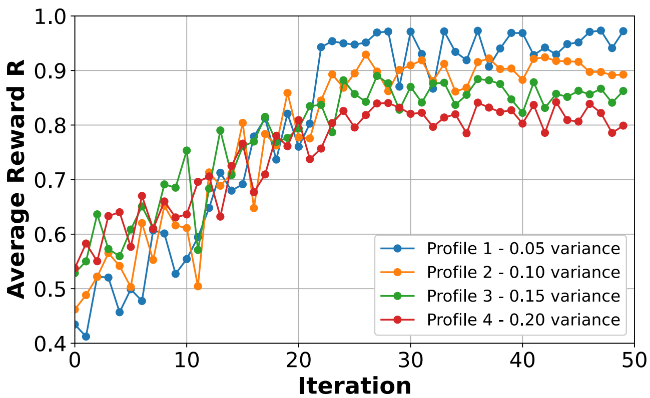

Fig. 4. Evolution of the average reward during Q-learning updates under different levels of noise.

name, category, GPS coordinates, opening-time window, average visit duration, accessibility flags, and popularity indicators (e.g., rating/review aggregates). As an example, the record for "Bastione Saint Remy" contains some of the most relevant fields in the dataset, including the PoI name, the category (*historical sites and monuments*), the average rating (4.6), the number of ratings (27 475), and the opening-time information, stored as a weekly schedule (e.g., open 24 h for each day). These attributes are used by the scoring and scheduling components during itinerary construction.

The first set of simulations assesses the convergence behavior of the tourist profiling module. To address the impact of hyperparameter settings, we first conduct an analysis of the parameters of the Q-learning update, namely the learning rate ϕ, the discount factor γ, and the exponential smoothing coefficient λ. The analysis is performed on a reference user profile, where the ground-truth preference vector is set randomly, and the simulated feedback at each interaction is obtained by perturbing with additive Gaussian noise (standard deviation 0.10). For each hyperparameter configuration, the interaction process is repeated 30 times with different random seeds, and we report the reward at the last interaction (Table III). The reward is computed as the similarity between the proposed action and the feedback distribution, and used to update the profile state via the Q-learning update rule. The results show that λ has the strongest impact, consistently with its role in controlling the stability in the profile update, while α and γ have a secondary but measurable effect through the Q-learning update. Based on Table III, the best-performing hyperparameter configuration is selected and used in the subsequent study.

TABLE IV
PERFORMANCE COMPARISON AMONG ITINERARY PLANNING ALGORITHMS

	Fitness	**Distance [km]**	**Invalid PoIs**	**Time [s]**	**Profile Similarity [%]**
Ant Colony Optimization (ACO)	3.10	2.44	2.14	1.13	54.35
Simulated Annealing (SA)	-3.60	4.91	2.42	0.62	90.14
Tabu Search (TS)	-2.48	3.90	2.30	0.73	88.21
Iterated Local Search (ILS)	-3.13	3.77	2.41	0.60	90.91
Proposed Genetic Algorithm (GA)	2.67	5.99	1.08	0.82	97.20

After selecting the best hyperparameter configuration, we evaluate the convergence of the profiling module under heterogeneous user behaviors. To this end, four synthetic tourist profiles are defined, each expressing a strong preference for a single PoI category. For each profile, we simulate a sequence of 50 interactions. At each interaction, the system selects an action (a distribution over PoI categories) following the ϵ-greedy policy, while the tourist feedback is generated by perturbing the corresponding ground-truth preference distribution with Gaussian noise, representing imperfect feedback. Since the profile is represented as a normalized vector of category percentages, this perturbation emulates the variability in how a tourist expresses preferences across interactions while preserving a valid distribution over categories. In particular, we consider four noise variances (from low to high), which correspond to different degrees of heterogeneity: lower variance captures more close feedback across interactions, whereas higher variance emulates less predictable choices. Moreover, the exploration probability ϵ starts from 1 and decreases linearly to 0.05 during the first 25 interactions, remaining fixed afterward. Fig. 4 reports the evolution of the reward across the interactions for the four profiles. In all cases, the reward increases in the initial phase, when exploration is higher, and the profile is still adapting, and then stabilizes as the estimated preference distribution becomes consistent with the user behavior. The curves show a clear convergence trend after the first few tens of interactions, while maintaining a stable reward level in the remaining steps, indicating that the profiling module can adjust the preference distribution with limited interaction history. Small differences among the profiles can be observed in terms of convergence speed and final reward. The results confirm that the proposed profiling strategy is able to refine user preferences over time under heterogeneous behaviors, providing a stable distribution output that can be used by the itinerary optimization module.

The second set of simulations evaluates the effectiveness of the proposed GA against several alternative optimization approaches: ant colony optimization (ACO) [29], simulated annealing (SA) [30], Tabu search (TS) [31], and iterated local search (ILS) [32]. Since the proposed system accounts for multiple heterogeneous requirements, the comparison is performed against optimization algorithms that can be instantiated under the same objective structure, whereas baselines from the literature would require substantial redesign and additional training data to cover all components; anyway, the selected algorithms are widely adopted in state-of-the-art itinerary planning proposals and are commonly used as reference optimizers in this research area. As this work focuses on the system architecture rather than on novel optimization heuristics, learning-based methods are excluded due to their dependence on large labeled datasets and limited applicability in such optimization scenarios without any previous tourist interactions. To ensure a fair comparison, we simulate 50 independent scenarios, each corresponding to a 3-day tourist experience. Regarding the fitness aggregation, the weights were set to reflect the desired priority among objectives. In particular, we assigned the highest importance to PoI relevance and profile matching, together with feasibility (opening-time compliance), while keeping travel distance as a secondary objective. Concretely, the adopted setting follows a normalized weight split of approximately 28.6% for relevance, 14.3% for travel distance, 28.6% for infeasible visits (opening-time related penalty), and 28.6% for profile mismatch. For every simulation, a new tourist profile is generated using a Dirichlet distribution to ensure valid probability vectors representing interests across different PoI categories. All algorithms operate under identical conditions and are evaluated using the fitness function defined in Section IV. The GA evolves a population of candidate itineraries through standard operators (selection, crossover, and mutation), progressively refining solutions to meet multiple objectives. Selection retains the highest-fitness individuals at each generation. Crossover is implemented by recombining two parent itineraries through a single cut point (swapping the second part of the visit sequence), while mutation randomly swaps two PoIs within an itinerary with a given probability. Infeasible visits (e.g., PoIs falling outside opening hours) are naturally penalized by the fitness term, which drives the evolutionary search toward executable routes. The GA evolves a population of candidate itineraries through standard operators (selection, crossover, and mutation), progressively refining solutions to meet multiple objectives. ACO simulates the behavior of ants by constructing itineraries based on a combination of desirability and pheromone trails. The pheromone levels are updated depending on the quality of previous solutions, thus reinforcing the most promising directions in the search space. SA mimics the physical process of annealing by exploring neighboring solutions and occasionally accepting worse ones according to a decaying temperature function, which allows the algorithm to escape local optima in early phases while gradually stabilizing. On the other hand, TS relies on local search with a memory structure (the tabu list) that prevents cycling and allows exploration of new areas in the search space. Finally, ILS repeatedly applies local search with perturbations to diversify the search and overcome local optima. Table IV presents the results of the comparison, reporting the average performance across all simulation runs. The Table includes the fitness value, the average distance between PoIs (in kilometers), the number of invalid PoIs per itinerary, average execution time (in seconds), and the profile similarity expressed as a percentage. Results show that the proposed GA outperforms competing methods across most metrics. It achieves the second-best fitness score (2.67), indicating an effective tradeoff among the optimization objectives, and clearly obtains the highest profile similarity (97.20%), confirming its ability to generate itineraries well aligned with user interests. Although the total distance covered (5.99 km) is slightly higher than that of ACO, the GA compensates with a significantly lower number of invalid PoIs (1.08). In contrast, ACO shows good performance in terms of minimizing distance but fails to preserve user alignment, with profile similarity dropping to 54.35%. SA, TS, and ILS exhibit moderate to low fitness values, with higher rates of invalid PoIs and lower alignment with user preferences.

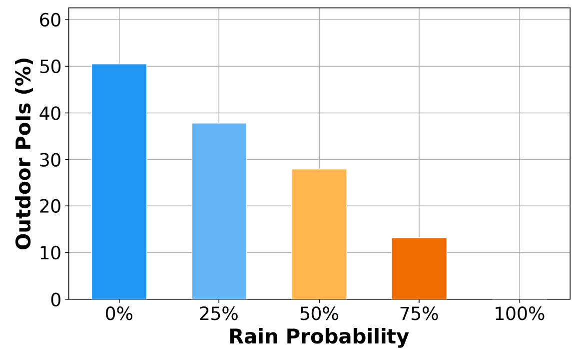

Fig. 5. Impact of weather on the percentage of outdoor PoIs selected in the itinerary.

To evaluate the impact of IoT inputs on the generated itineraries, a simulation is conducted to assess how real-time weather conditions affect the composition of the routes. In the proposal, weather data retrieved from IoT sources influences the accessibility parameter T_i of each PoI, which contributes to the priority score P_i defined in (4). Specifically, outdoor PoIs, belonging to categories such as beaches, parks and gardens, scenic spots, and several historical sites and monuments, are treated as inaccessible under rainy conditions, since their visitability is dependent on weather. To illustrate this behavior, five weather scenarios are simulated, corresponding to rain probability values of 0%, 25%, 50%, 75%, and 100%. For each scenario, 50 independent runs are executed, each with a generated tourist profile, and the percentage of outdoor PoIs selected in the final itinerary is measured. As shown in Fig. 5, the proportion of outdoor PoIs included in the itinerary decreases as the rain probability increases, dropping from 50.5% under clear conditions to 0% when rain is certain. The same adaptive behavior would arise in the presence of crowd data: highly congested PoIs would be flagged as inaccessible through the same accessibility mechanism, causing the GA to exclude them in favor of less crowded alternatives, thus improving the tourist experience.

Finally, to evaluate the scalability and the system's ability to handle multiple totems, a set of simulations using Azure is illustrated. Two distinct container configurations are tested: the first configuration is a container with 1 vCPU and 2 GB of RAM, and the second configuration is 2 vCPUs and 5 GB of RAM. For both configurations, four different setups are considered based on the number of replicas, in

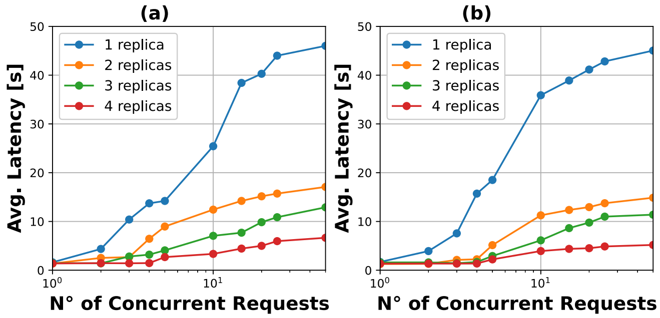

Fig. 6. Scalability analysis of the system with varying cloud configurations and replicas. (a) Performance for 1 vCPU, 2 GB RAM configuration. (b) Performance for 2 vCPUs, 5 GB RAM configuration.

TABLE V
CROSS-DATASET EVALUATION

Dataset	# POIs	# Categories	f_1	f_2	f_3	f_4	Distance [km]	Profile Similarity [%]
[33]	477	13	✓	✓	✗	✓	4.95	89.53
[34]	2341	13	✓	✓	✗	✓	6.64	86.23
[35]	877	10	✗	✓	✗	✓	4.12	95.71
[36]	20563	13	✗	✓	✗	✓	4.85	92.62

which replicas refer to multiple instances of the same container running in parallel. Each replica runs independently and handles incoming requests, which enables the system to distribute the load. As shown in Fig. 6, the left graph presents the performance for the first configuration, while the right graph shows the performance for the second one. In both cases, the average latency of the system is measured for different numbers of simultaneous requests. It can be observed that increasing the number of replicas significantly improves the system's ability to handle requests, as the latency for multiple concurrent requests decreases. However, in terms of request latency, the performance gain of the second configuration is limited, since the algorithm itself is already efficient, as shown in the previous simulations. Instead, the important factor influencing performance in this scenario is the number of replicas, as the system distributes the incoming requests across available instances. In general, the results demonstrate that the system is scalable. With four replicas in both configurations, the system can maintain average latency below 7 s for the first configuration and 6 s for the second one, even with 50 simultaneous requests and only four replicas.

C. Cross-Dataset Analysis

At last, this section investigates the portability of the system across publicly available PoIs' datasets. The goal is to assess whether the same optimization structure can be instantiated beyond the original dataset setting. To the best of our knowledge, none of the datasets of the state of the art is fully complete with respect to all the information required by the proposed formulation; therefore, the analysis in Table V focuses on the most well-known and complete datasets, and illustrates which optimization components can be tested, which ones can only be approximated, and which ones must not be considered.

Four city-scale datasets, namely Paris [33], London [34], Barcelona [35], and New York [36], have been considered. In all cases, geographic coordinates are provided for each POI; therefore, the distance-related component, that is, f_2, can be evaluated across datasets. Conversely, none of these sources provides opening-hour information in a form that can be exploited to determine whether a PoI is visitable within a given time slot; as a result, the visit component f_3 has not been assessed and PoIs have been treated as always visitable. With respect to the relevance component, Paris and London include review information that has been used to derive an approximate feature that is then normalized and combined with the fraction of positive versus negative comments. In contrast, Barcelona and New York do not provide reliable review or rating fields that can be mapped to the same notion of relevance; hence, the relevance component f_1 has not been considered for these two datasets. Finally, the profile component has been instantiated in all datasets, but with different levels of fidelity depending on the available categories. In Paris, categories are available and profile matching in f_4 has been performed directly. In the other datasets, profile matching has necessarily been approximated using the closest available categorical attribute: the annotated and not descriptive labels aligned to PoIs for London, the district categorization for Barcelona, and the facility type codes for New York.

Since these datasets contain a large number of PoIs, the generation is extended to 5 days to include more visits per itinerary and obtain more distinguished category distributions. The results in Table V show that route distances remain controlled even in larger cities, indicating that the optimizer tends to group visits into geographically daily patterns. At the same time, the profile similarity remains high across datasets, suggesting that the proposed optimization is robust in preserving the intended preference distribution even when relevance information is missing.

VI. CONCLUSION

This article has presented an itinerary planning system designed to enhance the tourist experience in smart cities by integrating user profiling with multiobjective optimization. The proposed solution combines reinforcement learning to adapt tourist profiles with a GA to generate optimized itineraries that account for preferences and various constraints, from the tourist perspective and of the smart city. The experimental evaluation has confirmed the efficiency of the approach. The profiling module, based on Q-learning, successfully converges toward tourists' interests, ensuring that itineraries improve over successive interactions. At the same time, the GA has demonstrated the best performance when compared with alternative methods, producing itineraries that achieve high similarity with user profiles while maintaining feasibility in terms of distance, time, and opening-hour constraints. The results also underline the scalability of the proposal. Simulations with different cloud configurations and multiple replicas show that the system can accommodate increased demand, either by deploying additional totems within the same city or by replicating the architecture in other urban contexts. Because portability only requires the integration of open city data, the proposal can be applied to other destinations without structural modifications.

It is important to acknowledge that the reinforcement learning profiling module has not been evaluated on public datasets, as existing benchmarks do not provide the interaction feedback required to evaluate the Q-learning module. As a consequence, the convergence analysis presented in this work relies on synthetic user interactions, where tourist feedback is simulated by perturbing a preference distribution with Gaussian noise. While this setup allows the characterization of the profiling dynamics under heterogeneous behaviors, it represents a limitation of the current evaluation. Future work will leverage real interaction data collected from the deployed totem to validate the profiling module in real conditions, through the real-world deployment of the totem in the city. However, the main contribution of this work lies not only in the individual techniques adopted but in the overall design of a complete system. The current deployment, implemented as a demonstrator of smart city technologies, serves as a proof of concept of how IoT can support tourism services. Ultimately, our system demonstrates how IoT-driven adaptive services can become a cornerstone of future smart tourism ecosystems, balancing personalization, scalability, and privacy.

REFERENCES

[1] D. C. Nguyen et al., "6G Internet of Things: A comprehensive survey," *IEEE Internet Things J.*, vol. 9, no. 1, pp. 359–383, Jan. 2022.

[2] O. Aouedi et al., "A survey on intelligent Internet of Things: Applications, security, privacy, and future directions," *IEEE Commun. Surveys Tuts.*, vol. 27, no. 2, pp. 1238–1292, Apr. 2025.

[3] J. Zhang, G. Wang, H. Chen, H. Huang, Y. Shi, and Q. Wang, "Internet of Things and extended reality in cultural heritage: A review on reconstruction and restoration, intelligent guided tour, and immersive experiences," *IEEE Internet Things J.*, vol. 12, no. 12, pp. 19018–19042, Jun. 2025.

[4] O. Olim, M. E. Balba, K. Khushvakt, and S. Muslimakhon, "IoT innovations for transforming the future of tourism industry: Towards smart tourism systems," *J. Intell. Syst. Internet Things*, vol. 14, no. 2, pp. 153–164, 2025.

[5] V. Tiwari, A. Mishra, and S. Tiwari, "Role of data safety and perceived privacy for acceptance of IoT-enabled technologies at smart tourism destinations," *Current Issues Tourism*, vol. 27, no. 19, pp. 3079–3094, Oct. 2024.

[6] S. Aggrawal, D. Teotia, Y. Yash, and S. Shukla, "Travel itinerary generation using AI and recommendation systems," in *Proc. 2nd Int. Conf. Comput. Intell., Commun. Technol. Netw. (CICTN)*, Feb. 2025, pp. 498–502.

[7] A. V. Shinde, D. D. Patil, and K. K. Tripathi, "A comprehensive survey on recommender systems techniques and challenges in big data analytics with IoT applications," *J. Law Sustain. Develop.*, vol. 11, no. 11, pp. e2243–e2243, 2023.

[8] K. Reddy and B. Sailesh, "Technology adoption in tourism and its effects on tourist behaviour," *J. Tourism Quart.*, vol. 6, nos. 3–4, pp. 53–73, 2024.

[9] M. Altulyan, L. Yao, X. Wang, C. Huang, S. S. Kanhere, and Q. Z. Sheng, "A survey on recommender systems for Internet of Things: Techniques, applications and future directions," *Comput. J.*, vol. 65, no. 8, pp. 2098–2132, Aug. 2022.

[10] J. L. Sarkar, A. Majumder, C. R. Panigrahi, S. Roy, and B. Pati, "Tourism recommendation system: A survey and future research directions," *Multimedia Tools Appl.*, vol. 82, no. 6, pp. 8983–9027, Mar. 2023.

[11] H. Ko, S. Lee, Y. Park, and A. Choi, "A survey of recommendation systems: Recommendation models, techniques, and application fields," *Electronics*, vol. 11, no. 1, p. 141, Jan. 2022.

[12] X. Yang, L. Zhang, and Z. Feng, "Personalized tourism recommendations and the E-tourism user experience," *J. Travel Res.*, vol. 63, no. 5, pp. 1183–1200, May 2024.

[13] C. Bin, T. Gu, Y. Sun, L. Chang, and L. Sun, "A travel route recommendation system based on smart phones and IoT environment," *Wireless Commun. Mobile Comput.*, vol. 2019, pp. 1–16, Jul. 2019.

[14] M. Tenemaza, S. Lujan-Mora, A. De Antonio, and J. Ramirez, "Improving itinerary recommendations for tourists through Meta-heuristic algorithms: An optimization proposal," *IEEE Access*, vol. 8, pp. 79003–79023, 2020.

[15] J. C. Cepeda-Pacheco and M. C. Domingo, "Deep learning and Internet of Things for tourist attraction recommendations in smart cities," *Neural Comput. Appl.*, vol. 34, no. 10, pp. 7691–7709, May 2022.

[16] G. Chen, "Tourism management strategies under the intelligent tourism IoT service platform," *Comput. Intell. Neurosci.*, vol. 2022, pp. 1–11, Apr. 2022.

[17] A. Majumder et al., "MERIT: Multi-itinerary tourist recommendation engine for industrial Internet of Things," in *Proc. IEEE Conf. Comput. Commun. Workshops (INFOCOM WKSHPS)*, May 2022, pp. 1–6.

[18] Y. Song and Y. He, "Toward an intelligent tourism recommendation system based on artificial intelligence and IoT using apriori algorithm," *Soft Comput.*, pp. 19159–19177, Mar. 2026.

[19] C.-L. Ho, W.-L. Chen, and C.-H. Ou, "Constructing a personalized travel itinerary recommender system with the Internet of Things," *Wireless Netw.*, vol. 30, no. 7, pp. 6555–6567, Oct. 2024.

[20] S. Arora, S. Pargaien, F. Khan, A. Mer, and A. Misra, "Itinerary planning for tourists using Internet of Things," in *Proc. 5th Int. Conf. Image Process. Capsule Netw. (ICIPCN)*, Jul. 2024, pp. 703–708.

[21] M. I. Mubarak and Z. K. A. Baizal, "Tourism route recommendation using reinforcement learning," in *Proc. IEEE 8th Int. Conf. Converg. Technol. (I2CT)*, Apr. 2023, pp. 1–7.

[22] A. Dalla Vecchia, S. Migliorini, E. Quintarelli, M. Gambini, and A. Belussi, "Promoting sustainable tourism by recommending sequences of attractions with deep reinforcement learning," *Inf. Technol. Tourism*, vol. 26, no. 3, pp. 449–484, Sep. 2024.

[23] S. Serengil and A. Özpinar, "A benchmark of facial recognition pipelines and co-usability performances of modules," *J. Inf. Technol.*, vol. 17, no. 2, pp. 95–107, Apr. 2024. [Online]. Available: https://dergipark.org.tr/en/pub/gazibtd/issue/84331/1399077

[24] S. Larsen, K. Wolff, R. Doran, and T. Øgaard, "What makes tourist experiences interesting," *Frontiers Psychol.*, vol. 10, p. 1603, Aug. 2019.

[25] D. Soldatenko and E. Backer, "A content analysis of cross-cultural motivational studies in tourism relating to nationalities," *J. Hospitality Tourism Manage.*, vol. 38, pp. 122–139, Mar. 2019.

[26] C. Hutto and E. Gilbert, "VADER: A parsimonious rule-based model for sentiment analysis of social media text," in *Proc. Int. AAAI Conf. Web Social Media*, vol. 8, May 2014, pp. 216–225.

[27] S. Katoch, S. S. Chauhan, and V. Kumar, "A review on genetic algorithm: Past, present, and future," *Multimedia Tools Appl.*, vol. 80, no. 5, pp. 8091–8126, Feb. 2021.

[28] B. Alhijawi and A. Awajan, "Genetic algorithms: Theory, genetic operators, solutions, and applications," *Evol. Intell.*, vol. 17, no. 3, pp. 1245–1256, Jun. 2024.

[29] A. Rezvanian, S. M. Vahidipour, and A. Sadollah, "An overview of ant colony optimization algorithms for dynamic optimization problems," in *Optimization Algorithms–Classics and Recent Advances*, M. Andriychuk and A. Sadollah, Eds., London, U.K.: IntechOpen, 2023, ch. 12, doi: 10.5772/intechopen.111839.

[30] T. Guilmeau, É. Chouzenoux, and V. Elvira, "Simulated annealing: A review and a new scheme," in *Proc. IEEE Stat. Signal Process. Workshop (SSP)*, Aug. 2021, pp. 101–105.

[31] S. M. Almufti et al., "Overview of metaheuristic algorithms," *Polaris Global J. Scholarly Res. Trends*, vol. 2, no. 2, pp. 10–32, 2023.

[32] M. Sarhani, S. Voß, and R. Jovanovic, "Initialization of metaheuristics: Comprehensive review, critical analysis, and research directions," *Int. Trans. Oper. Res.*, vol. 30, no. 6, pp. 3361–3397, Nov. 2023.

[33] R. Hermoso, S. Ilarri, and R. Trillo-Lado, "Enriched tourism dataset Paris (POIS)," Tech. Rep., 2024. [Online]. Available: https://data.mendeley.com/datasets/vh4g4g2322/1

[34] R. Hermoso, S. Ilarri, and R. Trillo-Lado, "Enriched tourism dataset London (POIS)," Tech. Rep., 2024. [Online]. Available: https://data.mendeley.com/datasets/gw9hjn4v65/1

[35] Ajuntament de Barcelona. (2015). *Cultural Interest Points in the City of Barcelona (Punts-Informacio-Turistica)*. [Online]. Available: https://opendata-ajuntament.barcelona.cat/data/en/dataset/punts-informacio-turistica

[36] (2026). *CommonPlace*. [Online]. Available: https://data.cityofnewyork.us/City-Government/CommonPlace/rxuy-2muj

Claudio Marche (Member, IEEE) received the M.Sc. degree (Hons.) in telecommunication engineering and the Ph.D. degree in electronic and computer engineering from the University of Cagliari, Cagliari, Italy, in 2018 and 2023, respectively.

Since graduation, he has been working as a Researcher with the Department of Electrical and Electronic Engineering, University of Cagliari. He is currently an Assistant Professor with the University of Cagliari, where he is a member of the Net4U Research Group. His current research interests include the Internet of Things (IoT), social networks, cloud-edge computing, and trust management.

Dr. Marche is a member of the Editorial Board for IEEE INTERNET OF THINGS JOURNAL and *Frontiers in Energy Efficiency*.

Vlad Popescu (Member, IEEE) received the M.Sc. degree in electronics and computer engineering and the Ph.D. degree in telecommunications from Transilvania University of Braşov, Braşov, Romania, in 1999 and 2006, respectively.

In 2000, he was with Malmö University, Malmö, Sweden, specializing in multimedia applications for four months. From 2001 to 2002, he was a Research Fellow with RWTH Aachen University, Aachen, Germany, where he finished the experimental part of his Ph.D. studies on wireless communication in underground environments from 2004 to 2005. Since 2009, he has been a Visiting Professor with the Department of Electrical and Electronic Engineering, University of Cagliari, Cagliari, Italy. Since 2000, he has been an Associate Professor with the Department of Electronics and Computers, Transilvania University of Braşov. He also collaborates closely with the Department of Electrical and Electronic Engineering, University of Cagliari, both on the research and teaching levels. His main research interests include telecommunications, cognitive radio systems, multimedia applications, and data acquisition.

Dr. Popescu is a member of IEEE BTS and IEEE AESS and is the Chair of the IEEE BTS Chapter in Romania.

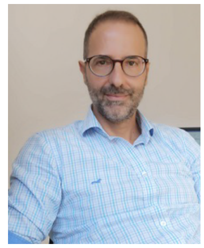

Luigi Atzori (Senior Member, IEEE) received the Ph.D. degree from University of Cagliari, Cagliari, Italy, in 2000.

Since April 2025, he has been the Head of the Department of Electrical and Electronic Engineering, University of Cagliari, Cagliari. He is currently a Professor of telecommunications with the University of Cagliari, where he leads the activities of the Networks for Humans (Net4U) Laboratory, with around 20 affiliated researchers. His research interests include the Internet of Things (IoT), with particular reference to the design of effective algorithms for the realization of social networks among connected devices to develop the Social IoT paradigm. He is the Co-Founder of two spinoffs in the areas of IoT solutions for people's mobility. His works on the IoT and the Social IoT have had a great impact, with more than 20K citations. His interests also fall in the area of the quality of experience (QoE), with particular application to the management of services and resources in new-generation networks for multimedia communications. Lately, he has also applied the study of the QoE to the IoT services.

Dr. Atzori is an Associate Editor of IEEE OPEN JOURNAL OF THE COMMUNICATIONS SOCIETY and *Digital Communications and Networks* (Springer).

Michele Nitti (Senior Member, IEEE) has been an Associate Professor with the University of Cagliari, Cagliari, Italy, since 2015. Moreover, he is currently the Co-Founder of an academic spin-off (GreenShare s.r.l.), which works in the mobility sector, Cagliari. His main research interests include the Internet of Things (IoT), particularly in the creation of a network infrastructure to allow objects to organize themselves according to a trustworthy social structure in smart city scenarios.

Dr. Nitti is a member of the Editorial Board for IEEE INTERNET OF THINGS JOURNAL, *Computer Networks* (Elsevier), and IoT (MDPI). He has been ranked in the "Top 2% World Ranking list of Scientists" for the years 2019–2023 and has received more than 4500 citations (source Google Scholar). He has been/is involved in the organization of several conferences, such as the Vertical and Topical Program Co-Chair of IEEE WF-IoT 2022, the General Chair of IEEE IoT V&T Summit 2021 on Tourism, and the Technical Program Co-Chair of IEEE BMSB 2017.

Open Access funding provided by 'Università degli Studi di Cagliari' within the CRUI CARE Agreement