

LLM-based Generation of Personalized, Context-aware City Tourist Itineraries: A User Study with GPT Trip Planner

Elena L. González-Sanz¹, Iván Cantador^{1,*} and Alejandro Bellogín¹

¹*Escuela Politécnica Superior, Universidad Autónoma de Madrid, Madrid, Spain*

Abstract

Personalized, context-aware trip planning is a critical challenge in e-tourism, requiring systems that dynamically adapt to individual preferences and situational factors. Recent advancements in Large Language Models (LLMs) present promising opportunities to address this challenge. In this paper, we aim to evaluate LLM capabilities in generating user-adapted city tourist itineraries through GPT Trip Planner (GPT-TP), a system that interacts iteratively with a GPT-4 LLM via prompting. GPT-TP recommends tailored itineraries based on user-defined interests and constraints, while retrieving comprehensive information about relevant tourist attractions. A user study with 30 participants was conducted to assess the system's effectiveness on 60 trip plans for 20 cities, focusing on the satisfaction of freely defined travel profiles. The evaluation considered key aspects such as the accuracy, relevance, and coverage of the plans' content, as well as the coherence, utility, and originality of the itineraries. The obtained results provided valuable insights into the strengths and limitations of leveraging LLMs for personalized, context-aware trip planning. Moreover, received user feedback let us define practical guidelines for designing adaptive itineraries in real-world e-tourism applications. To facilitate further research and ensure reproducibility, the source code of GPT-TP and the questionnaires of the user study are made publicly available.

Keywords

E-tourism, trip plans, tourism recommender systems, large language models, GPT

1. Introduction

E-tourism has long been recognized as a transformative force in the travel industry, fundamentally reshaping how individuals organize and experience trips [1]. The widespread adoption of digital platforms has streamlined access to tourist information, empowering users with sophisticated tools for travel decision-making and itinerary generation. The role of technology in enhancing user experiences is well-documented, with smart systems employing advanced algorithms and data-driven insights to optimize the efficiency and quality of tourist trips [2, 3, 4].

A key feature in this evolution is personalization [5, 6], which aims to tailor experiences to meet the travelers' unique interests and needs. Personalized systems in e-tourism thus address the diversity of user expectations, offering recommendations based on personal travel preferences and histories [7, 8]. This approach enables the generation of adaptive trip plans that enhance traveler satisfaction [9]. However, many traditional systems often struggle with incorporating contextual awareness and real-time adaptability, by accommodating factors such as schedules, budgets, and local events [9, 10, 11, 12].

Even current commercial trip planning applications, despite their widespread adoption and intuitive interfaces, exhibit significant limitations in terms of personalization and contextualization [13]. These applications typically offer standardized itineraries or generic recommendations, which often fail to capture the nuanced preferences of individual users or to adapt dynamically to contextual factors. As a result, the gap between the promise of personalized, context-aware tourism systems and the capabilities of state-of-the-art commercial solutions remains considerable.

Workshop on Recommenders in Tourism (RecTour 2025), September 22, 2025, co-located with the 19th ACM Conference on Recommender Systems, Prague, Czech Republic.

*Corresponding author.

✉ elenal.gonzalez@estudiante.uam.es (E. L. González-Sanz); ivan.cantador@uam.es (I. Cantador); alejandro.bellogin@uam.es (A. Bellogín)

ORCID 0000-0001-6663-4231 (I. Cantador); 0000-0001-6368-2510 (A. Bellogín)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Recent advancements in Large Language Models (LLMs) and generative artificial intelligence (GenAI) offer encouraging opportunities for personalized and context-aware trip planning. These models, aimed to understand and produce human-like text, have shown considerable potential in user modeling [14], information retrieval [15, 16], and recommender systems [17], among other fields. In the context of e-tourism, LLMs have been applied to tasks such as providing travel information, answering tourist inquiries, suggesting itineraries, and generating synthetic travel queries that emulate diverse user personas [18, 19, 20, 21]. Going beyond these applications, the ability of LLMs to process complex prompts responding to both user preferences and contextual conditions makes them a compelling tool for building tailored tourist itineraries [22].

Despite their potential, exploiting LLMs in dynamic and adaptive itineraries remains relatively underexplored [13]. While studies have investigated the use of LLMs for providing recommendations in e-tourism, there is a lack of comprehensive research on their capabilities to generate sophisticated trip plans [20]. In this sense, empirical evidence regarding the effectiveness –e.g., in terms of accuracy, relevance and comprehensiveness– of LLM-generated content, and the users’ satisfaction and trust is still limited [23]. To address these gaps, further research is required to evaluate the practical application of LLM-supported systems in e-tourism [24].

In this paper, we present GTP Tour Planner (GTP-TP), a system that leverages an LLM for generating personalized, context-aware city trip plans. This system interacts iteratively and via prompting with a GPT-4 model [25] to dynamically generate a structured trip plan for a city, incorporating user-defined preferences and contextual factors such as travelers’ attributes and constraints, attraction schedules, and travel distances. GTP-TP is thus aimed to generate highly valuable trip plan for given users and situations.

Similarly to previous work [9, 24], we conducted a user study involving 30 participants, who assessed GTP-TP’s capabilities to generate 60 city trip plans for 20 cities worldwide. The study focused on evaluating the personalization and contextualization of the generated trip plans, but also was intended to measure the accuracy, relevance, and coverage of the plans’ tourist attractions and information, as well as the perceived coherence, utility, and originality of their itineraries. The achieved results and received user feedback highlight strengths and limitations of using LLMs in trip planning, providing valuable insights into their potential in e-tourism applications, as we present in the form of practical guidelines.

The research questions driving our study were:

- **RQ1.** Can LLMs (specifically GPT-4o-mini) generate city trip plans that effectively account for a user’s personal preferences and contextual factors?
- **RQ2.** How accurate, relevant and comprehensive is the content of such trip plans regarding important tourist attractions and information of a city?
- **RQ3.** To what extent do the trip plans follow coherent itineraries that are both useful and original?

The remainder of the paper is organized as follows. Section 2 surveys related work on the use of LLMs in e-tourism and trip planning. Next, Section 3 introduces the GTP-TP system, and Section 4 delves into the personalized and context-aware trip plan generation process within the system. Finally, Sections 5 and 6 report on the conducted user study, and Section 7 concludes with principal findings of our work and potential avenues for future research.

2. Related Work

The exploitation of LLMs and GenAI for e-tourism is a recent research topic, whose studies are scarce and preliminary. In this section, we survey related work, without digging into the relatively extensive literature on trip planning [2, 4, 5, 6, 9, 10, 11, 12, 26, 27].

2.1. LLMs in E-tourism

LLMs can play a crucial role in supporting users throughout the three main stages of travel experiences: pre-trip, en-route, and post-trip [18, 20, 22]. In the pre-trip stage, LLMs can assist in researching destinations, comparing options, and generating personalized itineraries based on user preferences, constraints, and contextual data. During the en-route stage, LLMs can provide real-time updates, recommend nearby points of interest (POIs), and adapt travel plans in response to dynamic changes, such as weather conditions or user feedback. Finally, in the post-trip stage, LLMs can summarize experiences, assist in sharing trip highlights (e.g., via blogs or social media), and provide actionable insights for future travel planning [28].

Integrating LLMs into e-tourism applications is thus transforming how travelers plan and experience their journeys. LLM-supported systems have been proposed to generate custom itineraries, providing real-time recommendations, and summarizing travel experiences in a personalized manner [18]. Additionally, their ability to simulate personas, such as local residents or travel experts, enhances the authenticity and relatability of their guidance, and could increase user engagement [20].

In the tourism industry, LLMs are also valuable for customer service and operational tasks. Businesses can leverage LLMs to improve booking assistance, handle multilingual customer interactions, and automate back-office operations such as content creation, customer feedback analysis, and marketing campaigns [18, 23]. These applications boost efficiency and facilitate personalization, fostering customer loyalty [18].

Despite their effectiveness, LLMs raise challenges in e-tourism. Hallucinations remain a critical issue, as LLMs can generate inaccurate information, risking user trust [29]. Additionally, their reliance on outdated training data limits their utility in dynamic contexts [23]. Ethical concerns, such as biases in outputs and potential misuse of sensitive data, also require attention [18]. Future research should address these limitations by integrating external databases and specialized services and tools [23].

2.2. LLMs for Trip Planning

LLMs have shown significant potential in generating adaptive travel plans. These models are particularly effective at addressing user preferences and contextual factors through iterative interactions, enabling tasks like multi-day route optimization and POI recommendation [30]. The Map GPT Playground system showcases how LLMs can handle complex, multi-location queries, ensuring spatial and temporal consistency in travel plans [31]. The integration of LLM-based GenAI with Internet-of-Things (IoT) data supports real-time, personalized travel planning by incorporating dynamic environmental data [22]. This integration increases the contextual relevance of itineraries, adapting plans to evolving user needs and external conditions. In this context, our GPT-TP system allows establishing enables users to specify not only personal preferences, but also contextual factors to generate tailored trip plans.

Recent works have explored complementary directions for LLM-supported travel planning. In [19], Li investigates how to design a GUI that effectively supports trip planning with ChatGPT. Through a questionnaire-based study, the author proposes a form-based interface, which is well aligned with our own design in GPT-TP. Notably, our system extends this approach by allowing the specification of more nuanced preferences, including dietary restrictions, mobility needs, and temporal limitations. Moreover, unlike Li's work, which focuses on GUI design, our contribution evaluates the effectiveness of personalized itinerary generation based on structured user profiles and contextual information.

Other studies have examined the integration of LLMs with traditional planning frameworks. For instance, TRIP-PAL [32] combines GPT-4 with automated planners to optimize POI selection under minimal personalization input. The system only considers the destination city, number of POIs, and time constraints, using GPT-4 or an automated solver to produce a travel plan. Evaluation is conducted based on the utility (popularity) of selected POIs, with no in-depth consideration of user preferences or contextual dynamics. This contrasts with our approach, which emphasizes comprehensive personalization and contextual relevance in the generation of travel plans.




Figure 1: Screenshots of GPT-TP showing a map and POIs of Sydney, Australia.

Figure 2: Screenshots of GPT-TP showing part of the travel profile definition form.

Similarly, the LLM-Modulo framework [33] employs GPT-3.5 and GPT-4 to create travel plans using limited personalization parameters, such as destination, number of travelers, and budget. The study focuses on validating generated plans through critic-based agents that assess JSON structure completeness and POI diversity. While such framework emphasizes the robustness and validity of plan structure, it does not address user-specific preferences like accommodation type or dietary needs, which are core to our system’s design. Our work, in contrast, aims to evaluate how effectively an LLM can generate coherent and valuable itineraries grounded in detailed user and contextual profiles.

Despite the promise of LLMs in this domain, key challenges remain. One of the most critical issues is the risk of hallucinations, where models generate inaccurate or misleading content, potentially undermining user trust [30]. Moreover, the reliance on static or outdated training data can impair the temporal and contextual relevance of recommendations [31]. Addressing these limitations is crucial to improving the reliability and practical usability of LLM-based trip planning tools.

Through our user study with GPT-TP, we aim to assess the ability of a state-of-the-art LLM to generate travel plans that are not only aligned with user-defined constraints, but also accurate, coherent, comprehensive, and valuable—capturing both relevance and originality under diverse user and environmental conditions.

3. GPT-TP

GPT-TP is an open-source web application¹, leveraging the Render cloud platform to simplify deployment, hosting and scaling. The backend of GPT-TP is implemented in Python using the Django framework and relies on a PostgreSQL database. Its modular software architecture supports essential functionalities, including the management of user and travel profiles, and the LLM-supported acquisition of city POIs and POI metadata, and trip plan generation. A key feature of the backend is its integration with an external LLM via the LangChain library. Specifically, GPT-TP uses this library alongside the OpenAI API to make LLM prompt-based requests to the GPT-4o-mini model. Additionally, the backend

¹GPT-TP source code, <https://github.com/Elenaluciasanz/TourGPT>

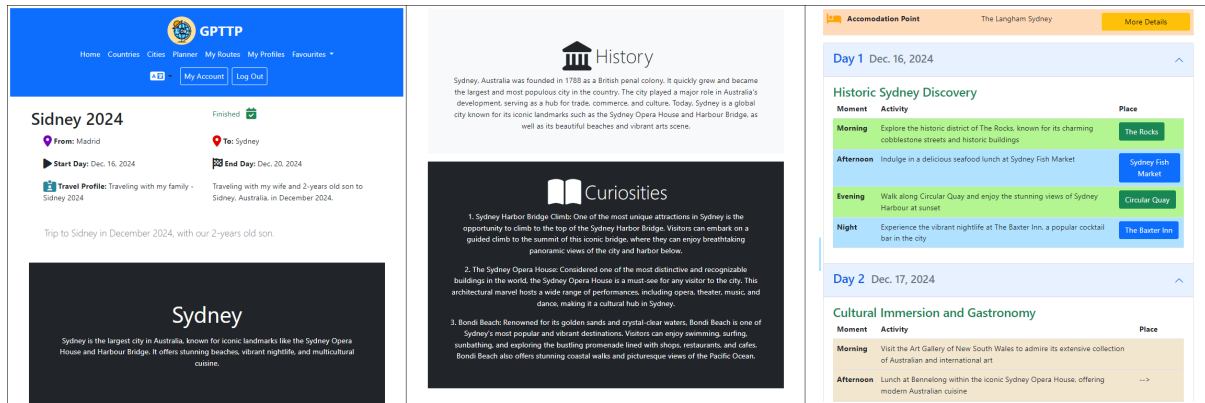


Figure 3: Screenshots of GPT-TP showing part of a generated trip plan for Sydney, Australia.

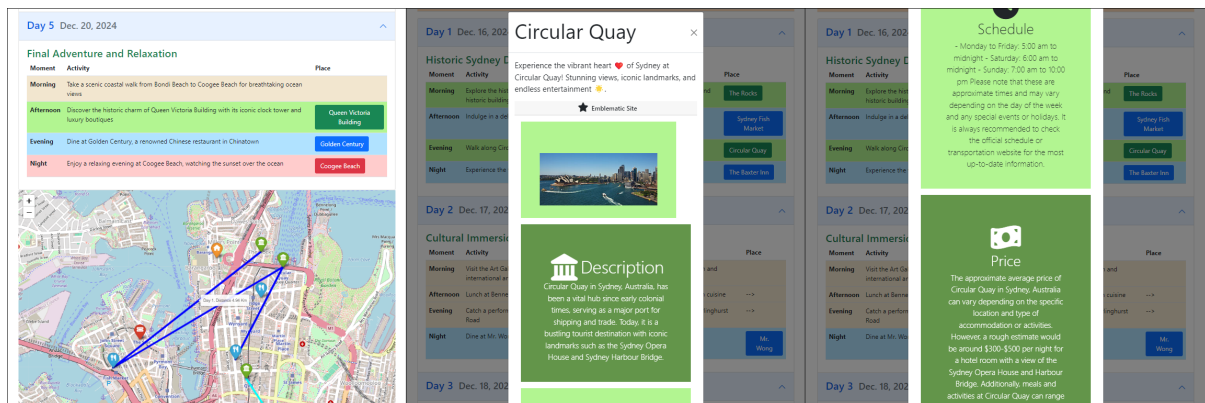


Figure 4: Screenshots of GPT-TP showing metadata of Circular Quay, a point of interest in Sydney, Australia.

incorporates several Python libraries to enhance its functionality. Geopy is used for geocoding and mapping geographical data, Google Translate provides automatic translation for multilingual support, and Requests interacts with the DBpedia endpoint via remote SPARQL queries to retrieve POI photos.

On the client side, GPT-TP features a responsive graphical user interface designed for a seamless and user-friendly experience. Built with Bootstrap, its frontend ensures compatibility across various devices and screen sizes. The interface uses AJAX for dynamic updates, allowing users to interact with the system without requiring full page reloads. To enhance personalization, GPT-TP integrates interactive maps powered by the Folium library, enabling users to visualize their itineraries and explore POIs directly on the maps.

As shown in Figure 1, the interface allows users to request POIs for a given city, and visualize them on an interactive map. The POIs are classified into different categories (i.e., *attraction*, *accommodation*, *entertainment*, and *dining*) and subcategories (e.g., emblematic site, museum, or park/garden under the *attraction* category). These categories are visually differentiated by colors on the interface.

The interface also includes a menu with access to a form where the user can request a trip plan for a city within a given date range, and considering a travel profile that the user defines in a separate window (Figure 2). Once generated, a trip plan is presented in a structured format. This includes a short introduction, historical insights, and curiosities about the city, followed by a day-by-day itinerary with POI recommendations of accommodation choices, attractions, entertainment venues, and dining places (Figure 3). The plan includes a final explanation detailing how it satisfies the input travel profile. Besides, the user can select a specific POI to view detailed information about it in a pop-up dialog window (Figure 4).

4. Trip Plan Generation

This section presents two core components of the GPT-TP system: the structure of the input travel profiles, capturing the user’s preferences and contextual constraints (Subsection 4.1), and the structure of the output trip plans, providing detailed, personalized, and context-aware itineraries (Subsection 4.2).

4.1. Travel Profiles

Listing 1 shows a fragment of the prompt used by GPT-TP to request a trip plan to the LLM. It encapsulates essential details about the travelers, including demographics, interests and needs. It specifies the destination city, the duration of the trip, and the composition of the traveling group, such as the number of adults and children, with children further categorized by age groups. The prompt also describes the purpose of the trip, whether it is for leisure, family bonding, or cultural exploration, and highlights specific POIs such as landmarks, and entertainment places.

The prompt concludes with a travel profile of practical considerations, such as the travelers’ budget, ranging from low to high, and the desired level of adventure, from relaxed activities to high-energy experiences. Observations and special requirements, such as accessibility needs or a preference for relaxed and kid-friendly locations, are also included. Finally, the profile allows specifying POIs to avoid, respecting particular user or system constraints.

4.2. Trip Plans

Listing 2 presents an example of a generated JSON output for a three-day trip to Sydney, Australia. The trip’s plan consists of a comprehensive itinerary designed to align with the input travel profile. It first includes a historical introduction and several curiosities of the city, and an accommodation recommendation.

The itinerary is organized by day and further divided into time slots, such as morning, afternoon, evening, and night, providing a structured sequence of activities. For each activity, the itinerary specifies the location, type, and an estimated budget level. Each day begins with a thematic header, such as “Beaches and Relaxation” or “Nature and Wildlife Fun,” which summarizes the planned activities and their overarching focus. A final explanation accompanies the itinerary, offering a concise justification of how the trip plan caters to the preferences and constraints outlined in the travel profile. This ensures transparency in how the system incorporates user-defined inputs into the final output.

5. Experimental Setting

We next detail the experimental setting of our user study, describing the followed evaluation methodology (Subsection 5.1) and the considered evaluation measures (Subsection 5.2).

5.1. Evaluation Methodology

The study was divided into three stages. In the first stage, the participants of the study were recruited, with invitations extended to family members, friends, and colleagues of the authors. They received a link to an online questionnaire which they completed confirming their participation. The questionnaire provided detailed information about the study’s purpose and assured participants of the anonymity and privacy of their data. In the questionnaire, participants were asked to provide their email addresses for communication purposes. Furthermore, they were requested to list 5–6 cities they had previously visited as tourists and another 5–6 cities they had not visited but wished to explore in the future.

Once all acceptance responses were received, participants were randomly assigned to one of two groups. A control group used GPT-TP with personalization and contextualization prompts disabled, while an experimental group used GPT-TP with these features enabled. For each participant, two cities were selected from their preference lists: one from the visited cities and the other from the non-visited

```

**Context:**
You are tasked with creating a personalized trip plan tailored to the interests, preferences, needs, and constraints of a provided travel profile. The itinerary should be detailed, optimized for a destination city, and accommodate to an input travel profile.

**Input description:**
To generate the trip plan you will be given:
1. City: The destination city.
2. Duration of stay: The number of days of the trip.
3. Travel profile: Detailed information about the travelers, including:
    - Adults: Number of adults.
    - Children: Number of children, categorized by age (baby: 0-2 years, child: 3-12 years, teen: 13-17 years).
    ...
4. Points forbidden: List of points of interest to avoid in the destination city, if applicable.

**Output instructions:**
Return a JSON object containing:
1. Accommodation ...
2. Daily itinerary ...
3. Explanation: A brief justification for why the itinerary is suitable for the input travel profile, highlighting how its interests and needs were considered in the plan.

**Example input:**
{
  "duration_of_stay": 3,
  "city": "Madrid, Spain",
  "travel_profile": {
    "adults": 2,
    "children": {"baby": 1, "child": 1},
    "purpose": "Family Vacation",
    "adventure_level": "Low",
    "interests": {"cultural": ["Historical landmarks", "Parks/Gardens"], "entertainment": ["Outdoor spaces", "Theme parks"], "gastronomy": ["Traditional restaurants"]},
    "budget": "Medium",
    "observations": "We prefer relaxed activities, historical exploration, and kid-friendly locations."
  },
  "forbidden_points": {"attractions": ["Plaza Mayor", "Retiro Park"], "entertainment": [], "gastronomy": ["Restaurante Casa Paco"], "accommodation": ["Riu Plaza Espana"]}
}

**Output JSON structure:**
...

```

Listing 1: Fragment of the prompt for trip plan generation. It has to be completed with the travel profile defined by the user.

cities. For this selection, global criteria were satisfied: the whole set of cities had to represent all continents, each city had to be selected at least once as a “visited city” and at least once as a “non-visited city,” and there had to be a balanced distribution of evaluations for each city in both categories. Additionally, to reduce biases, in each group, half of the participants began by evaluating a “visited city,” while the other half started with a “non-visited city.”

In the second stage, participants received an email containing their assigned user ID and password, a link to a study registration questionnaire² for collecting anonymized personal demographic and tourist behavior data, the names of their two assigned cities, and links to GPT-TP and a post-task evaluation questionnaire. The email also had detailed instructions for completing the study task.

²The study registration questionnaire is accessible at <https://forms.gle/XGMgEQRf6r67teb37>

```

{
  "accommodation": {"point": "Meriton Suites Waterloo", "subcategory": "hotel", "budget": "medium"},
  "itinerary": {
    "Day 1": {
      "header": "Historic and Iconic Sydney",
      "morning": {
        "activity": "Take a guided tour of the iconic Sydney Opera House",
        "point": "Sydney Opera House",
        "category": "attraction",
        "subcategory": "emblematic site",
        "budget": "medium"
      },
      "afternoon": {
        "activity": "Have lunch at a waterfront restaurant with views of the Opera House", "point": "Opera Kitchen", ...
      },
      "evening": {
        "activity": "Explore the historic streets and boutique shops of The Rocks", "point": "The Rocks", ...
      },
      "night": {
        "activity": "Dinner at a family-friendly restaurant overlooking the harbour", "point": "Quay Restaurant", ...
      }
    },
    "Day 2": {
      "header": "Beaches and Relaxation",
      "morning": {
        "activity": "Spend a fun morning at Bondi Beach building sandcastles", "point": "Bondi Beach", ...
      },
      "afternoon": {
        "activity": "Lunch at a beachfront cafe with kid-friendly options", "point": "Bondi Icebergs Club", ...
      },
      ...
    },
    "Day 3": {
      "header": "Nature and Wildlife Fun",
      "morning": {
        "activity": "Explore the lush greenery and wildlife at Taronga Zoo", ...
      }
    },
    ...
  },
  "explanation": "This itinerary balances family-friendly activities with a medium-budget preference. It includes iconic landmarks, outdoor attractions, and kid-friendly dining options, ensuring all ages enjoy the trip. Budget constraints were respected by mixing low-cost and medium-cost activities."
}

```

Listing 2: Fragment of a generated JSON output for a 3-days trip to Sidney, Australia.

Each participant was required to log in to GPT-TP using her credentials, and define a travel profile reflecting personal preferences and constraints for the requested city trip plans. Next, the participant had to use the system's travel planner to generate a trip plan for each of her two cities, freely choosing trip durations of 3–5 days. The system personalized and contextualized the recommended plans according to the travel profiles³. After the generation of a plan, the participant was instructed to review all provided information thoroughly, verifying details such as POIs, locations, prices, and opening and closing hours using external online resources.

³As explained in Section 6.1, in our study, GPT-TP only applied personalization and contextualization trip generation prompts for the members of an experimental group.

Once these tasks were completed for each trip plan, in the third stage, the participant had to fill out the post-task evaluation questionnaire. This questionnaire was aimed to evaluate several metrics, which are discussed in the subsequent subsection.

5.2. Evaluation Measures

Aimed at addressing the research questions stated in Section 1, the evaluation questionnaire⁴ was designed to assess multiple criteria related to the quality of the generated itineraries. In particular, participants were asked to rate:

- The relevance of the tourist attractions included in the plan.
- The coverage of the tourist attractions included in the plan considering the input duration.
- The accuracy (correctness, truthfulness) of the information (locations, visit times, prices) about the tourist attractions included in the plan.
- The originality of the plan.
- The coherence of the plan in terms of the types of places visited each day and time of day.
- The coherence of the plan in terms of visit and travel times.
- The satisfaction of your personal preferences for certain types of attractions and entertainment established in the travel profile.
- The satisfaction of your travel restrictions and mobility established constraints.
- The likelihood of following the suggested itinerary to visit the city under the established conditions.

The rationale behind the questionnaire was to evaluate the system across complementary dimensions:

- The GPT-TP's ability to personalize and contextualize trip plans, corresponding to the degree to which user-defined interests, preferences, and constraints were correctly satisfied.
- The accuracy, relevance, and coverage of the information provided about cities and POIs, ensuring that itineraries contained correct and sufficiently comprehensive content.
- The coherence, utility, and originality of the proposed itineraries, assessing whether generated trips were feasible, logically organized, and offered novel or valuable experiences.

The questionnaire employed a 1–5 Likert scale, with each response rating accompanied by a brief description. For example, the following questionnaire item evaluated the coverage of key POIs in a city within a trip plan, along with descriptions of the lowest- and highest-rating response options:

On a scale of 1 to 5, considering the duration you have set for the trip, how would you rate the coverage of the city's key points of interest included in the plan?

1. Unacceptable. The plan includes none or very few of the city's key points of interest.
2. Insufficient. The selection is limited and misses several key attractions that should be included.
3. Acceptable. The coverage is reasonable, although it could be more comprehensive.
4. Good. The plan adequately covers the key attractions of the city.
5. Excellent. The plan provides comprehensive and extensive coverage, including (almost) all the city's key points of interest.

6. User Study

Our user study took place in December 2024 and January 2025. In the following, we describe the individuals who participated in the study (Subsection 6.1), analyze the results of their evaluations (Subsection 6.2), and discuss limitations and guidelines derived from their feedback (Subsection 6.3).

⁴The evaluation questionnaire is accessible at <https://forms.gle/5Wb8LAtpZD9RxhiK8>

6.1. Participants

With informed agreement from participants, we recruited individuals ensuring a varied representation across demographics. A total of 30 people participated in the study, comprising 15 males and 15 females, with ages ranging 18-24 (26.9%), 35-44 (53.8%), 45-54 (15.4%) and 55-64 (3.8%) years old. Participants' educational levels were diverse: vocational training (15.4%), high school (7.7%), university studies (42.3%), master's degree (3.8%), and doctorate (30.8%).

In the registration questionnaire, 84.6% of participants reported traveling as tourists 1–2 times per year, and the remaining 15.4% between 3–5 times annually. They also indicated a broad spectrum of tourism preferences: relaxation trip (76.9%), cultural travel (61.5%), leisure travel (26.9%), shopping trip (15.4%), adventure or sports travel (11.5%), and gastronomic tourism (11.5%). For their upcoming trips, 15.4% of participants were planning to travel alone, 46.2% with a partner, 46.2% with family or friends including children, and 76.9% with family or friends and no children. Expected budgets varied as well, with 46.2% indicating economic, 30.8% economic or mid-range, and 23.1% mid-range budgets. This diversity in tourist profiles was deemed suitable for the scope of our study.

In general, prior to traveling as tourists to a city, 34.6% of participants usually gather little to no information –perhaps only a basic list of the city's main POIs. Meanwhile, 65.4% collect basic data about popular POIs, 26.9% gather detailed information beyond the city's main attractions, and only 3.8% collect exhaustive information, seeking lesser-known data. These findings suggest that the GPT-TP system's design is well-suited, providing standout information about cities and POIs without delving into extensive details as traditional travel guides commonly provide.

A notable finding from the questionnaire responses was the clear preference for digital resources when planning trips. Among them, platforms and applications such as Google Maps, TripAdvisor, Booking, and Expedia were the most frequently used, reported by 50.0% of participants. Social media like Instagram, Pinterest, and YouTube were also popular, used by 34.6%. In contrast, specialized tourism-focused digital resources showed a lower level of adoption: 26.9% of participants reported using travel blogs or forums, while only 3.8% used advanced digital trip-planning platforms like Google Travel, TripIt, or Sygic Travel. Traditional resources had varying degrees of popularity: printed travel guides such as Lonely Planet were consulted by 26.9% of participants, while travel agencies were used by just 3.8%. Interestingly, personal recommendations played a significant role in travel planning, with 50.0% of participants considering advice from friends or family when organizing their trips. These insights highlight the need to evaluate the perceived usefulness and appeal of systems like GPT-TP, which could bridge the gap between broad-use and specialized e-tourism platforms.

6.2. Results

In the user study, participants accessed GPT-TP using their personal devices, including desktop computers (18.2%), laptops (50.0%), tablets (9.1%), and mobile phones (22.7%), to perform the experiment tasks. No significant differences were observed in the experiment results based on the type of device used.

Participants freely used GPT-TP to define travel profiles to request the generation of 60 trip plans (with an average duration of 4.3 days) for 20 cities worldwide. These cities were Amsterdam, Athens, Barcelona, Copenhagen, London, Lisbon, Madrid, Paris, Prague, Rome, Tenerife, and Venice in Europe; New York, and San Francisco in North America, Rio de Janeiro in South America; Seoul, Shanghai, and Tokyo in Asia; Cairo in Africa; and Sydney in Australia.

Each participant used the system to request trip plans for two cities: one familiar (visited) and one unfamiliar (non-visited). In the control group, user-defined travel plans were excluded from the input prompts, while in the experimental group, these plans were included to enhance personalization and contextualization.

Table 1 summarizes the study results. The performance of the experimental group illustrates how incorporating user input into the LLM prompts significantly improves the quality of the generated trip plans across various evaluation criteria. To analyze these results in depth, we structure our assessment around the three research questions introduced in Section 1.

Table 1

Average (1-5 scale) ratings given by participants of the user study. Green highlights indicate the highest ratings, while red highlights represent the lowest ratings across categories. Results are compared between the control and experimental groups, and further broken down by non-visited and visited cities.

	Control group		Experimental group	
	<i>Non-visited cities</i>	<i>Visited cities</i>	<i>Non-visited cities</i>	<i>Visited cities</i>
Satisfaction of preferences	3,67	3,46	4.13	4.11
context	3.42	3.69	4.19	4.00
Content accuracy	3.92	3.62	3.47	4.33
relevance	3.92	3.54	4.13	4.44
coverage	3.58	3.08	4.06	4.11
Itinerary POI coherence	3.83	3.31	4.13	4.00
time coherence	3.00	3.31	3.65	3.78
utility	3.50	3.23	3.71	3.78
originality	3.50	3.62	3.56	3.90

RQ1. Can LLMs generate city trip plans that effectively account for a user’s personal preferences and contextual factors? The results show that GPT-TP successfully integrated user-defined travel profiles into its trip plans. Participants in the experimental group rated the satisfaction of user preferences highly, with scores of 4.13 for unvisited cities and 4.11 for visited cities. These ratings suggest that participants perceived the system’s personalization as effective across both familiar and unfamiliar cities. Similarly, context satisfaction received positive evaluations, with scores of 4.19 for unvisited cities and 4.00 for visited cities, further supporting the GPT-TP’s ability to adapt to situational conditions from user input. In comparison, the control group expressed lower satisfaction, with scores of 3.67 and 3.42 for user preferences and contextual factors in unvisited cities, and 3.46 and 3.69 in visited cities.

RQ2. How accurate, relevant, and comprehensive is the content of such trip plans regarding important tourist attractions and city information? Content quality was another area where GPT-TP performed well, particularly in the experimental group. Participants validated the accuracy of the trip plans, especially for visited cities, which scored 4.33. However, the lower score for unvisited cities (3.47) suggests that GPT-TP’s performance could benefit from improved handling of unfamiliar contexts. Despite this, participants appreciated the relevance of the included POIs, with high ratings for unknown (4.13) and known (4.44) cities. Content coverage was similarly well-rated, with participants reporting that the plans included most relevant POIs, reflected in scores of 4.06 for unvisited cities and 4.11 for visited cities. By contrast, the control group scored lower across all dimensions, especially for visited cities, where participants likely had higher expectations.

RQ3. To what extent do the trip plans follow coherent itineraries that are both useful and original? The coherence and utility of the itineraries generated by GPT-TP also received favorable evaluations. Participants rated POI coherence highly, with scores of 4.13 for unvisited cities and 4.00 for visited cities. However, time coherence received slightly lower ratings of 3.65 and 3.78, respectively, suggesting that while the system generally organized POIs in a logical sequence, there is room for improvement in time allocation. The utility of the generated itineraries, particularly in terms of personalization and contextualization, scored 3.78 for visited cities. Originality was perceived as moderate, scoring 3.90 for visited cities, indicating potential for more innovative itinerary designs. As with other research questions, the control group assigned consistently lower ratings to all measures of itinerary quality.

Overall, the experimental group consistently outperformed the control group, showing the positive impact of including user-defined travel plans in the input LLM prompts. This approach enhanced GPT-TP’s ability to personalize and contextualize trip plans effectively, resulting in higher satisfaction across most evaluation criteria. Nevertheless, the lower ratings for unvisited cities reveal a challenge in adapting to less familiar contexts. Additionally, aspects such as time coherence and originality, while positively evaluated, highlight areas for potential improvement.

To address these limitations, future iterations of GPT-TP could incorporate features that provide transparent explanations for the system’s recommendations. For instance, the system could explain why specific POIs were selected, how time allocations balance exploration with efficiency, or how the grouping of POIs reflects proximity, theme, or user preferences. This functionality could build user trust and further enhance the perceived quality of the generated plans.

Moreover, improving the handling of unvisited cities may involve integrating external data sources, such as updated travel databases or user feedback loops, to enhance the relevance and coverage of information in less familiar contexts. By addressing these challenges through transparent communication and iterative system refinement, GPT-TP can further solidify its position as a reliable and engaging tool for adaptive trip planning.

6.3. Limitations and Guidelines

The feedback collected from participants in the user study highlighted potential improvements for the generated trip plans, offering valuable insights for developing more effective LLM-based trip planning systems. Hence, the received suggestions allowed us to provide practical guidelines for refining the planning process.

A key issue identified was the existence of certain **hallucinations** on the geographical location of POIs in plan slots. To address this, it is essential to implement curated LLM prompts in a more structured fashion. A potential approach may involve a three-step process: first, identifying tourist districts or key areas within the target city; second, assigning these districts to specific time slots based on contextual factors such as travel times, opening hours, and user preferences; and third, retrieving POIs for each district while adhering to user-defined constraints. Integrating reasoning-oriented techniques may further enhance the reliability of the generated plans by enabling the system to validate its choices systematically [7, 30, 31].

Another improvement involves optimizing **time and distance constraints** within itineraries. Some participants indicated that the plans underutilized the available time, leaving room for more POIs to be included. To overcome this, LLM-based systems could incorporate detailed travel distance and time estimates between POIs [9], supplemented with transportation options like public transit, walking routes, or ride-sharing services. Explicitly presenting these details in the plans would improve user trust and strengthen the feasibility of the itineraries. In this context, it could also be valuable to introduce a user preference parameter that characterizes the desired travel style —e.g., relaxed, balanced, or intensive. This would allow the system to adjust the number of POIs and time allocations accordingly, ensuring itineraries better reflect individual travel expectations. For all cases, providing explanations of how time allocations were determined would enhance the transparency of the system’s recommendations.

Participants also highlighted the importance of offering **recommendations tailored to specific POIs**. For example, users might benefit from suggestions for nearby dining options during meal times or activity suggestions in certain areas of the city [7, 9]. To implement this, systems could generate personalized, context-aware sub-lists of options for each POI and time slot. Such functionality would significantly enhance the practicality of the plans.

Another mentioned limitation was certain over-reliance on popular POIs, which limited the **diversity in the itineraries**. To address this, diversification techniques can be employed, such as boosting user-defined preferences to prioritize lesser-known attractions or further balancing recommendations to include a mix of iconic and hidden gems [4, 5]. Encouraging users to specify their desired level of exploration —ranging from mainstream to off-the-beaten-path experiences— could guide the system in tailoring the itinerary.

Richer **contextual adaptability** emerged as another area for enhancement. Effective trip planning should consider real-time factors like weather conditions, local events, and visitor density at POIs [11]. Integrating data from external sources, such as weather APIs and event databases, could enable further adaption of itineraries. For instance, recommending indoor activities during adverse weather or highlighting festivals and cultural events during the users’ visit would greatly improve the relevance and appeal of the plans.

Finally, a few participants emphasized the importance of addressing **accessibility** and **sustainability** in trip planning, as discussed in [8, 22]. Systems should strive to include accessibility information, like wheelchair-friendly routes and POIs with facilities for individuals with disabilities. Similarly, promoting eco-friendly travel options –e.g., walking paths, bike rentals, and low-emission transportation– aligns with growing demand for sustainable tourism. Moreover, reducing the focus on popular POIs could also help alleviate issues related to tourist overcrowding, a growing concern in urban tourism management [34] and POI recommender systems [35]. In this context, systems could integrate sustainability metrics, enabling users to select itineraries that minimize environmental impact while maximizing enjoyment.

By adopting these practical guidelines –structured POI retrieval and trip planning, time and distance optimization, diversified recommendations, contextual adaptability, and support for accessibility and sustainability– developers could create more reliable, inclusive, and user-centric trip planning systems. These improvements not only address the limitations identified in our user study, but also set the stage for delivering truly personalized and engaging travel experiences powered by LLMs.

7. Conclusions

In this paper, we have presented GPT-TP, a novel system leveraging an LLM for personalized and context-aware trip planning. By iteratively prompting a GPT-4 model, the system builds trip plans tailored to user-defined travel preferences and situational constraints.

This has enabled us to explore how personalization influences the quality and usefulness of AI-generated trip itineraries, filling a research gap in the assessment of LLMs in itinerary generation. Related work also applied LLM prompting for automatic travel plan generation, but did not delve into the evaluation of generated plans as we do in our research.

Our conducted user study highlighted the GPT-TP’s ability to effectively personalize and contextualize the plans, according to accuracy, relevance, coherence, and originality criteria. Although additional LLMs and **evaluation metrics** should be considered, such as user satisfaction and trust, and system accessibility and usability [9, 24, 36], the obtained findings underscore the transformative potential of LLMs in e-tourism, showcasing their ability to enhance the user experience by offering adaptive and flexible trip planning.

Accompanying these results, there are several opportunities for further research and development. One important avenue is the integration of specialized **external services and data sources** to improve the accuracy and enhance the descriptions and metadata of POI recommendations in generated trips [37]. This could include linking the system to real-time data feeds for transportation schedules, weather updates, and event listings, as well as leveraging knowledge graphs for deeper cultural and historical insights. In this sense, incorporating Retrieval-Augmented Generation (RAG) [38] methods may be a promising direction. By combining LLMs with external search engines, the system could reduce reliance on model-internal knowledge, thereby improving factual accuracy and mitigating hallucinations. A RAG approach could also ensure that recommendations remain up-to-date and contextually precise.

Another potential improvement involves adding a **conversational interface** [30, 39], which would enable users to interact with the system as if it were a human guide or tourism expert. A dialogue-based communication could support iterative refinement of itineraries, allowing users to better and dynamically adjust preferences and constraints or resolve ambiguities in real-time, enhancing the system’s adaptability and usability. Besides, incorporating a virtual agent capable of delivering explanations under different, configurable roles –such as a local resident, a professional tour guide, or an expert in art or history– would add a layer of engagement and contextual depth [20]. These personalized narratives would cater to diverse user preferences, making the generated trip plans more relatable and valuable. The addition of a voice interface could further enhance the human-computer interactions, ultimately increasing the users’ satisfaction with their travel experience [18].

In addition to the promising outcomes and open research issues, it is important to acknowledge several limitations of our work. First, our study represents a preliminary, exploratory evaluation, aimed primarily at obtaining initial insights into user perceptions of LLM-generated trip plans. While GPT-TP demonstrated certain ability to generate coherent and contextually relevant itineraries, current LLMs inherently lack true reasoning and planning capabilities. As such, we cannot expect them to consistently produce itineraries that fully satisfy spatial or temporal coherence, optimize routes, consider distances and transport modes, or balance popularity and diversity of POIs without extensive human guidance.

Moreover, the standard criteria for classical itinerary planners –such as ensuring accessibility, sustainability, real-time contextual adaptation, and consequently multi-objective optimization– remain challenging for current LLMs, even when combined with RAG methods or sophisticated prompt engineering. Our study focused on subjective user experience in an uncontrolled environment and did not include benchmarking against traditional, rule-based planners or route optimization algorithms. Therefore, while GPT-TP provides valuable insights into perceived usefulness and user satisfaction, its performance should not be interpreted as a comprehensive assessment of itinerary optimization.

Nonetheless, we argue that the presented user study and findings, and the open-source implementation of GPT-TP offer a valuable foundation for future research. Researchers can leverage our publicly available web application to explore the integration of LLMs with classical reasoning and planning frameworks, multi-agent systems, or hybrid architectures that combine data-driven suggestions with optimization-based itinerary refinement. Such extensions could address current limitations in spatial and temporal coherence while preserving the adaptive, context-aware advantages of LLM-generated content.

Despite the constraints, in our humble opinion, our findings highlight the potential of LLMs to enhance the user experience in e-tourism applications. The flexibility, responsiveness, and personalization offered by GPT-TP illustrate a promising direction in which LLMs can complement existing planning tools, particularly in providing engaging, human-like guidance and explanations tailored to individual travel profiles. We anticipate that continued research in this direction –particularly through hybrid LLM-planners, probably through multi-agent interactions [40]– could unlock more robust, accurate, and contextually rich itinerary generation in the near future.

Acknowledgments

This work was supported by Grant PID2022-139131NB-I00 funded by MCIN/AEI/10.13039/501100011033 and by “ERDF, a way of making Europe.” The authors are grateful to the participants of the user study presented in this paper for their time and valuable insights. The authors also wish to thank the reviewers for their constructive comments and suggestions, as well as the RecTour organizers for granting the opportunity to present the work at the workshop.

Declaration on Generative AI

During the preparation of this work, the authors used GPT-4 in order to: Paraphrase and reword. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication’s content.

References

- [1] C. D. Huang, J. Goo, K. Nam, C. W. Yoo, Smart tourism technologies in travel planning: the role of exploration and exploitation, *Information & Management* 54 (2017) 757–770. doi:10.1016/j.im.2016.11.010.
- [2] S. Kotiloglu, T. Lappas, K. Pelechris, P. Repoussis, Personalized multi-period tour recommendations, *Tourism Management* 62 (2017) 76–88. doi:10.1016/j.tourman.2017.03.005.

- [3] R. A. Hamid, A. S. Albahri, J. K. Alwan, Z. Al-Qaysi, O. S. Albahri, A. A. Zaidan, A. Alnoor, A. H. Alamoodi, B. Zaidan, How smart is e-tourism? a systematic review of smart tourism recommendation system applying data management, *Computer Science Review* 39 (2021) 100337. doi:10.1016/j.cosrev.2020.100337.
- [4] Z. Yu, H. Xu, Z. Yang, B. Guo, Personalized travel package with multi-point-of-interest recommendation based on crowdsourced user footprints, *IEEE Transactions on Human-Machine Systems* 46 (2015) 151–158. doi:10.1109/THMS.2015.2446953.
- [5] I. Benouaret, D. Lenne, Recommending diverse and personalized travel packages, in: *Proceedings of the 28th International Conference on Database and Expert Systems Applications (DEXA '17)*, Part II, Springer, 2017, pp. 325–339. doi:10.1007/978-3-319-64471-4_26.
- [6] P. Thiengburanatham, S. Cang, H. Yu, Overview of personalized travel recommendation systems, in: *Proceedings of the 22nd International Conference on Automation and Computing (ICAC '16)*, IEEE, 2016, pp. 415–422. doi:10.1109/ICoAC.2016.7604955.
- [7] J. Borràs, A. Moreno, A. Valls, Intelligent tourism recommender systems: a survey, *Expert Systems with Applications* 41 (2014) 7370–7389. doi:10.1016/j.eswa.2014.06.007.
- [8] K. Chaudhari, A. Thakkar, A comprehensive survey on travel recommender systems, *Archives of Computational Methods in Engineering* 27 (2020) 1545–1571. doi:10.1007/s11831-019-09363-7.
- [9] H.-S. Chiang, T.-C. Huang, User-adapted travel planning system for personalized schedule recommendation, *Information Fusion* 21 (2015) 3–17. doi:10.1016/j.inffus.2013.05.011.
- [10] M. Kolahkaj, A. Harounabadi, A. Nikravanshalmani, R. Chinipardaz, A hybrid context-aware approach for e-tourism package recommendation based on asymmetric similarity measurement and sequential pattern mining, *Electronic Commerce Research and Applications* 42 (2020) 100978. doi:10.1016/j.eierap.2020.100978.
- [11] A. Majid, L. Chen, G. Chen, H. T. Mirza, I. Hussain, J. Woodward, A context-aware personalized travel recommendation system based on geotagged social media data mining, *International Journal of Geographical Information Science* 27 (2013) 662–684. doi:10.1080/13658816.2012.696649.
- [12] S. Renjith, A. Sreekumar, M. Jathavedan, An extensive study on the evolution of context-aware personalized travel recommender systems, *Information Processing & Management* 57 (2020) 102078. doi:10.1016/j.ipm.2019.102078.
- [13] M. Gupta, R. Dias, N. Jain, R. R. Dhamija, R. V. Bidwe, S. Mishra, G. Deshmukh, Travel with Generator AI: a novel approach to itinerary creation, in: *Proceedings of the 8th International Conference on Computing, Communication, Control and Automation (ICCUBEA '24)*, IEEE, 2024, pp. 1–6. doi:https://doi.org/10.1109/ICCUBEA61740.2024.10775161.
- [14] Z. Tan, M. Jiang, User modeling in the era of large language models: current research and future directions, *arXiv preprint arXiv:2312.11518* (2023). doi:10.48550/arXiv.2312.11518.
- [15] Q. Ai, T. Bai, Z. Cao, Y. Chang, J. Chen, Z. Chen, Z. Cheng, S. Dong, Z. Dou, F. Feng, et al., Information retrieval meets large language models: a strategic report from Chinese IR community, *AI Open* 4 (2023) 80–90. doi:10.48550/arXiv.2307.09751.
- [16] C. Zhai, Large language models and future of information retrieval: opportunities and challenges, in: *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*, ACM, 2024, pp. 481–490. doi:10.1145/3626772.3657848.
- [17] L. Wu, Z. Zheng, Z. Qiu, H. Wang, H. Gu, T. Shen, C. Qin, C. Zhu, H. Zhu, Q. Liu, et al., A survey on large language models for recommendation, *World Wide Web* 27 (2024) 60. doi:10.1007/s11280-024-01291-2.
- [18] I. Carvalho, S. Ivanov, ChatGPT for tourism: applications, benefits and risks, *Tourism Review* 79 (2024) 290–303. doi:10.1108/TR-02-2023-0088.
- [19] B. Li, EverywhereGPT: an AI travel planning assistant based on ChatGPT, in: *Proceedings of the 4th International Conference on Artificial Intelligence and Computer Engineering*, ACM, 2023, pp. 995–1003. doi:10.1145/3652628.3652793.
- [20] I. A. Wong, Q. L. Lian, D. Sun, Autonomous travel decision-making: An early glimpse into ChatGPT and generative AI, *Journal of Hospitality and Tourism Management* 56 (2023) 253–263. doi:10.1016/j.jhtm.2023.06.022.

- [21] A. Banerjee, A. Satish, F. N. Aisyah, W. Wörndl, Y. Deldjoo, SynthTRIPs: a knowledge-grounded framework for benchmark data generation for personalized tourism recommenders, in: *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, ACM, 2025, pp. 3743–3752. doi:<https://doi.org/10.1145/3726302.3730321>.
- [22] P. Suanpang, P. Pothipassa, Integrating Generative AI and IoT for sustainable smart tourism destinations, *Sustainability* 16 (2024) 7435. doi:[10.3390/su16177435](https://doi.org/10.3390/su16177435).
- [23] L. Mich, R. Garigliano, ChatGPT for e-tourism: a technological perspective, *Information Technology & Tourism* 25 (2023) 1–12. doi:[10.1007/s40558-023-00248-x](https://doi.org/10.1007/s40558-023-00248-x).
- [24] F. Ali, B. Yasar, L. Ali, S. Dogan, Antecedents and consequences of travelers’ trust towards personalized travel recommendations offered by chatgpt, *International Journal of Hospitality Management* 114 (2023) 103588. doi:[10.1016/j.ijhm.2023.103588](https://doi.org/10.1016/j.ijhm.2023.103588).
- [25] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al., GPT-4 technical report, *arXiv preprint arXiv:2303.08774* (2023). doi:[10.48550/arXiv.2303.08774](https://doi.org/10.48550/arXiv.2303.08774).
- [26] D. Gavalas, V. Kasapakis, C. Konstantopoulos, G. Pantziou, N. Vathis, C. Zaroliagis, The eCOMPASS multimodal tourist tour planner, *Expert systems with Applications* 42 (2015) 7303–7316. doi:<https://doi.org/10.1016/j.eswa.2015.05.046>.
- [27] P. Vansteenwegen, W. Souffriau, G. V. Berghe, D. Van Oudheusden, The City Trip Planner: an expert system for tourists, *Expert Systems with Applications* 38 (2011) 6540–6546. doi:[10.1016/j.eswa.2010.11.085](https://doi.org/10.1016/j.eswa.2010.11.085).
- [28] D. Mladenović, M. Beheshti, T. Kolar, E. Ismagilova, Y. K. Dwivedi, Synthetic WOM? The emergence of generative artificial intelligence-induced recommendations, *Journal of Computer Information Systems* (2024) 1–18. doi:[10.1080/08874417.2024.2409252](https://doi.org/10.1080/08874417.2024.2409252).
- [29] J. Christensen, J. M. Hansen, P. Wilson, Understanding the role and impact of Generative Artificial Intelligence (AI) hallucination within consumers’ tourism decision-making processes, *Current Issues in Tourism* (2024) 1–16. doi:[10.1080/13683500.2023.2300032](https://doi.org/10.1080/13683500.2023.2300032).
- [30] J. Lin, N. Tomlin, J. Andreas, J. Eisner, Decision-oriented dialogue for human-AI collaboration, *Transactions of the Association for Computational Linguistics* 12 (2024) 892–911. doi:[10.1162/tac1_a_00679](https://doi.org/10.1162/tac1_a_00679).
- [31] C. Zhang, A. Karatzoglou, H. Craig, D. Yankov, Map GPT playground: smart locations and routes with GPT, in: *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL ’23)*, ACM, 2023, pp. 1–4. doi:[10.1145/3589132.3625595](https://doi.org/10.1145/3589132.3625595).
- [32] T. de la Rosa, S. Gopalakrishnan, A. Pozanco, Z. Zeng, D. Borrajo, TRIP-PAL: travel planning with guarantees by combining large language models and automated planners, *arXiv preprint arXiv:2406.10196* (2024). doi:[10.48550/arXiv.2406.10196](https://doi.org/10.48550/arXiv.2406.10196).
- [33] A. Gundawar, M. Verma, L. Guan, K. Valmeekam, S. Bhambri, S. Kambhampati, Robust planning with LLM-Modulo framework: case study in travel planning, *arXiv preprint arXiv:2405.20625* (2024). doi:[10.48550/arXiv.2405.20625](https://doi.org/10.48550/arXiv.2405.20625).
- [34] H. Seraphin, P. Sheeran, M. Pilato, Over-tourism and the fall of Venice as a destination, *Journal of Destination Marketing & Management* 9 (2018) 374–376. doi:[10.1016/j.jdmm.2018.01.011](https://doi.org/10.1016/j.jdmm.2018.01.011).
- [35] S. Migliorini, D. Carra, A. Belussi, Distributing tourists among POIs with an adaptive trip recommendation system, *IEEE Transactions on Emerging Topics in Computing* 9 (2019) 1765–1779. doi:[10.1109/TETC.2019.2920484](https://doi.org/10.1109/TETC.2019.2920484).
- [36] D. P. Stergiou, A. Nella, ChatGPT and tourist decision-making: an accessibility–diagnosticity theory perspective, *International Journal of Tourism Research* 26 (2024) e2757. doi:[10.1002/jtr.2757](https://doi.org/10.1002/jtr.2757).
- [37] C.-C. Yu, H.-p. Chang, Personalized location-based recommendation services for tour planning in mobile tourism applications, in: *Proceedings of the 10th International Conference on E-Commerce and Web Technologies (EC-Web ’09)*, Springer, 2009, pp. 38–49. doi:[10.1007/978-3-642-03964-5_5](https://doi.org/10.1007/978-3-642-03964-5_5).
- [38] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, et al., Retrieval-augmented generation for knowledge-intensive NLP tasks, in:

Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS '20), volume 33, 2020, pp. 9459–9474. doi:10.48550/arXiv.2005.11401.

- [39] S. K. Dam, C. S. Hong, Y. Qiao, C. Zhang, A complete survey on LLM-based AI chatbots, arXiv preprint arXiv:2406.16937 (2024). doi:10.48550/arXiv.2406.16937.
- [40] B. Liu, J. Ge, J. Wang, Vaiage: a multi-agent solution to personalized travel planning, arXiv preprint arXiv:2505.10922 (2025). doi:10.48550/arXiv.2505.10922.