Integrating Temporal Planning and Knowledge Representation to Generate Personalized Touristic Itineraries

Silvia Gola[1], Donatella Capaldi[2], Alessandra Chivirì[3], Mohamed Ali Jaziri[3], Laura Leopardi[2], Saverio Giulio Malatesta[2], Irene Muci[3], Andrea Orlandini[1] Alessandro Umbrico[1], and Alberto Bucciero[3]

[1] CNR - Institute of Cognitive Sciences and Technologies
[2] Sapienza University - (DIGILAB)
[3] CNR - Institute of Heritage Science

Abstract. The *HERitage sMart social mEdia aSsistant* project offers innovative services enabling contextualized and multi-perspective, cross-cultural explorations of the rich and various cultural heritage of a territory. The project proposes the integration of Artificial Intelligence technologies to contextualize and personalize cultural paths according to users' interests and implicit/explicit relationships among tangible and intangible cultural entities. This work describes the designed AI-based architecture, the integration of Ontology-based Knowledge Representation and Reasoning, and Automated Planning to achieve the needed levels of contextualization and customization.

Keywords: Cultural Heritage · Knowledge Representation and Reasoning · Automated Planning · Customization

1 Introduction

The tourism industry is undergoing a profound transformation driven by the diffusion of smart technologies. Smart applications have emerged as powerful tools capable of enhancing the travel experience by offering personalized and dynamic touristic itineraries. This evolution is particularly significant in cultural heritage, where the richness and diversity of historical, artistic, and cultural assets offer many opportunities for tailored visitor experiences [13, 16, 8]. The diffusion of apps on smartphones offers enriched narrative experiences, allowing visitors to engage with cultural heritage in innovative ways. Through interactive content, users can gain deeper insights into historical contexts and cultural stories, fostering a more meaningful connection with the explored locations.

This paper presents the main features of a smart application developed as part of a research initiative promoted by Regione Lazio. The *HERitage sMart social mEdia aSsistant* (HerMeS) project, a joint effort of the National Research Council of Italy and La Sapienza University, aims to offer tools and innovative services to promote Lazio's Cultural Heritage through advanced AI and ICT methodologies and technologies. HerMeS' mission is, therefore, to enhance the fruition of cultural heritage through AI-generated touristic itineraries, tailored to users' needs and interests. The main outcome is an AI-enhanced smartphone application allowing different actors (i.e., tourists, citizens, economic operators, and public administration) to share experiences, feedback, and services.

The HerMeS app combines the socialization of the cultural experience with the development of AI technologies. This AI system takes into account users' interests (for instance, Nature, Archaeology, Eno-gastronomy, etc.) and needs (for example, visiting time, area of interest, special needs, etc.) to create a personalized tourist itinerary that maximizes the visitor's experience. Such a technology also offers useful information for defining targeted intervention strategies to the Public Administration and economic operators, which can lead to the development of solutions for territorial growth.

The touristic itinerary generation problem was examined from different perspectives and investigated through diverse AI techniques [3]. The most common solutions pursue a search and route generation approach (e.g., [14]) or recommendation techniques (e.g., [17, 9]) or machine learning (e.g., [4]) to maximize users' satisfaction of visiting a given destination. User preferences are usually considered limited to preferred points of interest and whole journey duration ([10]) or selecting a list of locations based on several criteria such as mandatory visits, tour duration, and endpoints of the tour. Similar to our approach, several works leverage AI-based solutions like, e.g., [1] to adapt touristic itineraries to users' preferences. Works generally formulate users' preferences and features of points of interest as mathematical problems solved through optimization algorithms capable of considering multiple (possibly conflicting) constraints and objectives (e.g., [11]). However, our work seems to pursue an original approach by leveraging an ontology-based knowledge base and contextual temporal planning to support thematic reasoning and personalized generation of cultural paths (i.e., users' itineraries). Furthermore, HerMeS adopted an interdisciplinary approach leveraging the partners' experience in pushing the innovation of knowledge, conservation, and fruition of Cultural Heritage.

The long-term objective of our work is to pave the way toward innovative applications that can transform cultural tourism, promote sustainable practices, and enrich individual experiences. In this regard, our research aims to contribute to the ongoing integration of technology and cultural heritage, highlighting the benefits and challenges of adopting smart solutions.

2 A Recommendation System for Cultural Itineraries

HerMeS provides a range of AI-based functionalities to enhance the enjoyment and exploration of Cultural Heritage, assisting stakeholders, especially tourists, in connecting their heterogeneous needs and interests. This is done through a bottom-up participatory model and advanced IT technologies, including AI algorithms, that analyze several variables (e.g., user preferences, historical data, and current trends) to generate personalized itineraries enriched with valuable

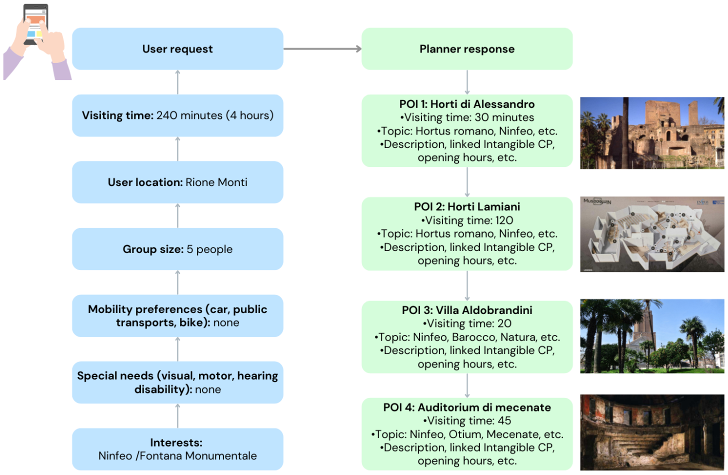

Fig. 1. Sample of HerMeS app flow

touristic information. The app recommends touristic itineraries, tailored to users' interests, that combine cultural sites (tangible cultural objects) with ephemeral experiences (intangible cultural objects) to promote off-road, unconventional, and hidden aspects of Lazio's cultural heritage.

To this end, an ICT infrastructure has been defined [5] to develop an app and a central database that serves the client apps installed on smartphones and provides access to its intelligent functionalities. The central database collects knowledge from different and highly heterogeneous sources, while a recommendation system selects a set of *cultural items* and proposes personalized itineraries based on user preferences. To support the recommendation system, a Knowledge Base (KB), based on ArCO [6], was designed to characterize a wide set of information concerning cultural heritage, such as geographical location, mobility information, type of cultural object (tangible vs. intangible), cultural object thematic descriptions, cultural object data properties (visiting time, inclusive accessibility, visiting hours, visiting price, etc.), relationships between cultural objects (a semantic relationship describing a close correlation between cultural objects), and correlations with cultural topics.

2.1 Extending the ArCO Ontology

HerMeS relies on the ArCO ontology framework, a network of 13 ontologies describing the domain of cultural heritage [4] [6, 7]. The key advantage of the ArCO ontology is its modularity, which supports flexible integration and usage within HerMeS. ArCO defines general concepts and properties that are suitable for interpreting pieces of knowledge and integrating existing thesauri e.g., PICO 4.1

---

[4] `http://wit.istc.cnr.it/arco`

[5]. However, while ArCO focuses on the representation of Cultural Objects, with a traditionally descriptive approach, what we wanted to achieve with HerMeS was the representation of Cultural Places - territorial structures with cultural significance, characterized by a stratification of tangible and intangible cultural objects. For this reason, HerMeS extends ArCO concepts to support a structured (and layered) description of a territory identifying parts (areas) that are relevant from a heritage perspective.

Among the main extensions to ArCO are: the introduction of new classes (Territorial Unit, Topographic Context, Monumental Unit, Monumental Complex, Cultural Property Description, etc.); the refinement of existing classes (Cultural Property Residual, Intangible Cultural Property, Topic); the introduction of new data properties (Visiting Time, Inclusive Accessibility, etc.), and; the integration of the PROV-O ontology [6] [15] to track the POI's provenance. The resulting formalism supports compositional descriptions of tangible and intangible cultural objects and contextual correlations with topics. The introduced taxonomy of topics supports a thematic indexing of cultural objects. We can thus easily retrieve cultural entities from general topics to more specific ones during the reasoning phase. Such a structure is crucial to flexibly personalize touristic itineraries according to users' interests, specified in terms of preferred topics (e.g., Figure 2).

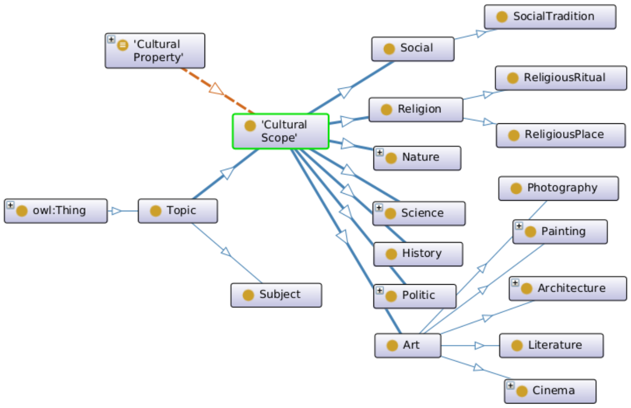

Fig. 2. Excerpt of the HerMeS taxonomy of topics.

HerMeS combines the semantic representation of cultural entities with automated planning to tailor user-specific itineraries. It requires synergetic reasoning on thematic cultural objects and related technical data of the visit like visit duration, previous visits, expected number of visitors, and potential congestion.

---

[5] `https://www.vocabularyserver.com/pico/it/index.php`
[6] `https://www.w3.org/TR/prov-o`

The implemented decision-making process addresses a numeric planning problem aiming to optimize user experience by balancing the number of cultural entities visited with the temporal constraints of the visit.

2.2 Dataset collection for HerMeS

A key element of the HerMeS project is to generate touristic itineraries that contain tangible cultural heritage (physical cultural places) and intangibles (festivities, traditions, proverbs, legends, etc.). With the collaboration of DigiLab Sapienza and ISPC teams, we managed to create a dataset of the tangible and intangible cultural heritage of two districts in the historical center of Rome (the Rioni Monti and Esquilino), which we used to populate our ontological framework. We obtained a knowledge graph of 100 cultural places: 76 tangibles (including 35 from Rione Esquilino and 41 from Rione Monti) and 24 intangibles. Figure 3 aggregates the modeled POIs by considering their geographic distribution over the territory. The intensity level characterizes the expected visit duration of the POIs (aggregated by geographic areas).

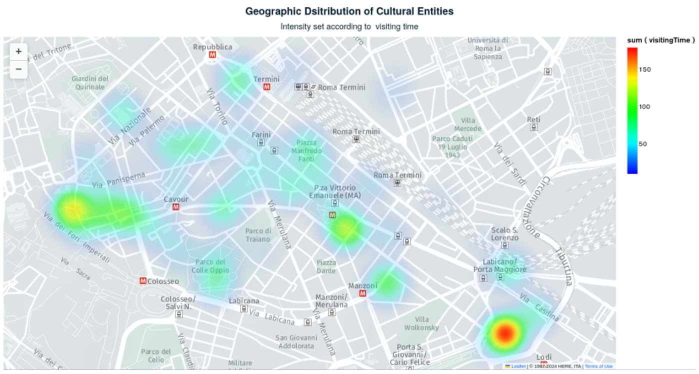

Fig. 3. Distribution of POIs considered for the Rione Monti and Esquilino in Rome.

3 AI-based Cultural Heritage Exploration

The implemented HerMeS services are the result of a pipeline of AI-based modules that retrieve and contextualize cultural heritage knowledge according to the requests and related preferences specified by the users. Figure 4 shows the functional flow implementing the personalization of users' itineraries. First, the user sends a trip request using the HerMeS app. A request encapsulates the user's interests (i.e., a list of preferred topics), and visit preferences (i.e., geographic area, duration of the visit, accessibility, and mobility requirements).

Each request is intercepted by HerMeS REST API and forwarded to the back end AI-based reasoning components. The semantic-based recommendation component retrieves information about the cultural objects that match the specified interests and preferences. It relies on HerMeS ontology to extract a contextualized view of cultural entities given by the thematic correlations with intangibles and compositional relationships with other tangibles.

The extracted set of tangibles is used to build and refine a travel dataset containing information about the expected travel distance between any pair of tangibles according to different mobility preferences (e.g., bus, metro, foot). Such a dataset is incrementally populated and refined by integrating third-party APIs (e.g., Distance Matrix API from Google) that provide reliable and updated mobility data. The travel-time dataset and the outcome of the semantic-based recommender system are input to the component that automatically generates the planning problem specification. The temporal planning component then synthesizes a *cultural path* by explicitly reasoning about temporal requirements (i.e., the total time available for the visit and the visiting time of each tangible) and the travel time of alternative sequences of visited tangibles. The resulting optimized (and personalized) cultural path is sent back to the user issuing the (synchronous) trip request as the response. The HerMeS app then interacts with the user by showing the planned tangibles and contextual associated information (i.e., correlated tangibles and intangibles).

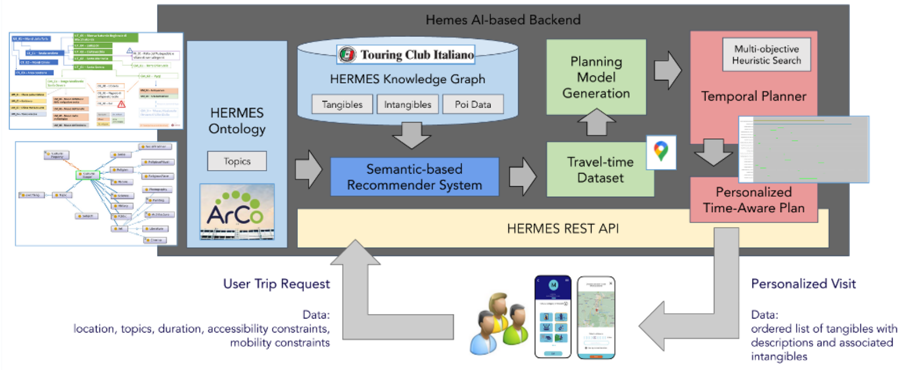

Fig. 4. The structure of the AI-based pipeline implemented by HerMeS back-end.

4 Personalizing Explorations through Planning

Integrated decision-making skills allow HerMeS to reason about temporal requirements and users' interests and synthesize consistent cultural paths. The developed temporal planner evaluates alternative sequences of cultural entities by considering the expected duration of the entire visit, the expected visit time of each cultural entity, and the expected travel time between pairs of entities. The optimal sequence is selected, encapsulated into a list of aggregated POIs, and returned to the user. This section introduces the temporal planning formalism used and then delves into the details of the algorithm developed to handle the *Visit Planning Problem*.

4.1 Reasoning on Time and Causality

The planning module of HerMeS relies on the timeline-based planning formalism introduced in [12]. Briefly, a timeline-based specification describes valid behaviors of domain features to be controlled over time. Given a description (i.e., a model), a timeline-based planning process synthesizes a set of flexible behaviors, i.e., timelines. The timelines describe how the modeled domain features should evolve to correctly realize the desired behaviors (i.e., the flexible sequences of states and actions each domain feature should respectively assume or perform).

More formally, a state variable is a tuple $SV = \langle V, T, D \rangle$ describing the set of valid behaviors of a domain feature:

- V is a set of values $v_i \in V$ representing states or actions the feature can assume or perform over time;
- $T : V \rightarrow 2^V$ is a state transition function describing possible successors on a timeline and thus valid transitions for each value $v \in V$;
- $D : V \rightarrow \mathbf{T} \times \mathbf{T}$ is a duration function specifying expected duration bounds, expressed in some temporal domain $\mathbf{T}$ (typically $\mathbf{N}^+$), for each value $v_i \in V$.

Temporal flexibility is crucial to deal with temporal uncertainty and robust execution of plans. A flexible timeline for a state variable sv_i is a sequence of (flexible) temporal intervals called tokens. Together these tokens describe an envelope of valid temporal behaviors of a domain feature. If $SV_i = \langle V, T, D \rangle$ is a state variable, a token x_j for the variable has the form $x_j = \langle v_k, [e_j, e_j^{'}], [d_j, d_j^{'}] \rangle$ where $v_k \in V$ is the value assumed by the token x_j, $[e_j,\ e_j^{'}]$ is the end-time interval of x_j (with $e < e'$) and $[d_j, d_j^{'}]$ is the minimum and maximum duration of x_j. The planned duration of a token should be consistent with the duration bounds of the associated value v_k. A timeline is a continuous sequence of tokens describing the behavior of a domain feature from a temporal origin to (at least) a desired planning horizon H. The start-time interval of a token is not explicitly represented since it coincides with the end-time interval of the previous token in the timeline i.e., the first token of a timeline starts the temporal origin [0, 0]. A timeline FTL_i for a state variable $SV_i = \langle V, T, D \rangle$ is a continuous and finite sequence of tokens of the form

$$x_1 = \left(v_1, [e_1, e_1^i], [d_1, d_1^{'}]\right),\ ...,\ x_m = \left(v_m, [e_m, e_m^i], [d_m, d_m^{'}]\right)$$

where v_1, .., $v_m \in V$ and for all J=1, ..., m-1, $v_{j+1} \in T\,(v_j)$. Denoting with start (x_j) the computed start time interval of a token x_j then, for all j = 1, ..., m-1, $[e_j, e_j^{'}] = \text{start}\,(x_{j+1})$.

State variables specify valid behaviors of domain features (i.e., local consistency). Complex behaviors however require the coordination of the behaviors of different state variables. Additional constraints are therefore necessary to coordinate simultaneous behaviors of state variables (i.e., enforce global consistency). *Synchronization rules* specify such constraints, necessary to synthesize valid plans (i.e., complex behaviors of state variable achieving desired goals). A synchronization rule has the form

$$a_0[SV_0 = v_0] \rightarrow \; a_1[SV_1 = v_1], ..., a_n[SV_n = v_n].\mathcal{C}$$

where every $a_i[SV_i = v_i]$ is a token variable denoting a temporal interval in which a state variable SV_i assumes the value v_i. The left-hand part of the synchronization rule ($a_0[SV_0 = v_0]$) is called the trigger. The set $\mathcal{C}$ specifies temporal relations between token variables. Synchronization rules with the same trigger are treated as disjunctions and represent alternative constraints that should hold between different sets of token variables.

The HerMeS planner has been developed as extension of the open-source framework PLATINUm [19, 18]. It integrates a novel search heuristics and solving procedure to generate cultural paths recursively. The next sections delve into the details of the modeled planning problem and the developed solving procedure.

4.2 Preference-Aware Visit Planning

The *Visit Planning Problem* consists of deciding the sequence of cultural entities that best fit users' interests and constraints among known tangible cultural properties. Planning choices are made among tangible cultural properties only. Correlated intangibles and tangibles are aggregated dynamically into the final POI structures sent back to the app.

Following the pipeline depicted in Figure 4, the temporal planning component receives dynamically generated problem specification as input and synthesizes an optimal temporal plan representing a personalized visit. The problem consists of building a timeline describing the personalized visit for a user. Planning decisions concern the incremental definition of the cultural path of a user (i.e., the user's timeline). Each incremental step selects the next tangible to insert into the visit according to the visit time of the next tangible and the travel time from the previous tangible in the timeline. The total duration of the planned visit should not exceed the input duration specified by the user. To achieve this, the developed heuristic search minimizes the visit's *coverage*. Figure 5 shows a conceptual view of the planning choices made during the iterative synthesis of a user timeline.

Planning choices concern the decision of the next tangible to add to the user timeline (*where to go next?*). Such choices are modeled through synchronization rules modeling alternative ways of instantiating a *Visit* token on the user timeline. It is worth underscoring that the number of such synchronization rules dynamically varies depending on the number of tangibles inferred by the knowledge-reasoning components. Therefore each visit choice has a branching

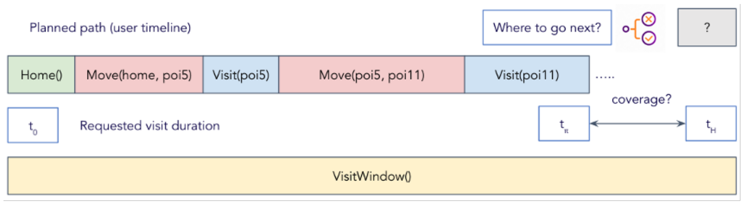

Fig. 5. Planning choices for deciding the cultural visit (i.e., user timeline).

factor equal to the number of tangibles that are considered relevant to user interests (i.e., user-selected topics). This modeling choice leads to a high branching factor of the search space. However, it is necessary to give the planner full flexibility in the synthesis of user cultural paths. A specifically designed heuristic supports these choices by evaluating alternative tangibles for visits (i.e., alternative tangibles that can be added to the timelines in the next iterations), and minimizing the differences between the requested time of the visit (i.e., the time horizon, t_H , of the planning problem) and the expected total time of the planned visit π. Interestingly, heuristic evaluation would consider alternative travel times between consecutive tangibles based on user preferences about mobility.

Algorithm 1 Domain independent iterative refinement of timelines.

```
function SOLVE($\mathcal{P}$, $\mathcal{S}$, $\mathcal{H}$)
    $\pi \leftarrow InitialPlan(\mathcal{P})$, $F_{pc} \leftarrow \emptyset$
    while $\neg IsSolution(\pi)$ do
        $\Phi^* = \{\phi_1^*, ..., \phi_m^*\} \leftarrow DetectFlaws(\pi, \mathcal{H})$
        for $\phi_i^* \in \Phi^*$ do                              ▷ Compute flaw solutions
            $N_{\phi_i^*} = \{n_1, ..., n_t\} \leftarrow HandleFlaw(\phi_i^*, \pi)$
            if $N_{\phi_i^*} = \emptyset$ then                     ▷ Unsolvable flaws
                $Backtrack(\pi, Dequeue(F_{pc}))$
            for $n_j \in N_{\phi_i^*}$ do                      ▷ Branching for each solution
                $F_{pc} \leftarrow Enqueue(n_j, \mathcal{S})$
        if $\neg IsEmpty(F_{pc})$ then                        ▷ Iterative refinement
            $\pi \leftarrow Refine(\pi, Dequeue(F_{pc}))$
        else
            return $Failure$                  ▷ No plan to explore and no solution found
    return $\pi$
```

More specifically, the solving procedure of a timeline-based planner iteratively refines a set of partially instantiated timelines. A dedicated data structure (the fringe) collects all the alternative partial plans that constitute the search space. At each iteration, the planner extracts the most promising partial plan from the fringe and analyzes the consistency of the timelines and the *goal conditions*. If the planner finds some flaws (i.e., conditions affecting the completeness and consistency of the timelines) the current plan is not a solution, and some operations should be made to solve the flaws and refine the plan. Each possible refinement (i.e., solution to a flaw) determines an alternative partial plan that is collected into the fringe. Algorithm 1 summarizes this general procedure implemented by PLATINUm [19, 18].

Depending on the characteristics of the domains, the solving procedure could be extended through dedicated heuristics. Heuristics are generally necessary to make better choices while expanding the search space and refining timelines. Such choices affect the quality of plans and the solving efficiency by analyzing qualities of partial plans to explore (i.e., possible solutions collected in the fringe) and discriminating among detected flaws of a refined plan.

In the *Visit Planning Problem*, a challenge concerns a missing *goal condition* suitable to identify a solution plan. The intrinsic recursive nature of the planning choices sketched in Figure 5 requires to “keep open” the possibility of adding a new step (i.e., tangible) to the visit. The planner cannot know how many tangibles a user can visit within the specified visit window. The total number depends on the duration of each visit and its schedule (i.e., the planned sequences of the tangibles) which determines different travel times. Therefore, the planning process should add tangibles to the visit incrementally until the partial plan meets a minimum quality condition (i.e., the *goal condition*). In this case, we set the *quality threshold* on coverage to 80% (i.e., at least 80% of the temporal window should be filled by the visit). Consequently, plans with no flaws but with the coverage below a certain threshold are discarded allowing the search to continue towards better solutions. Furthermore, we developed a new heuristic comparing partial plans based on their coverage. To synthesize reliable visits we encapsulated an evaluation criterion to polarize choices toward plans with less tangibles. Namely, we polarized planning choices towards plans that covered the visiting window with fewer steps.

Figure 6 depicts statistical data obtained within a solving instance. The *Fringe Size* clearly shows the recursive nature of the planning problem. The number of partial plans populating the fringe increases constantly which contrasts the typical behavior of seeing the fringe decreasing when planning choices move close to a solution. Despite this challenge, the planner addressed trip requests from users effectively and efficiently. A total number of 258 requests were issued during testing. The average response time of the planner was in the order of minutes (20 seconds at the lowest, 196 seconds at the highest) which was considered feasible for the application.

5 Conclusions

The HerMeS project allowed us to demonstrate the potential of a system that combines knowledge reasoning and planning in the context of cultural heritage. The knowledge graph enables us to represent complex data in a structured manner. Each point of interest (POI) is tagged with a series of properties and connected to other POIs in a network of semantic relations, allowing for inferential

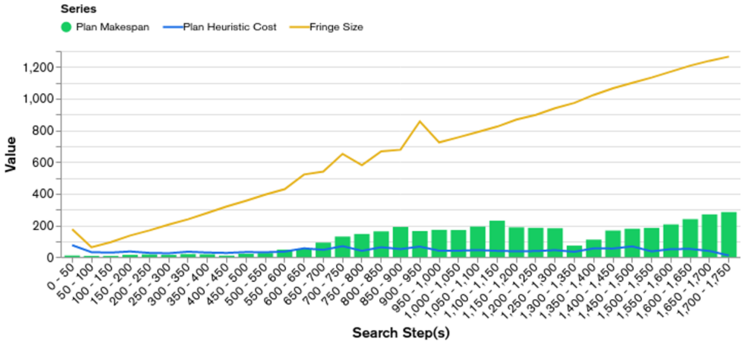

Fig. 6. Planning choices within the synthesis of a visit timeline.

reasoning. However, the digitization and tagging of cultural heritage datasets with tourism-related metadata remain open issues. In this initial phase of HerMeS, which represents a proof of concept, we opted for manual data collection to showcase the potential of this tool with high-quality data. However, this approach is not scalable if we decide to expand the geographical area of reference. Automating data collection and tagging is, therefore, a future research direction.

This proof of concept shows that the HerMeS system is highly flexible and can be integrated with other frameworks. HerMeS's planning algorithm has shown effective decisions concerning cultural paths. A central aspect was the effective reasoning on the time necessary for visiting tangibles and the expected traveling times between consecutive visits. The designed heuristic concretely evaluates alternative paths selecting the ones that achieve high quality. In this regard, planning technology supports flexible reasoning that can easily be tailored to different scenarios by integrating and evaluating different quality metrics of plans. Namely, the decision-making component could be extended to adapt planning choices to environmental contexts, and online knowledge about simultaneous visits planned on the territory. This latter aspect, combined with the capability of gathering real-time data from the environment (e.g., traffic, visiting queues, etc.) would strongly improve the *awareness* of planning choices.

Closing the loop between the planning process and the world state as in classical deliberative architectures [2] would strongly enhance the adaptability of planned paths and the experience of single users by better distributing the "cultural traffic". Future works would also focus on incorporating real-time data from other users' planned visits in the same area to better shape visits. The app is currently in its early stage of development. Usability tests to evaluate the effectiveness of the user interface will be conducted next.

Acknowledgments

The authors were partially supported by Regione Lazio and Lazio Innova within the HerMeS project (POR FESR LAZIO 2014-2020 Cod. A0375E0110).

References

1. Abbasi-Moud, Z., Vahdat-Nejad, H., Sadri, J.: Tourism recommendation system based on semantic clustering and sentiment analysis. Expert Systems with Applications **167** (2021)
2. Alami, R., Chatila, R., Fleury, S., Ghallab, M., Ingrand, F.: An architecture for autonomy. International Journal of Robotics Research, Special Issue on Integrated Architectures for Robot Control and Programming **17**(4), 315–337 (April 1998)
3. Aliano Filho, A., Morabito, R.: An effective approach for bi-objective multi-period touristic itinerary planning. Expert Systems with Applications **240** (2024)
4. Beraldi, P., De Maio, A., Olivito, F., Potrino, G., Straface, I., Violi, A.: A decision support system for trip tourism recommendation. International Journal of Transport Development and Integration **5** (2021)
5. Bucciero, A., Capaldi, D., Chirivi, A., Codella, M., Jaziri, M., Leopardi, L., Malatesta, S., Muci, I., Orlandini, A., Palombini, A., Pandurino, A., Panizzi, E., Umbrico, A.: Hermes: Heritage smart social media assistant. In: Extended Reality. Springer Nature Switzerland (2023)
6. Carriero, V.A., Gangemi, A., Mancinelli, M.L., Marinucci, L., Nuzzolese, A.G., Presutti, V.., Veninata, C.: Arco: The italian cultural heritage knowledge graph. In: Ghidini, C., Hartig, O., Maleshkova, M., Svátek, V.., Cruz, I., Hogan, A., Song, J., Lefrançois, M., Gandon, F. (eds.) "The Semantic Web – ISWC 2019". pp. 36–52. Springer International Publishing, Cham (2019)
7. Carriero, V.A., Gangemi, A., Mancinelli, M.L., Nuzzolese, A.G., Presutti, V., Veninata, C.: Pattern-based design applied to cultural heritage knowledge graphs. Semantic Web **12**, 313–357 (2021), 2
8. Ceccarelli, S., Cesta, A., Cortellessa, G., De Benedictis, R., Fracasso, F., Leopardi, L., Ligios, L., Lombardi, E., Malatesta, S.G., Oddi, A., Pagano, A., Palombini, A., Romagna, G., Sanzari, M., Schaerf, M.: Evaluating visitors' experience in museum: Comparing artificial intelligence and multi-partitioned analysis. Digital Applications in Archaeology and Cultural Heritage **33**, e00340 (2024)
9. Cesta, A., Cortellessa, G., De Benedictis, R., De Medio, C., Limongelli, C., Sciarrone, F., Tassarotti, G., Palombini, A.: Personalizing technology-enhanced learning for cultural visits. In: Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization. UMAP '20 Adjunct, Association for Computing Machinery (2020)
10. Chia-Ling, H., Wei-Lin, C., Chia-Ho, O.: Constructing a personalized travel itinerary recommender system with the internet of things. Wireless Networks (2023)
11. Choi, K.C., Sha, L., Chan-Tong, L., Angus, W., Philip, L., Benjamin, N., Ka-Meng, S.: Genetic algorithm for tourism route planning considering time constrains. International Journal of Engineering Trends and Technology **70**(3) (2022)
12. Cialdea Mayer, M., Orlandini, A., Umbrico, A.: Planning and execution with flexible timelines: a formal account. Acta Informatica **53**(6-8), 649–680 (2016)
13. Daga, E., Asprino, L., Damiano, R., Daquino, M., Agudo, B.D., Gangemi, A., Kuflik, T., Lieto, A., Maguire, M., Marras, A.M., Pandiani, D.M., Mulholland, P., Peroni, S., Pescarin, S., Wecker, A.: Integrating citizen experiences in cultural heritage archives: Requirements, state of the art, and challenges. J. Comput. Cult. Herit. **15**(1) (2022)
14. Expósito, A., Mancini, S., Brito, J., Moreno, J.A.: A fuzzy grasp for the tourist trip design with clustered pois. Expert Systems with Applications **127** (2019)

15. Lebo, T., Sahoo, S., McGuinness, D., Belhajjame, K., Cheney, J., Corsar, D., Garijo, D., Soiland-Reyes, S., Zednik, S., Zhao, J.: Prov-o: The prov ontology. W3C recommendation **30** (2013)
16. Silvestri, S., Tricomi, G., Bassolillo, S.R., De Benedictis, R., Ciampi, M.: An urban intelligence architecture for heterogeneous data and application integration, deployment and orchestration. Sensors **24**(7) (2024)
17. Tenemaza, M., Luján-Mora, S., De Antonio, A., Ramírez, J.: Improving itinerary recommendations for tourists through metaheuristic algorithms: An optimization proposal. IEEE Access **8** (2020)
18. Umbrico, A., Cesta, A., Cialdea Mayer, M., Orlandini, A.: Integrating resource management and timeline-based planning. In: The 28th International Conference on Automated Planning and Scheduling (ICAPS) (2018)
19. Umbrico, A., Cesta, A., Cialdea Mayer, M., Orlandini, A.: Platinum: A new framework for planning and acting. In: AI*IA 2017 Advances in Artificial Intelligence. pp. 498–512 (2017)