

Knowledge-based Personalization of Cultural Heritage Fruition through User Profiles: A proposal for the HERITAS Project

Saba Asrari^{1,*}

¹Ca' Foscari University of Venice – Department of Humanities, Dorsoduro 3246, Venice, 30123, Italy

Abstract

After democratizing Cultural Heritage fruition, the next step is personalizing it, to offer a more engaging experience to different users, especially those so far excluded from CH fruition. This is the objective of one of the tasks in HERITAS, a project funded by the Italian Ministry of Universities and Research. Given the complexity of both the cultural domain and the intended audiences, the Computer Science Unit of the University of Bari proposed the use of novel Knowledge Representation technologies to dynamically determine not only which contents should be presented to each user, but also how they should be organized. Building on this ground, this paper proposes to use the resulting personalized content organization to drive Virtual Reality and 3D front ends, providing tailored experiences for museums and exhibitions.

Keywords

Digital Humanities, Personalization, Knowledge Representation

1. Introduction & Motivations

Cultural Heritage (CH) has traditionally been considered as a non-essential or luxury good, reserved to an exclusive audience. More recently, it has started to be considered a public asset, whose fruition must be guaranteed to all citizens independently of their socio-economic status. Still, the perspective of fruition is individual (not aimed at groups of users) and standardized (the cultural objects are presented in the same way to all users, independently of their specificities). Workpackage 4 “HERITAGE4ALL (Human-centered Engagement and Research for Inclusive Technologies and Accessible Governance of Environments 4 Active Learning and Legacies): Connection, Interaction, Accessibility” of project HERITAS (Heritage Enhanced through Research, Innovation, Technology, and Accessible Solutions), funded by the Italian Ministry of Universities and Research and started in June 2026, aims at filling this gap, considering it as an important and powerful means of inclusion of specific categories of user so far excluded from CH fruition (such as migrants, older adults, children with special educational needs, and individuals who are neurodivergent or have psychiatric disorders).

HERITAS is dedicated to strengthening strategic research clusters, fostering collaboration among research entities, and developing skills for smart specialization, industrial transition, and entrepreneurship. WP 4 specifically addresses the issue of accessibility and interaction between communities and CH through an innovative and multidisciplinary approach, recognizing that full and informed access to cultural sites nowadays depends not only on physical factors but also on cognitive, social, and cultural dimensions. At the heart of all activities lies a systemic and innovative approach, centered on people and contexts, in line with the Faro Convention [1], based on the conviction that cultural landscapes are not mere backdrops, but inhabited spaces interwoven with memories, practices, and relationships. The goal is to build participatory learning environments by integrating a new concept of multidimensional accessibility—understood not only in a physical sense but also in cognitive, relational, and intercultural

4th Italian Workshop on Artificial Intelligence for Cultural Heritage (IAI4CH 2026, <https://ai4ch.di.unito.it/>), co-located with the 24th International Conference of the Italian Association for Artificial Intelligence (AIIA 2026). 6–9 October 2026, Perugia, Italy

*Corresponding author.

✉ 907931@stud.unive.it (S. Asrari)

ORCID 0009-0006-0185-195X (S. Asrari)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

terms. The activities will be carried out using a digital infrastructure that facilitates new forms of interaction among regions, people, and CH.

Among the objectives of WP4, this paper is especially concerned with:

- the design of personalized narrative tools and digital environments to foster mutual recognition between one's heritage of origin and adopted heritage, thereby strengthening the cultural inclusion of migrants;
- the development of a semantic knowledge base and digital applications (front-end and back-end) to enable accessible, personalized, and scalable engagement with cultural content for all, with special emphasis on interactivity and interoperability;
- the implementation of digital prototypes (serious games, virtual exhibitions, immersive environments, collaborative platforms, and training modules) with high potential for transferability to the cultural, educational, and social sectors, with the goal of promoting their adoption by museums, schools, nursing homes, cultural institutions, and businesses.

In particular, it will focus on the contribution of the UniBA Computer Science unit (UniBA-DIB) in HERITAS to support personalized fruition of museums and exhibitions. Tailoring the experience to specific users (or categories or groups of users) is essential for engaging them and better conveying the cultural message. We believe that personalization depends on a complex interplay between the user's background, perspectives, expectations, goals, and preferences. These facets must be related to many aspects of the cultural objects, including physical and intangible features, content, context, lifecycle and even usage. To effectively tackle such complexity, UniBA-DIB proposed the use of Knowledge Representation (KR) techniques, that can overcome the limitations of simple feature vectors and can ensure interoperability by design. Also, in our vision, personalization may act at several levels: description (focusing on different aspects of the objects to show), object (selecting different objects to show), organization (combining the objects differently to emphasize some relationships over others). So far, this last level has been approached by just proposing different visiting paths for a fixed environment. Leveraging the opportunities and flexibility that Digital Humanities (DH) technologies such as 3D and Virtual Reality can provide in this perspective, this paper proposes to automatically restructure the organization of a virtual museum or exhibition based on the KR personalized outcomes.

This paper is organized as follows. After presenting related work in the next section, we will describe our role and objectives in the project in Sections 3 and 4. Then, in Section 5 we will conclude the paper.

2. Background and Related Work

Cultural heritage is far more than a collection of static artifacts preserved behind glass; it is a living bridge for human connection, empathy, and cross-cultural inclusion. However, traditional museum exhibitions often struggle to connect with increasingly diverse global audiences. When visitors step into a gallery, they bring unique personal histories, cultural perspectives, and learning preferences. A single, rigid curatorial narrative can feel distant or inaccessible. To build spaces that truly welcome everyone, DH offers flexible, human-centered approaches that transform how cultural stories are shared. At the core of this transformation is blending artistic vision with computer science. Art provides the emotional resonance, spatial beauty, and semiotic depth, while computer science offers the structural framework needed to make these concepts dynamic [2].

At several institutional levels there is a commitment to support the development of CH-related tools. The Cultural Heritage Cloud¹ European Union initiative aims at creating a digital infrastructure to connect CH institutions and professionals across the EU. The requirements for the cloud infrastructure to support the broad spectrum of applications and uses of CH data, reported in [3], include full support to the relevant types of data, FAIR (Findable, Accessible, Interoperable, and Re-usable) content, and integration of suitable approaches and technologies to interact with CH data. The complexity of the

¹https://research-and-innovation.ec.europa.eu/research-area/social-sciences-and-humanities/cultural-heritage-and-cultural-and-creative-industries-ccis/cultural-heritage-cloud_en

field suggests to approach it at the knowledge level (inter-relation of information), rather than as just data. While most works on KR for CH deal with data schemas and vocabularies, standard ontological schemes like CIDOC-CRM [4, 5] serve as foundational reference models in this space, offering the semantic interoperability needed to structure CH knowledge bases and link them across domains. Other noteworthy initiatives about the development of ontologies, and associated knowledge graphs (KGs for short, knowledge bases organized according to graph models), also exist. E.g., the ArCo (Architecture of Knowledge)² initiative supported by the Italian Ministry of Culture is an ontology for, and a KG of, Italian CH. It currently reuses, and is aligned to standard CH-related schemes and ontologies: CIDOC-CRM, EDM, Cultural-ON, and OntoPiA. All these initiatives will be taken as a starting point for the development of the HERITAS ontology and KG, together with other resources such as Geonames, Wikidata, VIAF, etc. While powerful, knowledge-based solutions (based on the RDF graph model) departed from the industrial DB technologies (based on the LPG model), which are the reference for practical applications. Since the two approaches have complementary pros and cons, it would be desirable to fill this gap. Differently from most attempts in the literature skewed in favor of the RDF perspective, GraphBRAIN was the first framework to push for a native LPG-oriented approach [6], followed by other initiatives [7].

The proposed approach is closely related to knowledge-based and context-aware recommendation [8]. The latter considers the circumstances—location, time, available time, device, activity, social setting, etc.—in which an item is consumed [9]. The former, possibly in combination with the latter, exploits explicit knowledge about users, items, and domain, rather than relying primarily on ratings or behavioral similarity [10], and is particularly suitable when interaction data are sparse, items are complex, or recommendations must be explainable. In CH, they have been used to recommend artworks, points of interest, explanatory material, and personalized itineraries. [11] adapted only presentations and support to individual visitors, not the actual organization of the museum. [12] matched user profiles against semantically enriched museum collections and generated personalized museum tours, but just the recommendation, not the actual virtual visit. [13] integrated ontologies, mobile sensing, visitor profiles, and contextual constraints to support adaptive access to museums, archaeological sites, and cultural-tourism resources. Our approach does not just recommend items or tours; it aims at actually changing the museum structure and rendering it as a virtual experience. It places particular emphasis on using explicit, multidimensional user profiles—not only interaction histories—to select and adapt cultural content to users’ backgrounds, interests, goals, and presentation preferences.

3. The Knowledge Management Platform: GraphBRAIN

As said, providing tailored experiences to CH users is key for the effectiveness of cultural objects fruition. In turn, this requires connecting, combining and integrating information about many different aspects, which is what we call ‘knowledge’. Storing and handling knowledge, and ensuring interoperability among different systems and platforms, is the realm of the KR branch of Artificial Intelligence (AI), whose state of the art is represented by Semantic Web standards and technologies. However, some features of the Semantic Web approach are not fully satisfactory for application to the CH domain. E.g., they push for a Linked Open Data perspective that is partially incompatible with privacy preserving needs; their representations are not very intuitive and human-friendly; their storage and manipulation facilities are not in wide use in the practice; etc. [6] In past research we developed the GraphBRAIN framework exactly with the purpose of overcoming these drawbacks, while still pursuing the same objectives as the Semantic Web. So, we will adopt this framework in the project, and in this section we describe its main features and explain why they can be relevant to HERITAS.

GraphBRAIN is a general-purpose knowledge management system aimed at covering all stages and tasks in the lifecycle of a knowledge base, including knowledge acquisition, organization, and (personalized) fruition³. It is a proprietary technology, but its functions are embedded in an API which

²<http://wit.istc.cnr.it/arco/?lang=en>

³A demo of the system can be found at <http://digitalmind.di.uniba.it:8088/GraphBRAIN/>

will be released for free use, and exposed as REST services. In this way, any third-party application and organization will be able to exploit it. In particular, the tools developed in HERITAS (especially those described in the previous section) will embed it in order to make their behavior knowledge-aware and thus more flexible, powerful and personalized. From now on, by ‘GraphBRAIN user’ we will mean any entity using the framework, either single persons through interactive interfaces, or other applications through calls to the API or REST services. However, any access must be associated to a registered user, to track the interactions, but also to support privacy and personalization (as explained in the following).

GraphBRAIN joins a database management system for efficiently handling, mining and browsing the individuals, with an ontology level that allows it to carry out formal reasoning on the knowledge. It adopts the Labeled Property Graph (LPG) data model, where nodes (representing individuals) and arcs (representing relationships) may have labels (usually expressing their type) and associated attribute-value pairs. It uses state-of-the-art graph DBMSs [14] (it is currently compatible with the most famous solutions, Neo4j and FalkorDB), which are typically schema-less. This ensures great flexibility but does not allow to associate a clear semantics to the graph items, which is fundamental in KGs. GraphBRAIN wraps them, requiring all uses of the DB to be compliant with pre-specified data schemes, expressed in the form of ontologies (expressed in a proprietary format, specifically tailored for LPGs), and guarantees through the API that all accesses and uses of the KG are compliant with the ontologies.

Several ontologies can be applied to the same graph, representing different perspectives on the same knowledge, and combined with each other to connect knowledge across domains, extending in this way the scope of KR beyond the single perspectives. At the ontological level, the alignment and integration of the combined ontologies happens through their shared classes (those having the same name in the different ontologies); at the instance level, individuals act as bridges, allowing the users of a domain to reach information coming from other domains. This possibility is fundamental for HERITAS, where the aim is connecting knowledge from very different fields in a seamless way. Developing separately ontologies for the different aspects and domains, and then integrating them, will allow us to deal with the complexity and peculiarities of the field. GraphBRAIN provides a top-level ontology defining very general and highly reusable concepts and relationships (e.g., Person, Place; Person.wasIn.Place; etc.). This top-level ontology plays a crucial role to interconnect the domain-specific ontologies, ensuring an overall connected KG. A specific ontology will be developed for HERITAS, which will import the general ontology and possibly other existing ontologies of interest (e.g., we have available ontologies about libraries/archives/museums [15], education [16], etc.).

The knowledge base can be fed manually and collaboratively by users or automatically by knowledge extraction from documents and other kinds of resources (e.g., books or the Internet). Both ontologies and instances may be imported from, or exported to, the standard Semantic Web format OWL (Web Ontology Language) [17], in order to support their interoperability with existing resources and their reuse as linked open data (LOD) [18]. Instances may also have attachments, making GraphBRAIN a digital library, whose content is organized according to formal ontologies, fostering interoperability with other systems. Users may add, show, or delete attachments. Users may also add comments on entity or relationship instances (to provide suggestions or add information), or approve/disapprove them, their single attributes, and even the ontological items. This feedback is used to assign a trust value to the users and to improve personalization.

Users can search specific information in the KG with a predefined goal in mind, or browse it interactively (expanding or compressing node neighbors), letting the data drive the search and possibly finding relevant information in a serendipitous way. The ontology itself will be searchable knowledge, useful for understanding the domain. Users can also run on the available knowledge several advanced AI solutions based on Natural Language Processing (e.g., for expressing in natural language the formal knowledge extracted from the graph), Document Image Analysis (to automatically extract knowledge from existing documents), Information Retrieval (to find relevant information in the knowledge base), Knowledge Representation and Reasoning (to derive new knowledge from the available one in an explainable way), Network Analysis (to understand the relevance and possible connection of knowledge items), Graph/Data Mining (to discover relevant information in the available knowledge) and Machine Learning (to learn new knowledge from the knowledge base or profiles from user behaviors). The use

of qualitative approaches (not affected by frequency and explainable), in addition to quantitative ones, enforces diversity in the outcomes and trustworthiness of the system and of its outcomes.

If available, a user profile can be used to personalize the results of these algorithms, ensuring that each user obtains tailored information. This is a fundamental feature for HERITAS: satisfying the needs of many different users and kinds of stakeholders requires the ability to return different results to each of them. The user profile is based on the ontology; it can be provided from outside or built and maintained by the system, based on the user's interactions. A proposal for representing the user profile and using it for recommendation is described in [19].

While designed to support linked data and systems' interoperability, GraphBRAIN guarantees data protection, overcoming some limitations of both the collaborative approach and the LOD perspective. This is particularly relevant in CH, where some data may have restrictions related to Intellectual Property or Privacy. Also, some users might not want to disclose their sensitive information or part of the knowledge they provided. Privacy is strictly related to personalization, in that each user must decide what parameters can be exploited to build his model, and which functions can access which of these parameters. The administrators or the data owners may decide what information in the knowledge base, when, how and to whom can be disclosed, and again the API will filter user access to information based on the privacy settings. In particular, GraphBRAIN provides a very flexible approach to data visibility, by which the data owner may decide at several levels of granularity who is guaranteed access to each kind of information stored in the knowledge base. So, the GraphBRAIN KG is not available in its entirety as Linked Open Data. The portion meant to be public can be published as LOD, and transparently be perceived as such by standard Semantic Web technologies.

4. The DH Proposal

After describing the knowledge management framework that will drive personalized fruition of CH, we will now provide the details of the envisioned technical pipeline. Let us start from a general description of the System's Architecture and methodology:

- **Backend Architecture: Dynamic Semantic Querying.** Instead of displaying a single, static museum plaque, the KR engine operates as a dynamic, intelligent database backend. The system continuously queries and presents distinct cultural perspectives, historical migrations, and geometric layers based on what a specific visitor wants to explore. An artwork is no longer locked into one interpretation; computer science gives us the engine to hold these multi-layered stories, while artistic intent guides how those stories feel. By grounding technical architectures in artistic sensitivity, cultural heritage adapts to the human spirit rather than expecting visitors to conform to a rigid box.
- **Front-End Methodology: Canvas-to-Code Pipeline.** To bring abstract cultural data into the physical reality of an exhibition, front-end technologies act as the canvas where computer science meets artistic expression. Technologies like 3D modeling, Virtual Reality (VR), and spatial Digital Twins serve as visual extensions of human curiosity, letting visitors step inside an artistic vision [20]. The methodology bridges analog craft with digital precision through a clear pipeline:
 1. **Vector Grid Extraction:** Hand-drawn patterns or archival designs are deconstructed into vector-based node networks using parametric shape rules [21], preserving their authentic geometric proportions.
 2. **Volumetric 3D Modeling:** These vectors are imported into 3D modeling software like Blender to create volumetric spatial digital twins—living models that accurately replicate physical forms, surface textures, and light responses [2].

Special care will be devoted to user interaction mechanics. Through real-time 3D models, VR interfaces, and spatial digital twins built in Blender, visitors will have complete control over how they view the artwork:

- **Tactile & Spatial Exploration:** Visitors are no longer restricted by physical barriers. Through 3D interfaces and VR, they can freely rotate digital twin, adjust scales, and walk around complex architectural motifs, experiencing spatial rhythms from angles that match their personal curiosity [20].
- **Atmospheric & Environmental Lighting:** By simulating dynamic lighting conditions across surface textures, users can manually adjust environmental variables to watch how historic patterns respond to changing light—fostering a deeper appreciation for how light shapes visual perception [2].
- **Intuitive Interface Design:** Complex database queries and structural algorithms happen quietly in the background, keeping the front-end interface clean, interactive, and tailored to the individual viewer.

By wrapping technical computer science within thoughtful design, customized spatial digital twins turn museum visits into hands-on, memorable human experiences. To streamline production and reduce manual effort, the pipeline relies on semi-automated generation rather than fully manual 3D modeling. By converting vector motifs into 3D geometries using automated SVG import and parametric modifier stacks (such as Solidify and Array) in Blender, modular heritage patterns can be instantly generated and reused across different digital twins.

No two museum visitors experience art in the same way. Personalization is essential for truly engaging audiences and conveying cultural messages. When an exhibition tailors its story to match an individual’s background, curiosity, or artistic inclination, the cultural message resonates far more deeply. Merging the backend KR engine directly with the interactive 3D front-end enables real-time, customized visitor journeys:

- **Tailored Narrative Depth:** A design student interested in hands-on geometry can interactively tweak vector parameters to see how structural rules build a pattern [2, 21], while a historian views deep archival maps and scholarly documentation.
- **Co-Creating the Experience:** Customization transforms visitors from passive observers into active co-creators. Users interact with exhibition kiosks or VR setups to modify color palettes, ambient lighting, and geometric patterns in real time [2, 20], building a personal emotional connection to the artwork.
- **Adaptive Accessibility:** Personalization allows the gallery interface to automatically adapt—adjusting visual contrast, interaction speed, or narrative complexity to ensure every individual feels included.

The changes made by a user will affect only his own experience(s), and will not be visible by others. Co-creation happens in that, while these changes are stored in strict connection to the user, the most generally relevant of them will be selected (part manually, part automatically) and, after validation/moderation, integrated in the general experiences for all users.

To evaluate how effectively backend KR integrates with customized 3D spatial twin front-ends, a rigorous, mixed-methods evaluation methodology will be used. Measuring success will move beyond simple visitor counts; it will assess how technical customization drives meaningful learning outcomes, cross-cultural empathy, and sustained audience engagement.

The evaluation will operate across three interconnected layers:

- **Quantitative Interaction Telemetry:** The front-end 3D engine will log real-time user interaction metrics, tracking parameter adjustments (such as vector node tweaks or shader lighting shifts), dwell time per architectural motif, and navigation paths through the customized VR space.
- **Qualitative Emotional & Cognitive Assessment:** post-visit interviews, visual mapping, and reflection prompts, will help to capture the visitors’ sense of agency, cultural connection, and narrative comprehension.
- **Semantic Integrity Verification:** The underlying KR network will be tested for responsiveness, verifying how efficiently queries retrieve customized historical narratives without sacrificing data provenance.

The validation strategy and tools, especially concerning individuals with neurodivergent or psychiatric disorders, will be set up by the domain experts (psychologists and archaeologists) involved in the project.

5. Conclusions

Tailoring the Cultural Heritage fruition experience to different users, especially those so far excluded from CH fruition, is one of the objectives of the HERITAS project. The task is challenging due to the complexity of the field and the many aspects that must be jointly considered to obtain effective personalization. The Computer Science Unit of the University of Bari will develop Artificial Intelligence solutions based on the use of Knowledge Representation technologies to drive flexible and personalized fruition by Virtual Reality and 3D front-ends.

By combining customized front-end interactions with formal backend semantic modeling, the proposed methodology will deliver concrete theoretical and practical outcomes:

- **Deeper Engagement:** the aim is demonstrating that when visitors are given control over visual parameters—such as exploring lighting dynamics across a digital twin—their dwell time and narrative recall increase significantly compared to static exhibition displays [20].
- **Empowered Co-Creation:** the developed solutions will enable the transition of users from passive consumers to active participants, constructing personalized visual paths that connect historical patterns (such as Safavid or Bargello geometries) directly to their own cultural heritage [2].
- **Scalable Digital Humanities Framework:** the project outcomes will provide a reusable blueprint for museums and heritage sites, proving that AI (KR techniques) together with technical tools (parametric vector grids, and Blender PBR rendering) can be harnessed to serve inclusive, human-centered cultural missions.

Acknowledgments

The author would like to thank Stefano Ferilli for the useful discussions and for providing information about project HERITAS and the GraphBRAIN framework and platform.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

- [1] Council of Europe, Council of Europe Framework Convention on the Value of Cultural Heritage for Society, 2005. URL: <https://www.coe.int/en/web/conventions/full-list?module=treaty-detail&treatynum=199>, cETS No. 199.
- [2] S. Asrari, Integrating digitized cultural heritage into public and private space redesign using digital twins and traditional patterns to create immersive, intercultural environments that preserve identity while embracing innovation, Master's thesis, Ca' Foscari University of Venice, 2026.
- [3] European Commission: Directorate-General for Research and Innovation, P. Brunet, L. De Luca, E. Hyvönen, A. Joffres, et al., Report on a European collaborative cloud for cultural heritage – Ex-ante impact assessment, Publications Office of the European Union, 2022. URL: <https://data.europa.eu/doi/10.2777/64014>.
- [4] ISO/TC 46/SC 4, ISO 21127:2023 Information and documentation — A reference ontology for the interchange of cultural heritage information, 3rd ed., ISO, 2023. URL: <https://www.iso.org/committee/48798.html>.

- [5] M. Doerr, The CIDOC conceptual reference module: An ontological approach to semantic interoperability of metadata, *AI Magazine* 24 (2003) 75–92.
- [6] S. Ferilli, Integration strategy and tool between formal ontology and graph database technology, *Electronics* 10 (2021). URL: <https://www.mdpi.com/2079-9292/10/21/2616>.
- [7] R. Angles, A. Bonifati, S. Dumbrava, G. Fletcher, A. Green, J. Hidders, B. Li, L. Libkin, V. Marsault, W. Martens, F. Murlak, S. Plantikow, O. Savkovic, M. Schmidt, J. Sequeda, S. Staworko, D. Tomaszuk, H. Voigt, D. Vrgoc, M. Wu, D. Zivkovic, PG-Schema: Schemas for property graphs, *Proc. ACM Manag. Data* 1 (2023).
- [8] G. Adomavicius, A. Tuzhilin, Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions, *IEEE Transactions on Knowledge and Data Engineering* 17 (2005) 734–749. doi:10.1109/TKDE.2005.99.
- [9] G. Adomavicius, A. Tuzhilin, Context-aware recommender systems, in: F. Ricci, L. Rokach, B. Shapira, P. B. Kantor (Eds.), *Recommender Systems Handbook*, Springer, Boston, 2011, pp. 217–253. doi:10.1007/978-0-387-85820-3_7.
- [10] A. Felfernig, R. Burke, Constraint-based recommender systems: Technologies and research issues, in: *Proceedings of the 10th International Conference on Electronic Commerce*, ACM, New York, 2008, pp. 1–10. doi:10.1145/1409540.1409544.
- [11] O. Stock, M. Zancanaro, P. Busetta, C. Callaway, A. Krüger, M. Kruppa, T. Kuflik, E. Not, C. Rocchi, Adaptive, intelligent presentation of information for the museum visitor in PEACH, *User Model. User-Adapt. Interact.* 17 (2007) 257–304. doi:10.1007/s11257-007-9029-6.
- [12] Y. Wang, N. Stash, L. Aroyo, P. Gorgels, L. Rutledge, G. Schreiber, Recommendations based on semantically enriched museum collections, *J. Web Semant.* 6 (2008) 283–290. doi:10.1016/j.websem.2008.09.002.
- [13] M. Casillo, F. Colace, D. Conte, M. Lombardi, D. Santaniello, C. Valentino, Context-aware recommender systems and cultural heritage: A survey, *J. Ambient Intell. Humaniz. Comput.* 14 (2023) 3109–3127. doi:10.1007/s12652-021-03438-9.
- [14] I. Robinson, J. Webber, E. Eifrem, *Graph Databases*, 2nd ed., O’Reilly Media, 2015.
- [15] S. Ferilli, E. Bernasconi, D. Di Pierro, D. Redavid, A holistic ontology for digital libraries – the IFLA-compliant core, in: *Proceedings of the 2nd Workshop on Artificial Intelligence for Cultural Heritage (IAI4CH@AIIA2023)*, co-located with the XXII International Conference of the Italian Association for Artificial Intelligence (AI*IA 2023), volume 3536 of *CEUR Workshop Proceedings*, 2023, pp. 39–52.
- [16] S. Ferilli, E. Bernasconi, D. Di Pierro, D. Redavid, L. Loop, A holistic ontology for digital libraries – the educational extension, in: *Proceedings of the XX Conference on Information and Research science Connecting to Digital and Library science (IRCDL 2024)*, volume 3643 of *CEUR Workshop Proceedings*, 2024, pp. 74–86.
- [17] W3C, OWL 2 Web Ontology Language Document Overview, 2nd ed., W3C, 2012. URL: <https://www.w3.org/TR/owl2-overview/>.
- [18] T. Heath, C. Bizer, *Linked Data: Evolving the Web into a Global Data Space*, Synthesis Lectures on the Semantic Web, Morgan & Claypool Publishers, 2011.
- [19] S. Ferilli, Knowledge graph-based user models and personalized access for cultural heritage, in: *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization (UMAP Adjunct ’25)*, ACM, New York, 2025. doi:10.1145/3708319.3733654, 5 pages.
- [20] M. Carrozzino, M. Bergamasco, Beyond virtual museums: Experiencing immersive virtual reality in real museums, *Journal of Cultural Heritage* 11 (2010) 452–458.
- [21] G. Stiny, J. Gips, Shape grammars and the generative specification of painting and sculpture, in: *Information Processing 71: Proceedings of IFIP Congress 1971*, volume 71, North-Holland Publishing Company, 1972, pp. 1460–1465.