Interpretable Clusters for Representing Citizens' Sense of Belonging through Interaction with Cultural Heritage

GUILLERMO JIMÉNEZ-DÍAZ, Facultad de Informática, Universidad Complutense de Madrid, Spain
BELEN DIAZ-AGUDO, Facultad de Informática, Universidad Complutense de Madrid, Spain
LUIS EMILIO BRUNI, Department of Architecture, Design and Media Technology, Aalborg University, Denmark
NELE KADASTIK, Department of Architecture, Design and Media Technology, Aalborg University, Denmark
ANNA FOLLO, Fondazione Torino Musei, Italy
ROSSANA DAMIANO, Dipartimento di Informatica, Università di Torino, Italy
MANUEL STRIANI, Dipartimento di Informatica, Università di Torino and Computer Science Instituite - Università del Piemonte Orientale, Alessandria, Italy
ANGEL SANCHEZ-MARTIN, Facultad de Informática, Universidad Complutense de Madrid, Spain
ANTONIO LIETO, Università di Torino and Cognition Interaction and Intelligent Technologies Lab, DISPC, Università di Salerno, Italy

The EU H2020 project SPICE (Social cohesion, Participation, and Inclusion through Cultural Engagement) focuses on developing, designing, and implementing new methods and digital tools for citizen curation. This paper delineates several software tools developed within the project, presenting innovative approaches to represent and visualize citizens and communities resulting from their engagement with cultural heritage. Aligned with the central tenets of SPICE –particularly the notions of belonging and the Interpretation Reflection loop– the primary objective is to bolster citizens' participation and inclusion in fostering social cohesion. This paper describes how the SPICE tools can be utilized to guide the processes of interpretation and reflection on cultural heritage artefacts. The Community Model serves as a pivotal component, enabling the modeling of citizens and communities through the utilization of similarity functions for clustering citizens based on perspectives. The clustering algorithm is intricately crafted to generate coherent communities, iterating until all clusters are interpretable using demographic attributes, centroid-based representations, and similarity attributes. Authors posit that this model holds significant value in comprehending and structuring complex data within cultural heritage contexts.

Authors' Contact Information: Guillermo Jiménez-Díaz, gjimenez@ucm.es, Facultad de Informática, Universidad Complutense de Madrid, Spain; Belen Diaz-Agudo, belend@ucm.es, Facultad de Informática, Universidad Complutense de Madrid, Spain; Luis Emilio Bruni, leb@create.aau.dk, Department of Architecture, Design and Media Technology, Aalborg University, Denmark; Nele Kadastik, neka@create.aau.dk, Department of Architecture, Design and Media Technology, Aalborg University, Denmark; Anna Follo, anna.follo@fondazionetorinomusei.it, Fondazione Torino Musei, Italy; Rossana Damiano, rossana.damiano@unito.it, Dipartimento di Informatica, Università di Torino, Italy; Manuel Striani, manuel.striani@uniupo.it, Dipartimento di Informatica, Università di Torino and Computer Science Instituite - Università del Piemonte Orientale, Alessandria, Italy; Angel Sanchez-Martin, josanc16@ucm.es, Facultad de Informática, Universidad Complutense de Madrid, Spain; Antonio Lieto, alieto@unisa.it, Università di Torino and Cognition Interaction and Intelligent Technologies Lab, DISPC, Università di Salerno, Italy.

ACM 1556-4711/2024/5-ART
https://doi.org/10.1145/3665142

To exemplify our approach, the paper examines different attributes of individual citizens and citizen groups in the GAM (Galleria Civica d'Arte Moderna e Contemporanea) case study. Here, perspectives are delineated based on visitors' demographic attributes and their emotional responses when engaging with artworks. These perspectives are then visualized using the VISIR tool, facilitating the exploration and revelation of connections between citizens and communities, thereby bridging the realms of citizen space and cultural heritage space.

CCS Concepts: • **Human-centered computing → Visualization systems and tools**; *Collaborative and social computing*; • **Information systems →** *Collaborative and social computing systems and tools*; • **General and reference → Surveys and overviews**; • **Social and professional topics → User characteristics**.

Additional Key Words and Phrases: cultural heritage; citizen curation; narrative identity; heterarchy of values; clustering; community modelling; affective computing; sense of belonging; social cohesion; museum interaction

1 INTRODUCTION

The EU H2020 project SPICE (Social cohesion, Participation, and Inclusion through Cultural Engagement) focuses on developing, designing, and implementing new methods and digital tools for citizen curation. Through the development, design, and implementation of new methods and digital tools in the cultural heritage domain, SPICE focused, among other aspects, on suggesting innovative approaches for representing and visualizing citizens and communities, which resulted from citizens' engagements with cultural heritage. The research conducted in SPICE revolved around an understanding of *citizen curation* as "citizens applying curatorial methods to archival materials available in memory institutions in order to develop their own interpretations, share their own perspective and appreciate the perspectives of others" [5]. In this direction, citizen curation in SPICE refers to a process through which citizens are encouraged to develop and share their personal interpretations and reflections on cultural heritage objects and digitized artworks.

Previous work in this area includes, among others, EU-funded projects such as PLUGGY[1] (Pluggable Social Platform for Heritage Awareness and Participation), CHESS[2] (Cultural Heritage Experiences through Socio-personal interactions and Storytelling), EMOTIVE[3] (Emotive virtual cultural experiences through personalized storytelling), and GIFT[4] (Meaningful Personalization of Hybrid Virtual Museum Experiences Through Gifting and Appropriation).

PLUGGY (2016-2019), a European research and innovation action, introduced a novel networking platform for cultural heritage with the goal of actively involving European citizens in various cultural heritage activities. The primary focus was on creating interactive technologies to support crowdsourced curatorial processes related to cultural heritage [20]. The project's rationale was based on the observation that, although social platforms such as Facebook and Instagram can support the building of networks, they have not been fully explored for the promotion of cultural heritage nor for fostering the creation of long-lasting heritage communities. PLUGGY aimed to bridge this gap by developing necessary Information Communication Technology (ICT) tools for empowering end-users to create and share their stories and local knowledge about cultural heritage.

Similarly, addressing the limitations of conventional social media platforms in facilitating intimate sharing of museum experiences among visitors, the GIFT project (2016-2019) developed a suite of tools, frameworks, and design guidelines to assist museums in using digital technology for more meaningful and engaging visitor experiences. For instance, the tool Artcodes enables the creation of visual, scannable markers that users can design themselves. These markers can then be connected to digital assets, shared with others, and allow for the exploration and adaptation of others' experiences.

[1] https://www.pluggy-project.eu
[2] http://www.chessexperience.eu/
[3] https://emotiveproject.eu
[4] https://gifting.digital/

Furthermore, the CHESS project (2011-2014) explored the potential of new technologies and interactive digital storytelling (IDS) to create more personalized on-site visitor experiences. Its aim was to enable cultural heritage institutions to enhance engagement and accessibility of their collections through innovative cultural interactive experiences. With a specific focus on younger audiences, the goal was to provide experiences that could connect with visitors' interests and personal narratives. The CHESS project also laid the groundwork for a subsequent EU Research and Innovation (RIA) effort called EMOTIVE (2016-2019) which delved into methods and tools to support cultural and creative industries in developing Virtual Museums, emphasizing 'emotive storytelling' and exploring XR technologies.

The mentioned projects not only demonstrate the potential of digital technologies in supporting heritage institutions to make their collections more accessible and engaging, but also emphasize the crucial role of citizens in contributing to curatorial processes through their personal experiences and inputs. Thus, by encouraging such active participation from visitors, these initiatives not only shape the narratives around cultural artifacts but also promote museums and heritage institutions as dynamic platforms that can reflect diverse perspectives, fostering inclusion and social cohesion.

In SPICE, the motivation for involving the public in the curatorial process and encouraging them to share their personal interpretations of cultural artifacts and museum collections was to cultivate social cohesion and inclusion through a more participatory approach. This not only seeks to afford visitors an opportunity to actively engage in the museum experience, but also contributes to the exchange of diverse perspectives within and across citizen groups, fostering an enriched comprehension of cultural heritage.

The overall framework guiding the SPICE processes is supported by the continuous Interpretation Reflection Loop (IRL). The IRL serves as a model for linking the diverse interpretation and reflection activities presented by the system, embedded and distributed across various temporal phases and components of the SPICE digital platform. For example, the SPICE platform presents users with a range of activities, such as selecting artefacts, tagging, and sharing personal stories and opinions. These contributions are then analyzed by the SPICE technical tools to support reflection, for instance, by presenting users with similar or alternative perspectives, providing recommendations for new types of content, and/or suggesting additional activities [14]. In this way, citizen curation is regarded as a tool for enabling social cohesion and aims to encompass a wide range of audiences and use-cases in order to "promote inclusive participation and social cohesion in a variety of contexts" [5].

The SPICE project involved five case studies: the Irish Museum of Modern Art (IMMA) in Ireland, Design Museum Helsinki (DMH) in Finland, the Turin Civic Gallery of Modern and Contemporary Art (GAM) in Italy, the National Museum of Natural Sciences (MNCN) in Spain, and the Hecht Museum in Haifa, Israel. Within the project, each case was working with specific groups of citizens and end-user communities, including senior citizens, asylum seekers, communities of d/Deaf people, students, and different religious groups, among others. In this direction, the focus of this paper is to describe and illustrate how SPICE technologies have been used by the project's case studies to detect and explore the communities emerging from the citizens' input which are integral to the Interpretation-Reflection Loop.

These communities can, for instance, be used for the exploration of user interpretations of museum artefacts, for finding contents of interest, for reflective reasoning, and for exploring different dimensions of social cohesion. A critical aspect of the logic is carried out by the Community Model and the visualization tool called VISIR. These components are responsible for discovering new citizen communities and highlighting meaningful representations of citizens and citizen groups, which in turn, can reveal novel and surprising connections to "others". In this paper, the use of these tools is exemplified using one of the five SPICE case studies, namely the case of the Civic Gallery of Modern and Contemporary Art (Galleria Civica d'Arte Moderna e Contemporanea, GAM) in Turin, Italy. The case study, carried out in collaboration with the Turin Institute for the Deaf, revolves around the interpretation of artworks through emotional responses, with a specific focus on the d/Deaf community (see [19] for a detailed description of the case study).

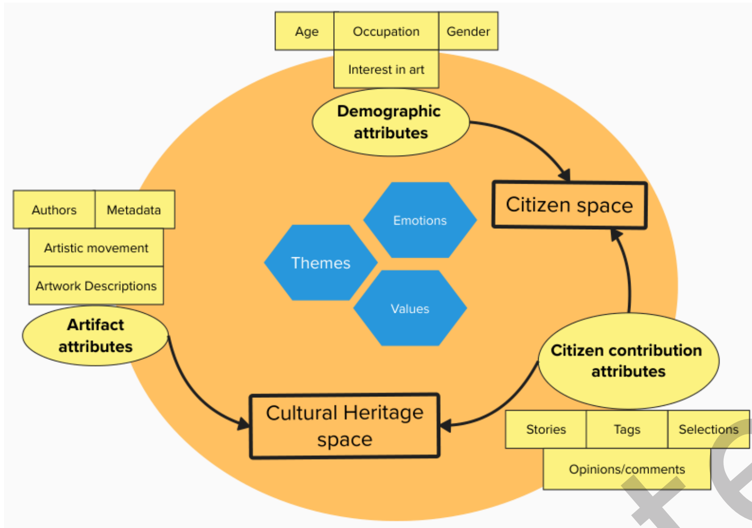

Fig. 1. Citizen space and cultural heritage space. Two mutually constraining but disjoint feature spaces, with citizen curation bringing both spaces together.

Section 2 describes the cultural heritage interpretation and reflection model underlying the SPICE project (called Interpretation Reflection Loop), summarizing the data workflow through the project infrastructure. In Section 3, we describe the technologies, algorithms, and tools used to detect and explore emerging communities and how the interpretation of these communities plays a key role in supporting the exchange of perspectives within and across citizen groups. Section 4 introduces the GAM case study and offers a description of its implementation including the GAM workflow architecture as well as the considered attributes and perspectives employed in the GAM case. In Section 4, we provide an example of interpretation and reflection within the GAM case study. Here, we provide the procedural details and present the findings derived from the analysis conducted with the VISIR tool. Lastly, in Section 5 we discuss the process, outcomes, and implications, and outline potential avenues for future research and development.

2 THE INTERPRETATION REFLECTION LOOP

The Interpretation Reflection Loop (IRL) serves as a model for guiding the embedded processes and activities of the SPICE digital platform that instantiates different stages of interpretation and reflection. The IRL framework is rooted in the idea of *citizen curation*, a process through which citizens are encouraged to develop and share (employing an integrated technological infrastructure [9]) their personal interpretations and reflections on cultural heritage objects and digitized artworks [5].

Citizen curation in SPICE aims to bridge the gap between what we characterize as *the cultural heritage space* and *the citizen space*. On one hand, the cultural heritage space includes the existing data on the cultural artefacts and the museums or cultural sites themselves (cultural artifact information, metadata, artists, artwork descriptions, etc.) and all the opportunities for accessing this heritage. On the other hand, the citizen space comprises both the data about the citizens as well as their curatorial contributions developed in the IRL process (see Figure 1). The model allows us to consider the many-to-many relationships that express the intertwining of the different attributes present in both the citizen space and the cultural heritage space. Through the case studies' user journeys, the two initially disjointed spaces enter into a fruitful relationship, putting in relation three types of data: citizens'

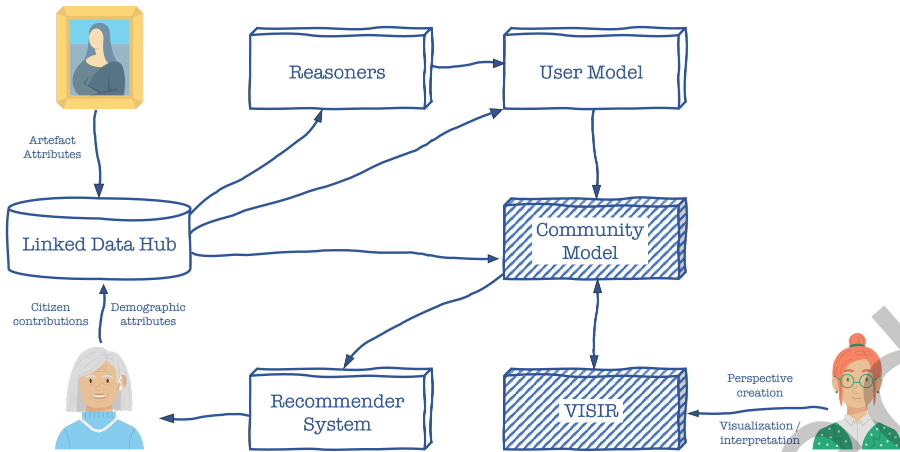

Fig. 2. SPICE workflow that supports the Interpretation Reflection Loop. Shading modules (Community Model and VISIR) will be described in next subsections.

demographics, citizens' curatorial contributions (i.e., selections, tags, narratives, emotional profiles, etc.) and cultural artefacts.

During the *Interpretation*, the SPICE platform presents users with a range of activities, such as selecting and tagging cultural artefacts, writing personal stories or opinions about them. The activities are presented to users through suitable interfaces that support the citizens in contributing and sharing meaningful interpretations of the artefacts they encounter. The resulting citizen contributions may include tags with evoked emotions or themes, personal stories, opinions or comments, inspired by the artefacts or artworks. Additionally, citizens may be invited to provide supplementary personal information such as age and gender.

The citizen contributions are then analysed by the SPICE technical tools (e.g. reasoners, semantic tools) to support *Reflection*. While the precise application of these tools will vary depending on the specifics of the case study implementation, this could encompass elements such as presenting users with analogous or alternate perspectives on cultural artefacts, recommendations for new content, and/or proposing additional activities [14]. In this context, the Community Model (see Section 3.2) analyzes and clusters the data to find the communities mingling in these two spaces, i.e. the citizen space and the cultural heritage space (Figure 1).

2.1 Technologies and Data Workflow supporting IRL

The objectives behind the SPICE digital platform are twofold. First, it sought to provide curators and researchers with a social laboratory for cultural analysis to map the ongoing cultural process. Secondly, the goal was to offer the participating citizens dynamic representations of themselves, allowing them to recognize themselves about the perspectives of others, thus fostering a greater sense of belonging. From a technical point of view, this platform was developed as a data infrastructure, built according to the requirements detailed in [9], to support the acquisition and management of dynamic data from a variety of sources including museum collection metadata and digital assets, social media events and user activities, systems' activities (e.g., recommendations, reasoning outputs), ontologies and linked data produced by pilot case studies.

We summarize the data workflow (sketched in Figure 2) through the SPICE infrastructure:

(1) Museum curators gather and store information about the museum artefacts, like the author, artistic movement, content depicted in an artwork, etc. (*Artefact Attributes*).

(2) Citizens interact with museum artefacts creating stories, tagging the artwork, commenting the evoked emotions, etc. (*Citizen Contribution Attributes*) through an app or a web page. Citizen can also be requested to provide other personal data (age, gender, etc. –*Demographic attributes*).
(3) Artefact, Demographic and Citizen Contribution data are stored into the SPICE **Linked Data Hub** (LDH) [2]. A layout of the system and its key components[5] is provided in Figure 3. This back-end component worked as infrastructure integrating, via API, a number of different applications and systems adopted for the IRL.
(4) **Reasoners** mine citizen contributions for emotions, values and iconographical subjects using multiple services in parallel (like DEGARI for emotions [18], eMFD API[6] for values [13] and ICONCLASS API[7] for artefacts, among others).
(5) **User Model** stores the enriched contributions (with added triples describing annotations about emotions, values and subjects) in JSON-LD format. It also combines this data with the citizen demographics stored in the LDH (UM-data).
(6) User Model sends UM-data to the **Community Model**, which is in charge of running clustering algorithms to discover implicit communities using different perspectives defined by curators. Thereafter, the clustering algorithms (with a particular similarity function) generate the resulting communities in an explicable way.
(7) Community Model sends clusterized and interpretable data to **VISIR** which is able to support the Interpretation-Reflection-Loop (IRL) by executing the perspectives visualisation, communities and users' introspection and inter-intra community visualisation. VISIR is also employed by the curators to create and configure new perspectives for discovering new citizen communities.
(8) Community-based and emotion-based **recommendations** of stories and artefacts are provided by the **Recommender System** for engaging citizens to reflect on their sense of belonging.

Further details of this technological framework are provided in [2, 6]. This paper focuses on the discovery and interpretation of the implicit communities created after analyzing citizen contributions, supported by the Community Model (section 3.2) and VISIR (section 3.3).

3 DISCOVERING AND INTERPRETING COMMUNITIES

In SPICE, citizen communities encompass citizens who share a common identity and a sense of belonging. Citizens can belong to one or many *explicit communities*, also called interest or target groups –a static *a priori* asserted category used as one of the inputs describing user features. Explicit communities line up with the museum's interests, identified personas or profiles that have been formalised in the project ontology, and represent user archetypes. Each SPICE case study focuses on specific groups of interest, e.g. students from a certain school, teachers, members of an association, senior citizens, asylum seekers, children with serious illnesses, d/Deaf people,

[5] The Linked Data Hub (LDH) Portal is made up of a number of components, the most visible being the front-end web portal. The portal provides a facility for data providers to create and catalogue datasets, manage data and API access as well as manage their privacy, licences and provenance. The web portal enables users to browse catalogues and registries of datasets and their corresponding metadata and data models. Sitting behind the web portal and driving most of its functionality is the LDH API that exposes a range of REST-based API functionality to support the production, management and consumption of data. This API is directly available to all, enabling developers to side-step the web portal and integrate SPICE LDH read/write operations with existing automated systems. The API also offers a range of extended functionality over and above that offered by the web portal such as: full read/write (CRUD) access to datasets for users with appropriate permissions; enhanced browsing capabilities; advanced querying and data filters (using both JSON-style queries and SPARQL); a read-only SPARQL endpoint. Datasets take the form predominantly of JSON documents or static files, so anything that can be encoded as a JSON string can be stored and accessed using the SPICE LDH's full range of features. All JSON documents pushed into the SPICE LDH via the API are also replicated as RDF to a graph database for read-only query access via SPARQL
[6] https://github.com/medianeuroscience/emfd
[7] https://iconclass.org/help/api

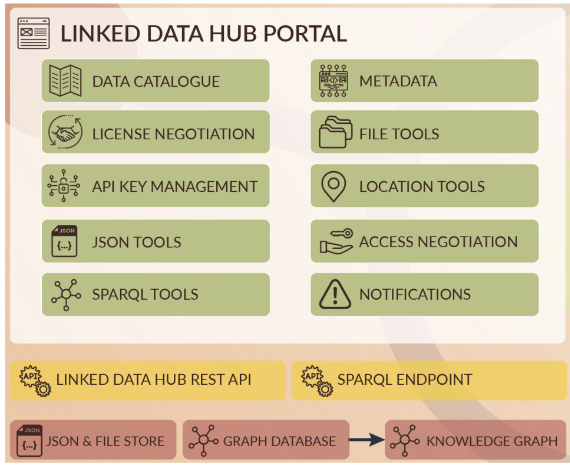

Fig. 3. Layout of the SPICE Platform integrating different systems and functionalities

and people from different religious and secular communities, among others. Most of these explicit communities are identified using *Demographic attributes*.

Besides, citizens may be inferred to belong to different *implicit communities*, i.e. emerging communities which are based on citizen attributes that are extracted from citizens' contributions when they interact with cultural artefacts, e.g. emotional responses, comments, and stories (see Figure 1).

In the IRL, the Community Model and the VISIR visualization tool, designed and implemented in the SPICE project, play a key role in supporting the exchange of perspectives within and across citizen groups by suggesting "new" unexpected groups that can arise from interactions with cultural artefacts. In the next section, we describe the Community Model and the VISIR tool which allow for the perspective configuration and the juxtaposing of two perspectives simultaneously together with the supporting explanations.

3.1 Previous work

One of the key aspects of the SPICE technologies supporting the Interpretation-Reflection loop (IRL) is the citizen clustering and community detection. Clustering techniques have been employed in cultural heritage for visitor segmentation to identify new market opportunities for museums [30] or to organize groups of visitors to preserve the cultural environment [26], among others. These works mostly employ the data in the citizen space but are restricted to the demographic attributes, obtained with questionnaires and surveys.

Other works have used clustering techniques to identify visitor behaviours in museums, such as involving which paths are most commonly followed by visitors or how the visitors are normally distributed, using proximity sensors [22]. In [22], focus was on identifying which museum artefacts were visited (i.e., when a visitor spent a minimum period of time close to the artefact) and the duration of the visitors' contemplation. In this context, both can be considered as citizen contribution attributes as they reflect the visitors' interest in the artwork, however, they also lack deeper attributes that our approaches employ, such as emotions and values extracted from citizen opinions and comments.

Social networks are vast platforms where people freely express their opinions on a wide range of topics and themes, including cultural heritage. For this reason, some works have employed community detection algorithms specifically implemented for social network analysis to gain insight about citizen interests about cultural heritage.

The work in [12] combines topic modelling with a community detection algorithm to identify communities of interest around Greek cultural heritage on Twitter. A similar dataset has been employed in [16] to test the performance of several community detection algorithms for finding communities of users with similar features. However, these features are considered in terms of Twitter metrics, instead of opinions or preferences, like in the approach described in this paper. Additionally, the authors have not found any work that combines both cultural and citizen space to discover new citizen communities.

Finally, although there is a large number of tools for visualizing clusters and communities, the majority of them are for general purposes, lacking the domain-specific explanations needed to interpret the formed clusters, which we consider a key aspect in the IRL process. Additionally, none of the revised works provide the capabilities to compare different perspectives–different clustering or community detection algorithms, using different parameters. In SPICE, the ability to compare different perspectives can be viewed as central in the process of reflection. This is particularly relevant when visitors and curators seek to reflect on a citizen's identity based on the discovered groups the citizen belongs to.

3.2 Community Model: Discovery and explanations

The Community Model provides the technological infrastructure[8] that enables to model citizens and communities based on citizens' characteristics, opinions, and preferences. Each community represents a group of citizens that are related to each other according to a given similarity criteria, but who are also sparsely connected to other groups. The Community Model analyses and clusters the data to find communities combining two spaces (see Figure 1): the citizen space (which includes both the data about the citizens and their curatorial contributions in the process) and the cultural heritage space (which includes the existing data on the cultural artefacts and the museums or cultural sites themselves).

Implicit communities are computed by clustering algorithms, which rely on similarity functions that use Artefact Attributes and Citizen Contribution Attributes. Consequently, different communities may be formed using different configurations, i.e. *perspectives*. Thus, in this context, defining a perspective is based on how the different attributes are shuffled, i.e., how users are brought together based on some established properties. This means that the same user can simultaneously be classified in different communities when using different perspectives. These perspectives can be generated by curators using VISIR (see Section 3.3.1).

Clustering algorithms are non-supervised learning methods to separate the data points contained in a dataset into groups (or clusters) in a way that the similarity between the data points in the same group is maximized (internal homogeneity) while the similarity between data points in different groups is minimized (separation). This way, similar data points are in the same cluster while the clusters themselves are dissimilar among them. Partitioning the dataset into clusters requires in most cases the use of similarity functions. Note that there is no clustering algorithm that can be universally used for every type of dataset. There are different factors that influence the selection of the clustering algorithm, like the size, sparsity, and dimensionality of the dataset, the type of features employed to describe the data points, the correlation among these data features, the outliers and noisiness of the data and the time complexity of the algorithm, and its stability. We refer interested readers to some of the many review papers on clustering algorithms [23, 32].

There are hundreds of different algorithms to solve the community detection task, each with its own understanding and definition of what a "community" is. See [8] for a similarity-based classification of community detection algorithms that compares more than seventy approaches. Note that the contribution of this paper is not on the definition of new clustering algorithms. Instead, we focus on how the resulting community model and

[8]Community Model infrastructure is designed and implemented in the context of the SPICE project. It is available at https://github.com/spice-h2020/spice-community-model and a live version of this API for the GAM case study described in Section 4 is available at https://spice.fdi.ucm.es/gam/

its associated explanations can help museum curators to make sense of the Interpretation Reflection Loop that was implemented in their particular cases.

In the following sections we will describe which similarity functions implemented in the Community Model can be employed in VISIR for the clustering and community detection algorithms, and how these algorithms are modified to generate interpretable communities.

3.2.1 Similarity functions. Similarity and its complementary notion of distance are essential to solving a broad range of problems in different AI domains and applications, including the problem of community detection since they can serve as an organizing principle by which individuals classify objects, form concepts, and generalize them.

We have compiled a catalog of similarity measures that are applicable to different types of attributes that are specific for the cultural heritage domain (colors, materials, emotions, themes, moral values, ...) [10]. Note that attributes can be simple, numeric, symbolic or textual, or, in some applications, they may use derived features obtained by inference, based on domain knowledge, may that be necessary. Also, when dealing with complex structures (such as graphs or first-order terms) similarity requires an assessment of their structural similarity [3]. The most relevant similarity functions employed by the Community Model are the following:

1. For Citizen Contribution attributes:
   (a) Emotion-based similarity [10] using Plutchik emotion theory [28].
   (b) Expert-based similarity for Moral Foundation Values [13].
2. For Artefact attributes:
   (a) Taxonomic similarity [3] for artefact content description using Iconclass [31], a taxonomy maintained for decades, and recently published as Linked Data[9] to support a variety of computational tasks [24, 29].
   (b) Taxonomic similarity for artefact materials. The material taxonomy[10] has been created ad-hoc based on the artefacts in GAM and DMH case studies and inspired by the taxonomy described in [21].
   (c) Year similarity. Although the year can be treated as a number and range and equality functions can be applied, we have also added functions for comparing decades and centuries.
   (d) Color similarity based on Delta E 2000 [25], that quantifies the difference between two colors as perceived by the human eye. Colors can be provided as part of the artefact description or extracted from artefact images [10].

3.2.2 Explainable clustering. Communities in computer science represent a fundamental concept for understanding and analyzing complex systems, particularly in the context of networks and data with inherent structural patterns. The detection and study of communities play a significant role in various computational and data-driven tasks [7]. The communities generated by a clustering algorithm should have a meaning for the expert who is analyzing the resulting clusters. If the communities are not meaningful the community detection process might be useless. For this reason, the Community Model not only generate communities based on clustering algorithms and similarity functions but also uses the later functions and the data attributes involved in the process to provide an explanation that describes the community.

The Community Model proposes an Explainable Community Detection process, which runs for different clustering algorithms (like K-medoids, Birch or OPTICS, among others) based on similarity that enriches the discovered communities using community explanations. These explanations can follow three complementary approaches:

[9]https://iconclass.org/help/lod
[10]This taxonomy is depicted in https://miro.com/app/board/uXjVP6eke_g=/?share_link_id=747726058527

- Based on demographic attribute values of the citizens who belong to a community. The explanation gathers the value distribution for each citizen attribute concerning the explicit community she belongs to.
- Centroid-based explanations. These explanations are based on creating a non-existent model citizen that acts as the representative citizen of this community. The data that characterizes this model citizen meets two requirements: it has the highest average similarity with all the other members of the community, and it belongs to the dominant value for each demographic attribute. To define this citizen, we choose the existing community member with the highest average similarity with other community members, also known as the medoid, and assign it the dominant values of the explicit citizen attributes.
- Based on the representative values of the similarity attributes involved in the clustering process (citizen contribution and artefact attributes). These explanations are created in two stages. First, the model stores information describing the relationship between the citizen similarity attributes involved in the clustering algorithm. Second, the model filters the previous information until only the citizen contributions used to compute the similarity between community members remain, reducing data noise and providing representative community interaction data as an explanation. Since these descriptions are usually too long or complex to be properly understood by human beings, a simplification process is applied.

For the last explanation approach, we distinguish between similarity attributes that do not follow a hierarchical structure and the ones that do, such as taxonomies and ontologies. In the first case, for each citizen-citizen comparison, we store in a list the attribute values involved in the computation of its distance from the other user and reduce it to the most frequent value. At the community explanation stage, we perform another simplification: For each citizen, we consider the values encoding its relationships with other community members and extract the most frequent values.

If the attributes used for the explanation are aligned with a hierarchical structure, like a taxonomy or an ontology, for each pair of attribute values used by the similarity functions, we store the entities with similar attributes and their common parents. At the community explanation stage, we pick the parent entities with the largest number of associated artefacts that are representative of this community.

The explanation process affects the community detection process and how the clustering algorithms are executed. The result of a clustering algorithm is considered explainable if the explanation process can explain the discovered communities according to the attributes employed by the similarity functions required by the clustering algorithms. We say that a community C is *explainable* if, for a given similarity threshold p and a set of similarity attributes S, at least $p\%$ of the community members are described by the same value in one or more of the similarity attributes s in S. For example, if the community is based on the similarity of the evoked emotions, the community would be considered as explainable only if at least $p\%$ of the community members are associated with the same emotion. If a community is not explainable, the clustering output is considered invalid. In this case, the clustering algorithm runs again with an increased number of clusters. If the number of clusters cannot be directly provided to the clustering algorithm, an approximation based on the available parameters is implemented to simulate it. This process continues until the number of clusters is equal to the number of citizens (in which case, all the citizens are classified as not belonging to any community) or an explainable output is yielded.

This way, the Explainable Community Detection algorithm proposed in the Community Model (see Algorithm 1) runs as follows:

1. Compute the similarity matrix between artefacts according to similarity functions to compare any pair of artefacts.
2. Compute the similarity matrix between citizens using similarity functions on their contributions. For each pair of citizens:
    (a) Compute the similarity between each pair of artefacts that both citizens have contributed to.
    (b) If the similarity is above a threshold, compute the similarity between the contribution attributes.

(c) Aggregate the contribution similarity to compute the final similarity between two citizens.
(3) Run the clustering algorithm using the similarity matrices computed above.
(4) This process is repeated, starting from 2 clusters up to a maximum of clusters equal to the number of citizens, until *all the obtained clusters are explainable*.

```
Data:
A = {a_1, ..., a_n} ; // List of artefacts
C = {c_1, ..., c_m} ; // List of citizens
I(c_j) = {i(a_1), ..., i(a_p)} ; // List of contributions for citizen c_j
Sim_A ; // Similarity function between artefacts
Sim_I ; // Similarity function between citizen contributions
// Compute Artefact similarity Matrix SM_A
SM_A ← Sim_A(a_i, a_j) ∀i, j|0 ≤ i < j < n;
// Compute Citizen Contribution similarity Matrix SM_I
foreach (c_i, c_j) ← c_i, c_j ∈ C; ∀i, j|0 ≤ i < j < m do
    SM_I[c_i, c_j] ← 0;
    foreach (a_k, a_l) ← i(a_k) ∈ I(c_i); i(a_k) ∈ I(c_j) do
        if SM_a[a_k, a_l]) > threshold then
            SM_I[c_i, c_j] ← agg(SM_I[c_i, c_j], Sim_I(i(a_k), i(a_k));
        end
    end
end
// Run clustering algorithm until the resulting clusters are interpretable
explainable ← false;
nClusters ← 2;
while (nClusters ≤ m)and(!explainable) do
    Clusters ← RunClustering(A, C, I, SM_A, SM_I, m);
    explainable ← true;
    foreach cl ∈ Clusters do
        if not Explainable(cl) then
            // Run cluster algorithm again
            explainable ← false;
            break;
        end
    end
end
return Clusters;
```

Algorithm 1: Explainable Community Detection algorithm

3.3 VISIR: Visualization and interpretation

Explainable clusters discovered by the Community Model are employed by the VISIR tool (VISualization for Interpretation and Reflection, see Figure 4) for the exploration of these emerging communities. VISIR is a web app[11] employed by museum curators to detect, visualize, and explain emerging communities based on previously defined perspectives. Thus, the visualization supports the introspection of the implicit communities according to each defined perspective. Each perspective instantiates a Community Model that categorizes users in implicit communities, so the visualization of a perspective will support the introspection of the explicit attributes shared between the implicit groups found.

Moreover, to assist curators in the analysis of these communities and perspectives, VISIR allows viewing two different perspectives with common citizens simultaneously, wherein each perspective visualizes the communities computed by the Community Model. This comparison ensures a reflection on the different viewpoints that arise from the citizen contributions, highlighting similarities, differences, and interrelations.

VISIR is not dependent on a specific clustering algorithm and relies on a configuration-visualization interactive loop. The different configurations and visualizations help to uncover latent information and to reveal citizen curation insights, usual and unusual distributions of citizens, local patterns, gaps, and outliers. VISIR allows the exploration of different perspectives thereby supporting the *interpretability* across target groups. Apart from the exploration, reflective reasoning, and social cohesion, communities are key to finding related contents of interest and avoiding the cold start situation as, due to similarity intra-community, a new user can be treated like other users from the same community.

3.3.1 Perspective configuration. Perspective configuration allows the curator to configure the Community Model around her interests by clustering the users in a set of implicit communities according to the selected configuration of interaction attributes and properties. For example, curators in GAM (see Section 4) will be interested in finding out what kind of users have "similar emotions when they interact with similar artworks".

The configuration of a perspective follows the process below:

- Feature selection and extraction: it consists on choosing the set of attributes from the data that will be employed by the clustering algorithm.
- Clustering algorithm selection or design: it consists of choosing the most suitable algorithm according to the features selected in the previous step. This step also implies the definition of the similarity measures and other criterion functions that the algorithm employs. VISIR offers a catalogue of clustering algorithms and similarity metrics implemented in the Community Model.
- Visualization and Explanation: it consists of reviewing and validating the clusters and the explanations created by the Community Model.

Note that this is not a one-shot process because cluster and explanation validation can trigger new trials and repetitions to find the most suitable communities.

3.3.2 Perspective Interpretation. The visualization used by VISIR lies in a simple 2D visual metaphor, where the citizens are represented by nodes who are confined in bounding boxes, which represent the implicit communities discovered using a defined perspective.

Explainable clustering algorithms described in section 3.2 are vital to help VISIR users to understand why and how data points are grouped together and make clustering results more transparent and interpretable. The goal of VISIR visualizations is to explain the sense of identity and belonging associated with the implicit communities and the in-group/out-group nature of this identity. Here we detail specific features that VISIR uses for visualizing the implicit communities and explanations provided by the Community Model:

- The visual aspect of a node. It is defined according to the attributes that characterize the explicit communities to which the subjects belong to. In the present version, two categorical demographic attributes can

[11]A live version with specific perspectives for this paper is available at https://gjimenezucm.github.io/SPICE-VISIR/. VISIR was designed and implemented as part of the SPICE project and its source code is available at https://github.com/spice-h2020/VISIR

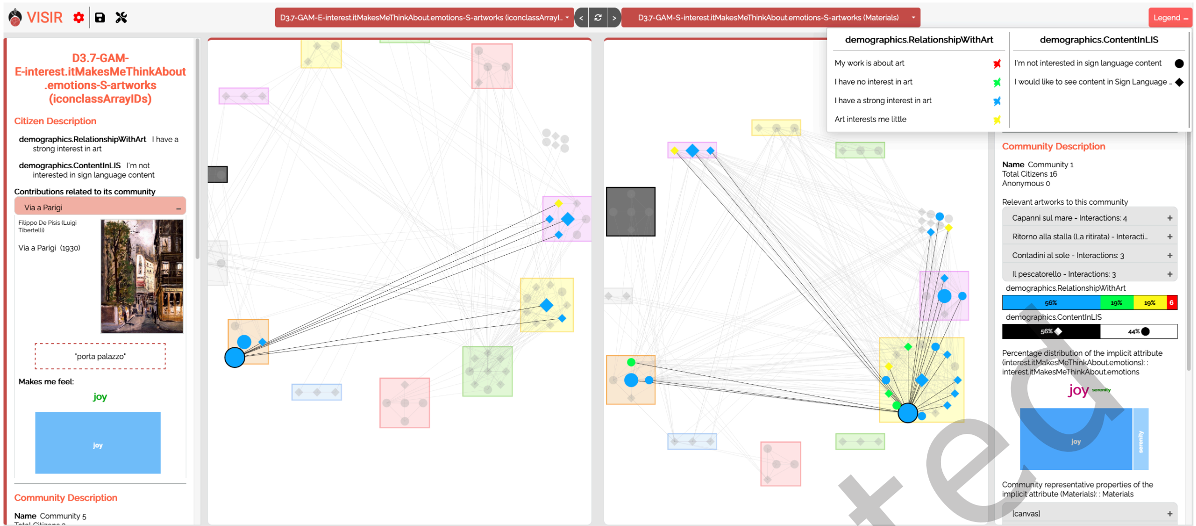

Fig. 4. Screenshot of VISIR main interface. Implicit communities emerged from two different perspectives can be simultaneously visualized and compared. The nodes visible in both perspectives are colorized according to the citizen's relationship with the art and the shape of the nodes represents the target group a citizen belongs to (d/Deaf community).

be employed simultaneously, one symbolized by the node shape and one by its color[12]. For example, the nodes in Figure 4 are colorized according to the citizen's relationship with the art and the shape of the nodes represents the target group a citizen belongs to in GAM case study (d/Deaf community).

- Legend. The legend accounts for the fact that people can be part of different communities simultaneously and work together with the previous feature. Additionally, a fine-grain exploration of these attributes is supported by interacting with legend elements, to filter the users who share specific dimensionality and, hence, belong to the same explicit community (see Figure 4, top right).
- Links connecting citizens. These elements help VISIR users to understand the compactness and separability of communities. We compute and present intra-cluster distances (within-cluster variance) and inter-cluster distances (between-cluster separations). Then, these distances are represented in VISIR by lines that connect nodes by similarity (according to the perspective).
- Special Explanation nodes. We describe each community using the centroid based-explanation (described in Section 3.2.2), creating a prototypical citizen of the community, termed a *representative citizen*, that is highlighted in the center.
- Interactive nodes. When clicking a node we see their description on the explanation panels on the sides (Figure 4 on the left). In this case, VISIR shows information about the citizen attributes represented by the selected node. Moreover, VISIR user can explore the citizen contributions that implies that this citizen belong to this community (see the "Contributions related to its community" panel on Figure 4. on the left) Finally, this node is highlighted in both perspectives, to see the different communities that the citizen can belong to (see highlighted node in Figure 4).
- Interactive communities. When clicking on a community bounding box, VISIR shows two types of explanations on a side panel:

[12]The use of two features has been recognized as simple and significant in experiments when compared with including other options like line colors, strokes, and dashed lines

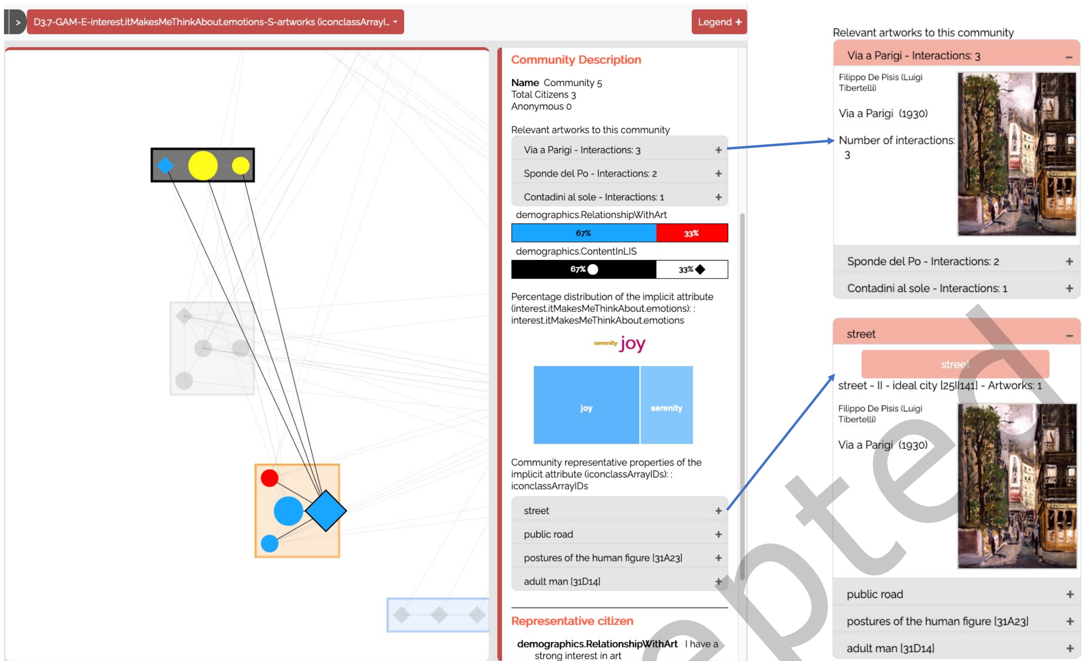

Fig. 5. Screenshot and detail when clicking on VISIR Explanation panel for describing discovered communities.

- Demographic community data: number of citizens and distributions of demographic attribute values of the citizens who belong to a community (see explanation types described in Section 3.2.2).
- Representative values of the similarity attributes involved in the clustering process, that contributed the most to the formation of clusters (generated by the Community Model, as described in Section 3.2.2). An example is shown in Figure 5, where the most representative emotions of the selected cluster (joy and serenity) are represented using a word cloud and a treemap. Additionally, the Iconclass concepts that are more relevant in the representative artefacts of the community (like "street" and "public road") are also detailed.

- Interactive explanations. Explanation panels have different interactive elements so VISIR user can click on them to visualize other communities that share some values in the same or different perspectives. For example, when clicking on the "Joy" element of the community explanation on Figure 5, communities that share this attribute are highlighted in the displayed perspectives.

For example, the nodes in Figure 4 are colorized according to the citizen's gender and the shape represents the citizen's relationship with art. A fine-grain exploration of these attributes is supported by an interactive legend (top right), to filter the users who share specific dimensionality hence, belong to the same explicit community. It allows exploring and getting insights into the distribution of each explicit community through the detected implicit communities based on contributions.

4 GAM: A CASE STUDY FROM SPICE

This section provides an illustrative example of the IRL workflow implemented within the GAM case study. It is important to emphasize that the data used in the GAM case study was obtained from a pilot study with a limited number of participants, focusing on specific attributes relevant to the case-specific context and the target group (d/Deaf community). It serves as an illustrative instance that highlights its application specifically in the GAM case study and the context of SPICE. We therefore describe the use of the Community Model and VISIR in the GAM case, and discuss some of the possible limitations of this example.

The Galleria Civica d'Arte Moderna e Contemporanea (GAM) in Turin, Italy is a museum operated by the Fondazione Torino Musei. With a collection of approximately 45,000 artworks, GAM encompasses works from the 19th Century to the present day. In SPICE, the GAMGame has been developed by the University of Turin in a co-design process with the museum and the Turin Institute for the Deaf as stakeholders. Designed and implemented as part of the SPICE project, the GAMGame web app enables anonymous creation and sharing of personal interpretations of artworks in narrative form, with a particular focus on d/Deaf teenagers. In the app, citizens are invited to create their own *stories* in the style of social media stories, using artworks from the GAM museum collection [19].

The GAMGame aims at reducing the impact of text in art interpretation and communication, replacing written text production –a barrier for speakers of Sign Languages– with visual storytelling and emotional interpretation. Moreover, the GAMGame is designed to support multiple-use scenarios, both in the museum and outside the museum, as well as before and/or after the museum visit. As illustrated in Figure 6, the platform allows users to put together their own selection of artworks from the museum collection, and annotate the selected artworks with emojis, tags, and short text comments. In the IRL, this stage corresponds to the Interpretation: in the GAMGame, its output is a set of visual narratives enriched with comments, tags and emojis (the Citizen Contribution Attributes in Figure 1). To trigger Reflection, users have the option to share their stories anonymously with the other users of the platform. Additionally, users can explore and interact with the stories created by others.

In order to enhance diversity in both the interpretation and reflection stages, in the GAMGame users receive recommendations of artworks associated with similar and opposite emotions; similarly, when browsing the stories featuring a given artwork, they receive recommendations about stories with similar and opposite emotions. For recommendations, the GAM game relies on the DEGARI system [17], an affective-based reasoning engine that uses Plutchik's model of emotions to track similarity and opposition relations over emotion types. Figure 7 shows the use of the emotions "mined" from the user-created stories to create affective-based, diversity-oriented recommendations of artworks. The GAMGame design and the use of DEGARI as sensemaking tool are described in [19].

4.1 GAM data collection

During the European Researchers' Night at University of Torino (November 1st, 2022), SPICE researchers collected 149 stories from 49 casual users who volunteered to use the GAMGame app.

Users logged in to the GAMGame using anonymous IDs for privacy reasons; age ranges were used so as to minimize the level of detail with respect to the experimental goals.[13] Figure 8 shows the pie chart statistics on the user attributes: Most of the users were aged between 20 and 30 years old (about 58.3%) (Figure 8 (a)); users were evenly distributed according to gender (about 43.8% of male and female)(Figure 8 (b));(52.1%) had a strong interest in receiving contents in Italian Sign Language (Lingua Italiana dei Segni (LIS)) (Figure 8 (c)); 37.5% users stated that they often visit museums and art galleries (Figure 8 (d)). Finally, (64.6%) reported a strong interest in art (Figure 8 (e)).

To identify the overlapping categories of attributes for exemplifying the use of VISIR in the context of the GAM case study, we consider the following attributes which can be used to derive the emerging clusters:

[13] The experiments conducted to develop and test the GAMGame were approved by the Ethical Committee of the University of Turin and all the user data were anonymously collected and stored according to the project data management plan (https://spiceh2020.eu/document/deliverable/D1.2.pdf)

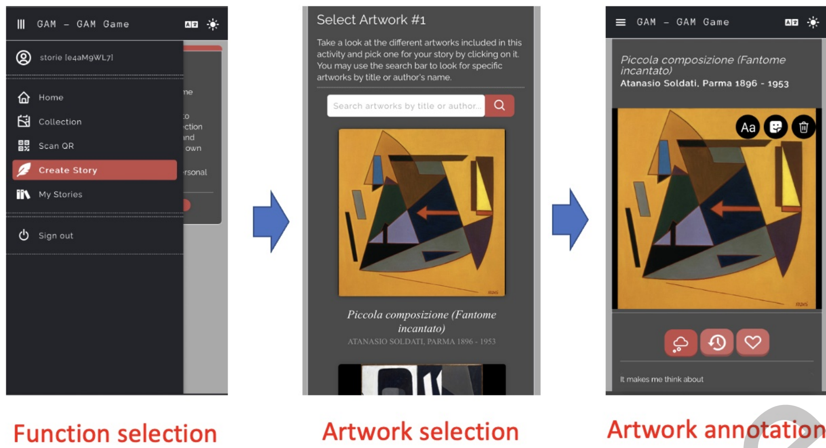

Fig. 6. Screenshots of the GAM Game application wherein a user is creating a story (left) by selecting an artwork (centre) and annotating it with emojis, tags and short comments (right).

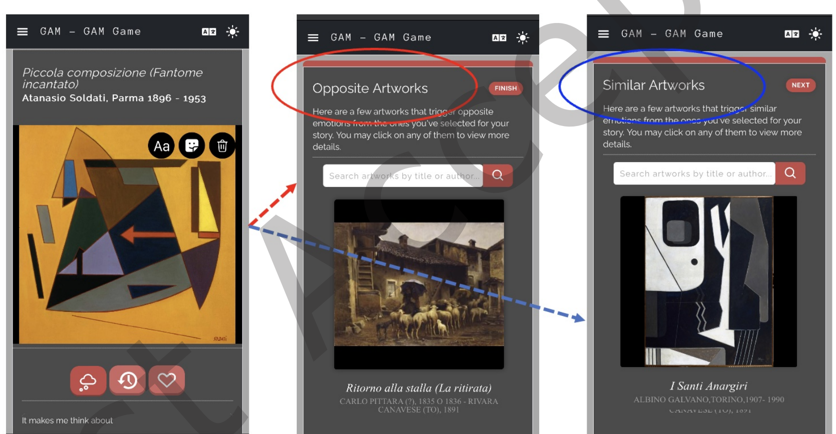

Fig. 7. Recommendations of artworks with similar and opposite emotions in the GAMGame.

- **Demographic attributes**: 5 different attributes were collected from the users of the GAMGame through an online form that they had to fill in before using the app: *relationship_with_arts* encodes the citizen's interest in art; *relationship_with_museums* encodes the frequency with which the citizen visit museums; the binary attribute *interest_in_LIS* encodes her interest in contents in Italian Sign Language (Lingua Italiana dei Segni, LIS), roughly corresponding to the belonging to the d/Deaf community; *gender* (male, female and non-binary) and *age range*.

- **Artefact attributes**: this set of attributes includes both standard museum catalogue metadata, such as *title*, *technique*, *dimension* (*size_height*, *size_width*, *size_depth*), *collection*, *year* (including *artwork_start_date* and *artwork_end_date*) and *type*, and a set of attributes added by the museum staff to investigate the

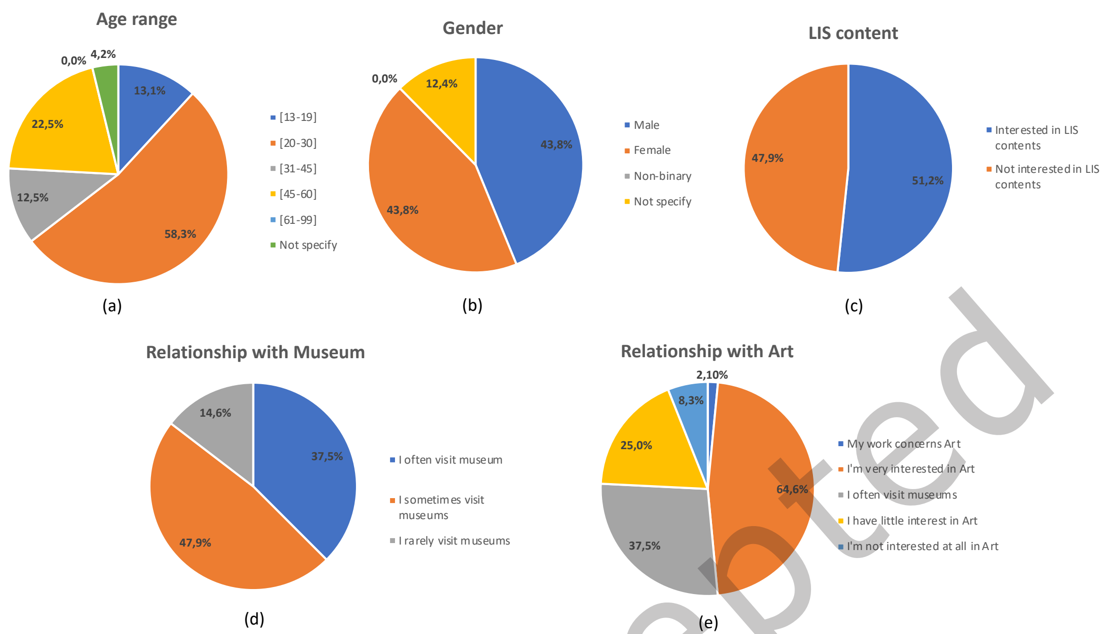

Fig. 8. Statistics on user attributes in GAMGame dataset

users' relationship with the artworks. In particular, based on the experience gathered during museum labs with schools and communities, the latter includes information on the artwork's *subject*, *materials*, and artistic movement (*artwork_artistic_movement*), since the visitors tend to engage with the subject depicted by the artwork, the artistic movement it belongs to, and to the materials that compose it. Since visitors are sometimes aware of the artist's main biographic data, specific attributes for the artist's nationality (*artist_country*), birth and death date (*artist_birth_date*, *artist_death_date*), and *gender* were also included in the Artefact Attributes.

To avoid arbitrariness in identifying the subjects depicted in the artworks, the curators resorted to an authoritative, standard classification of iconographic subjects, called Iconclass, to encode the subject of the artefacts. Widely employed in digital heritage research [1, 11], Iconclass organizes iconographic subjects into 9 main subject types (Abstract, Non-representational Art; Religion and Magic; Nature; Human Being, Man in General; Society, Civilization, Culture; Abstract Ideas and Concepts; History; Bible; Literature; Classical Mythology and Ancient History), further subdivided into more specific types of varying specificity (e.g., "adolescent, young woman, maiden" or "sea (seascape)", which are reoccurring subjects in the GAM collection), accompanied by keywords to facilitate their indexing and search.

- **Citizen contribution attributes**: for each artefact, a list of emotion labels was extracted from the curatorial notes, and from the stories generated by the users through the GAMGame. Curatorial notes and user-generated text in English and Italian (comments and tags) were fed into the semantic analysis and reasoning pipeline of SPICE, which relies on Plutchik's theory of emotions [28], as described in [4]. The pipeline encompasses two main tools: the SOPHIA emotion extractor which has a wide coverage in detecting basic emotions, while the DEGARI reasoning system can use basic emotions to identify

GAM Artefact	SOPHIA emotions	DEGARI emotions
 (Der) Matrose Fritz Múller aus Pieschen	Disapproval Disgust Interest **Joy** Serenity Sadness	Delight Love Optimism Curiosity **Joy** Pride
 Contadini al sole (Farmers in the sun)	Serenity Anticipation Joy Surprise Trust	Despair Disapproval Envy Pessimism Remorse Sentimentality
 Le tre finestre (The three windows)	Fear Sadness Interest **Joy** **Love** Serenity Surprise	Awe Curiosity Delight **Love** Hope **Joy**
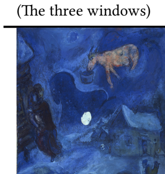 Dans mon pays	Fear Sadness Disgust Anticipation Surprise	Curiosity Delight Love Joy

Table 1. Simple and complex emotions extracted by the semantic analysis and reasoning pipeline in SPICE (overlapping emotions are shown in bold).

fine-grained, complex emotions. This combination is technically obtained by automatically updating knowledge graphs of the GAM artefacts. The examples in Table 1 show that the emotions extracted by DEGARI and SOPHIA for the GAM artworks "(Der) Matrose Fritz Múller aus Pieschen","Contadini al sole" ("Farmers in the sun"), "Le tre finestre" ("The three windows") and "Dans mon pays" which are complementary in nature.

The GAMGame installation created for the experiment included 56 artefacts from the GAM collection, selected by the curators with the aim of presenting the users with a variety of artworks by author, type, style and time period, thus avoiding any biases for specific artwork types. This choice is reflected by the value distributions of the artwork attributes, which encompass 12 different types of material, 26 artistic movements, 21 techniques, 4 artwork types and artists from 10 different countries.

4.2 Perspectives in the GAM case study

The *perspectives* to be used in the GAM case were defined based on the analysis of similar, different, and identical emotions when interacting with similar or the same artworks, using Plutchik similarity between emotions [10], theme similarity based on Iconclass data, and relationship with art. These attributes were indicated as relevant for the reflection on how the museum collection is interpreted in the visitors' stories by the museum curators, based on the feedback on their collection obtained through the GAMGame and their previous experience with visitor groups in the museum. All the perspectives use an explainable version of the generative agglomerative clustering algorithm [27]. For the different perspectives, in this example we have reflected on the d/Deaf/non-deaf categorization(*interest_in_LIS*), gender, and relationship with art (*relationship_with_art*) as the citizen demographic attributes (in the legend).

4.2.1 VISIR exploration. To support the unveiling of connections between citizens, communities, and the cultural heritage space of GAM, pertinent perspectives for the museum were chosen to be visualized and juxtaposed (in accordance with the museum's interests and the SPICE objectives). In relation to the goals of SPICE, the intention behind juxtaposing two or more perspectives was to explore possibilities for representing the "sense of belonging" as a dimension of social cohesion. Thus, by examining how citizens can belong to the same or different communities across different perspectives, the museum could uncover connections between citizens that may not have been evident solely through the analysis of demographic data and preliminary insights. Below we shortly describe the cluster visualizations emerging from these perspectives.

4.2.2 GAM Perspectives and emerging clusters. Perspective 1 considered *the same emotions in the same painting*, i.e. the painting "Dans mon Pays" by Marc Chagall [14]. On the other hand, Perspective 2 considered *the same emotions in similar paintings*, involving artworks classified as similar based on the iconclass classification[15]. The choice of the painting for perspective 1 was based on the preliminary insight that the painting was reported as attracting more interaction across both citizen groups (d/Deaf and non-deaf). On the other hand, perspective 2 was established to see if paintings classified as similar (based on iconclass) would attract the same type of emotional responses for citizens across different demographic attributes, i.e. interest in art and the LIS/non-LIS categorization.

Regarding the demographic attributes for the established perspectives, the GAM case study was specifically targeting the d/Deaf community, making the LIS/nonLIS parameter a focal point. Additionally, the museum considered attributes related to the visitors' relationship to art as central. This was thought to not only offer insights into the relationship of emotional responses and the level of interest in art but also to support the museum in engaging potential new audiences.

From perspective 1, in our example, two interesting clusters relating to emotional responses emerged: a cluster featuring "fear" (3 citizens) and a cluster featuring "joy" (3 citizens) Figure 9.

By juxtaposing the two perspectives, as outlined in Section 3.3.2, and considering the same citizens in the two perspectives simultaneously, we could observe linkages and shifts in terms of belonging to implicit groups. For example, the 3 citizens in the first implicit community in perspective 1 (community 1, "fear") all reappear in different clusters in the second perspective. However, with respect to the second cluster in the first perspective (community 2, "joy"), two citizens (Citizen 4, non-LIS/little interest in art and Citizen 6, LIS/strong interest in art), also fall in the same implicit community (community 5) in the second perspective (see Figure 9). These clusters were interesting not only due to the distinctively opposite emotional responses but also due to featuring citizens with different levels of interest in art and different preferences in terms of LIS.

5 CONCLUSIONS AND FUTURE WORK

As highlighted in the introduction, the SPICE project explored five unique museum cases, each targeting diverse audiences and involving distinct curatorial activities, digital interfaces, and user journeys. Shared across all the

[14]This perspective is available in VISIR (https://gjimenezucm.github.io/SPICE-VISIR/), selecting the perspective called `_JOCCH_sameEmotionsDansMonPays`

[15]This perspective is available in VISIR (https://gjimenezucm.github.io/SPICE-VISIR/), selecting the perspective called `_JOCCH_sameEmotionsSimArtworksIconclass`

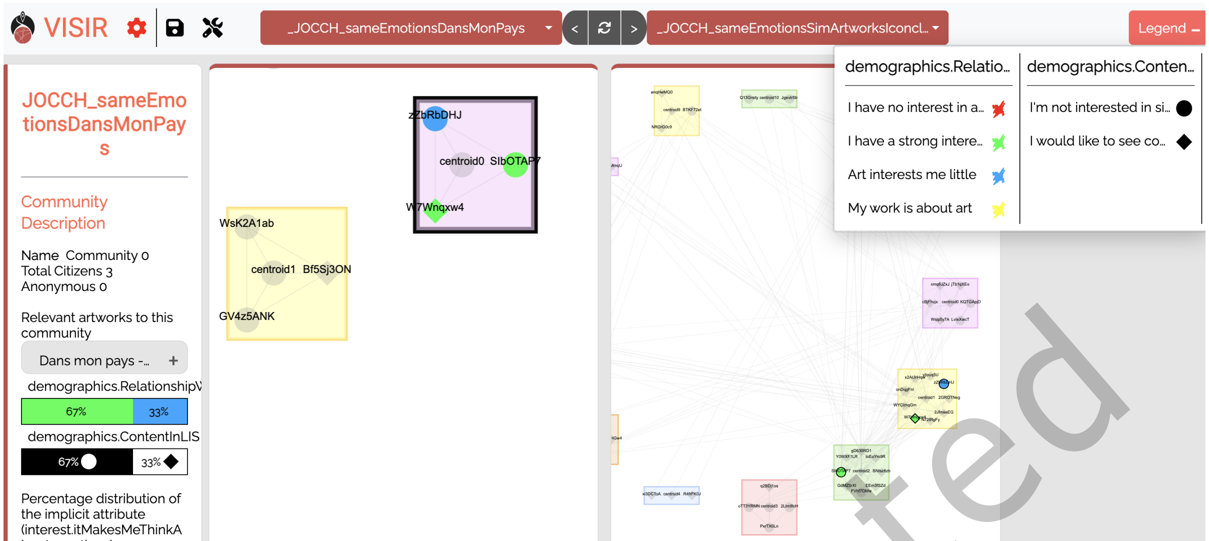

Fig. 9. Screenshot from the two perspectives in the VISIR tool. On the left, emerging clusters from the first perspective, (highlighted in color is 3 citizens in community 2, "joy"). On the right, emerging clusters from the second perspective, (highlighted in color are the citizens from community 2 that now emerge in new clusters in the second perspective).

case studies was the utilization of the SPICE technical platform, the idea of the Interpretation-Reflection Loop (IRL), and the application of the framework and tools presented in this paper. One of the aims of the IRL was to focus on the "sense of belonging" as an important dimension of social cohesion. In the paper, we have described how the SPICE technical tools integrated into the SPICE system can be combined to guide and facilitate such interpretation and reflection processes. The Community Model facilitates modeling citizens and communities utilizing similarity functions to cluster citizens together according to different perspectives. These communities are then explained using demographic attributes, centroid-based representations, and similarity attributes. The clustering algorithm is designed to create meaningful communities and iterates until all clusters are explainable. We claim that this model is valuable for understanding and organizing complex data in cultural heritage contexts and we have illustrated our approach by considering different attributes on individual citizens and citizen groups in the GAM (Galleria Civica d'Arte Moderna e Contemporanea) case study. Perspectives were defined based on the analysis of emotions when interacting with artworks and demographic attributes. These perspectives were visualized using the VISIR tool to explore connections between citizens, communities, and the cultural heritage space. Two key perspectives were examined: one focusing on the same emotions in the same painting and another on the same emotions in similar paintings. These perspectives revealed clusters related to similar emotional responses and insights to understand the connections between citizens, communities, and their interactions with cultural artefacts.

Future work includes running comparative analysis among the five SPICE cases and the evaluation of the clustering techniques involved in the community model. Common offline evaluation metrics are not feasible for our approach because community explainability is a feature more important than the compactness, the separation or the clustering overlapping, among others. Although the curators who used VISIR have informally validated the interest of the discovered communities, it is necessary to evaluate the impact of the categorization of the visitor on that communities on her reflection process. It should follow a user-centered evaluation approach and should test the impact of the IRL in complex human values like the sense of belonging, empathy or serendipity, among others. Moreover, VISIR visualizations may be shared with the museum visitors and users of the GAMGame, thus exploring their use –currently limited to the museum staff– for enhancing the sense of belonging of citizens to their explicit and implicit communities.

Finally, we are adapting methods, tools, and techniques from data storytelling practices for exploring how to visualize heterarchical relations more effectively. The creation of data stories, meta-stories, visualizations, or innovative recommending criteria can enhance reflection among the audience and within the museum community [15]. While these meaning-making tools are originally intended for facilitating the analysis and the representation of citizens among museum audiences and within the museum community, they can potentially be adapted to inspire similar tools in fields and domains such as pedagogy, social work, community building, conflict resolution, and advocacy initiatives.

ACKNOWLEDGMENTS

The research leading to these results/this publication has been partially funded by the European Union's Horizon 2020 research and innovation programme http://dx.doi.org/10.13039/501100007601 under grant agreement SPICE 870811. The publication reflects the author's views. The Research Executive Agency (REA) is not liable for any use that may be made of the information contained therein

REFERENCES

[1] Anila Angjeli, Antoine Isaac, Thierry Cloarec, Frédéric Martin, Lourens Van-der Meij, Henk Matthezing, and Stefan Schlobach. 2009. Semantic web and vocabulary interoperability: an experiment with illumination collections. *International cataloguing and bibliographic control* 38, 2 (2009), 25–29.

[2] Luigi Asprino, Enrico Daga, Aldo Gangemi, and Paul Mulholland. 2023. Knowledge Graph Construction with a Façade: A Unified Method to Access Heterogeneous Data Sources on the Web. *ACM Trans. Internet Technol.* 23, 1, Article 6 (feb 2023), 31 pages. https://doi.org/10.1145/3555312

[3] Pablo Beltrán-Ferruz, Belén Díaz-Agudo, and Oscar Lagerquist. 2006. Retrieval over Conceptual Structures. In *Advances in Case-Based Reasoning, 8th European Conference, ECCBR 2006 (Lecture Notes in Computer Science, Vol. 4106)*, Thomas Roth-Berghofer, Mehmet H. Göker, and H. Altay Güvenir (Eds.). Springer, 443–457. https://doi.org/10.1007/11805816_33

[4] Andrea Bolioli, Alessio Bosca, Rossana Damiano, Antonio Lieto, and Manuel Striani. 2022. A complementary account to emotion extraction and classification in cultural heritage based on the Plutchik's theory. In *UMAP '22: 30th ACM Conference on User Modeling, Adaptation and Personalization*. ACM, 374–382. https://doi.org/10.1145/3511047.3537659

[5] Luis Bruni, Enrico Daga, Rossana Damiano, Lily Diaz, Tsvi Kuflik, Antonio Lieto, Aldo Gangemi, Paul Mulholland, Silivio Peroni, Sofia Pescarin, and Alan Wecker. 2020. Towards Advanced Interfaces for Citizen Curation. In *AVI2CH Workshop on Advanced Visual Interfaces and Interactions in Cultural Heritage colocated with 2020 International Conference on Advanced Visual Interfaces (AVI 2020)*.

[6] Luis Emilio Bruni, Nele Kadastik, and Thomas Anthony Pedersen. 2022. SPICE Technical Report - D2.4 REVISED METHODS FOR REFLECTION. (2022). https://doi.org/10.13140/RG.2.2.28870.96321

[7] Iván Cantador and Pablo Castells. 2011. Extracting multilayered Communities of Interest from semantic user profiles: Application to group modeling and hybrid recommendations. *Computers in Human Behavior* 27, 4 (2011), 1321–1336. https://doi.org/10.1016/j.chb.2010.07.027

[8] Michele Coscia. 2019. Discovering Communities of Community Discovery. *CoRR* abs/1907.02277 (2019). arXiv:1907.02277 http://arxiv.org/abs/1907.02277

[9] Enrico Daga, Luigi Asprino, Rossana Damiano, Marilena Daquino, Belen Diaz Agudo, Aldo Gangemi, Tsvi Kuflik, Antonio Lieto, Mark Maguire, Anna Maria Marras, et al. 2022. Integrating citizen experiences in cultural heritage archives: requirements, state of the art, and challenges. *ACM Journal on Computing and Cultural Heritage (JOCCH)* 15, 1 (2022), 1–35. https://doi.org/10.1145/3477599

[10] Belén Díaz-Agudo, Guillermo Jimenez-Diaz, and Jose Luis Jorro-Aragoneses. 2021. User Evaluation to Measure the Perception of Similarity Measures in Artworks. In *Case-Based Reasoning Research and Development*, Antonio A. Sánchez-Ruiz and Michael W. Floyd (Eds.). Springer, 48–63. https://doi.org/10.1007/978-3-030-86957-1_4

[11] Chris Dijkshoorn, Lizzy Jongma, Lora Aroyo, Jacco Van Ossenbruggen, Guus Schreiber, Wesley Ter Weele, and Jan Wielemaker. 2018. The Rijksmuseum collection as linked data. *Semantic Web* 9, 2 (2018), 221–230.

[12] Elias Dritsas, Maria Trigka, Gerasimos Vonitsanos, Andreas Kanavos, and Phivos Mylonas. 2021. Aspect-Based Community Detection of Cultural Heritage Streaming Data. In *2021 12th International Conference on Information, Intelligence, Systems & Applications (IISA)*. 1–4. https://doi.org/10.1109/IISA52424.2021.9555549

[13] Frederic R. Hopp, Jacob T. Fisher, Devin Cornell, Richard Huskey, and René Weber. 2021. The extended Moral Foundations Dictionary (eMFD): Development and applications of a crowd-sourced approach to extracting moral intuitions from text. *Behavior Research Methods* 53, 1 (Feb. 2021), 232–246. https://doi.org/10.3758/s13428-020-01433-0

[14] Nele Kadastik, Luis Emilio Bruni, and Thomas Anthony Pedersen. 2023. Citizen Curation Methods for Interpretation and Reflection on Cultural Heritage: Insights from SPICE. In *UMAP '23 Adjunct: Adjunct Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization.* Association for Computing Machinery, United States, 367–373. https://doi.org/10.1145/3563359.3596668

[15] Nele Kadastik and Luis Emilio Bruni. 2023. Storifying Data for Museum Audiences. ExICE 2023 - Extended Intelligence for Cultural Engagement. https://spice-h2020.eu/conference/.

[16] Andreas Kanavos, Maria Trigka, Elias Dritsas, Gerasimos Vonitsanos, and Phivos Mylonas. 2021. Community Detection Algorithms for Cultural and Natural Heritage Data in Social Networks. In *Artificial Intelligence Applications and Innovations. AIAI 2021 IFIP WG 12.5 International Workshops (IFIP Advances in Information and Communication Technology)*, Ilias Maglogiannis, John Macintyre, and Lazaros Iliadis (Eds.). Springer, Cham, 395–406. https://doi.org/10.1007/978-3-030-79157-5_32

[17] Antonio Lieto, Gian Luca Pozzato, Manuel Striani, Stefano Zoia, and Rossana Damiano. 2023. Degari 2.0: A diversity-seeking, explainable, and affective art recommender for social inclusion. *Cognitive Systems Research* 77 (2023), 1–17. https://doi.org/10.1016/j.cogsys.2022.10.001

[18] Antonio Lieto, Gian Luca Pozzato, Stefano Zoia, Viviana Patti, and Rossana Damiano. 2021. A commonsense reasoning framework for explanatory emotion attribution, generation and re-classification. *Knowledge-Based Systems* 227 (2021), 107166. https://doi.org/10.1016/j.knosys.2021.107166

[19] Antonio Lieto, Manuel Striani, Cristina Gena, Enrico Dolza, Anna Maria Marras, Gian Luca Pozzato, and Rossana Damiano. 2024. A sensemaking system for grouping and suggesting stories from multiple affective viewpoints in museums. *Human–Computer Interaction* 39, 1-2 (2024), 109–143. https://doi.org/10.1080/07370024.2023.2242355

[20] Veranika Lim, Nikos Frangakis, Luis Molina Tanco, and Lorenzo Picinali. 2018. PLUGGY: A Pluggable Social Platform for Cultural Heritage Awareness and Participation. In *Advances in Digital Cultural Heritage*, Marinos Ioannides, João Martins, Roko Žarnić, and Veranika Lim (Eds.). Springer, 117–129.

[21] Mohammed M. Mabkhot, Ali Al-Samhan, and Lotfi Hidri. 2019. An Ontology-Enabled Case-Based Reasoning Decision Support System for Manufacturing Process Selection. *Advances in Materials Science and Engineering* 2019 (08 2019), 1–18. https://doi.org/10.1155/2019/2505183

[22] Claudio Martella, Armando Miraglia, Jeana Frost, Marco Cattani, and Maarten van Steen. 2017. Visualizing, clustering, and predicting the behavior of museum visitors. *Pervasive and Mobile Computing* 38 (2017), 430–443. https://doi.org/10.1016/j.pmcj.2016.08.011

[23] Vivek Mehta, Seema Bawa, and Jasmeet Singh. 2020. Analytical review of clustering techniques and proximity measures. *Artif. Intell. Rev.* 53, 8 (2020), 5995–6023. https://doi.org/10.1007/s10462-020-09840-7

[24] Federico Milani and Piero Fraternali. 2021. A dataset and a convolutional model for iconography classification in paintings. *Journal on Computing and Cultural Heritage (JOCCH)* 14, 4 (2021), 1–18. https://doi.org/10.1145/3458885

[25] Wojciech Mokrzycki and Maciej Tatol. 2011. Color difference Delta E - A survey. *Machine Graphics and Vision* 20 (04 2011), 383–411.

[26] Héctor Moreno Mendoza, Agustín Santana Talavera, and José Molina González. 2021. Formation of clusters in cultural heritage – strategies for optimizing resources in museums. *Journal of Cultural Heritage Management and Sustainable Development* 11, 4 (Jan. 2021), 580–595. https://doi.org/10.1108/JCHMSD-12-2019-0155 Publisher: Emerald Publishing Limited.

[27] Frank Nielsen. 2016. *Hierarchical Clustering*. Springer, 195–211. https://doi.org/10.1007/978-3-319-21903-5_8

[28] Robert Plutchik. 1984. Emotions: A general psychoevolutionary theory. *Approaches to emotion* 1984, 197-219 (1984), 2–4.

[29] Gjorgji Strezoski and Marcel Worring. 2018. Omniart: a large-scale artistic benchmark. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 14, 4 (2018), 1–21. https://doi.org/10.1145/3273022

[30] Chieh-Ching Tien. 2010. The formation and impact of museum clusters: two case studies in Taiwan. *Museum Management and Curatorship* 25, 1 (2010), 69–85. https://doi.org/10.1080/09647770903529434

[31] H van de Waal, LD Couprie, E Tholen, RH Fuchs, WH de Haan-van de Wiel, HEC Mazur-Contamine, and NC Sluijter-Seijffert. 1974. *Iconclass: An Iconographic Classification System: System: System, 2-3.*

[32] Dongkuan Xu and Yingjie Tian. 2015. A Comprehensive Survey of Clustering Algorithms. *Annals of Data Science* 2, 2 (2015), 165–193. https://doi.org/10.1007/s40745-015-0040-1