A Knowledge Graph of Values across Space and Time

Martin Ruskov[1], Maria Dagioglou[2], Marko Kokol[3], Stefano Montanelli[1] and Georgios Petasis[2]

[1]Università degli Studi di Milano, 20100 Milan, Italy
[2]Institute of Informatics and Telecommunications, NCSR "Demokritos", 15310 Athens, Greece
[3]Semantika Research, Semantika d.o.o., 2000 Maribor, Slovenia

Abstract
The moral values present in classical texts are only implicitly codified and are a form of intangible cultural heritage. In this paper, we focus on capturing the perceptions of these references of values, with the aim to study how they evolve over space and time. To this end, we present our approach that consists of a meta-model and two intertwined design processes for the creation of a knowledge graph capable of capturing an integrated representation of both values and perceptions. In particular, we illustrate how the input collected from different activities with the general public can feed the graph to represent multiple and possibly divergent perceptions of values in classical texts. Our meta-model allows for the integration of data from both established scientific techniques such as expert annotations on the one hand, and on the other from standard tests, social media activity, visual representations and games. The proposed approach gives practical means to make explicit both historical and current perceptions of values in classical works of art. Our approach and the resulting knowledge graph enable the comparative analysis of values and their perceptions over space and time.

Keywords
intangible cultural heritage, citizen curation, values across space and time, knowledge externalisation

1. Introduction

The moral values present in classical texts are an important part of our intangible cultural heritage, and in particular when studying a common European identity. Yet, being only implicitly codified in artifacts, in different regions and different time periods (what we refer to as "over space and time") the perceptions of these values diverge [1]. We propose a way to formalise, and thus make explicit, such perceptions by integrating them in a single knowledge graph,

2nd Italian Workshop on Artificial Intelligence for Cultural Heritage (IAI4CH 2023, https://ai4ch.di.unito.it/), co-located with the 22nd International Conference of the Italian Association for Artificial Intelligence (AIxIA 2023). 6-9 November 2023, Rome, Italy
martin.ruskov@unimi.it (M. Ruskov); mdagiogl@iit.demokritos.gr (M. Dagioglou); marko.kokol@semantika.eu (M. Kokol); stefano.montanelli@unimi.it (S. Montanelli); petasis@iit.demokritos.gr (G. Petasis)
https://islab.di.unimi.it/team/martin.ruskov@unimi.it (M. Ruskov);
https://www.iit.demokritos.gr/people/maria-dagioglou/ (M. Dagioglou); https://islab.di.unimi.it/montanelli (S. Montanelli); https://www.iit.demokritos.gr/people/georgios-petasis/ (G. Petasis)
0000-0001-5337-0636 (M. Ruskov); 0000-0002-3357-2844 (M. Dagioglou); 0000-0003-4292-989X (M. Kokol); 0000-0002-6594-6644 (S. Montanelli); 0000-0001-8150-6334 (G. Petasis)

thus facilitating their comparative study. In the VAST project we focus on the core values of the European Union. Accordingly, we pursue two goals: i) to make explicit (i.e. externalise) values implicitly present in literary texts from the past; and ii) to digitise values as they are perceived by the general public today. In particular, we focus on studying the transformations of values from three important periods in European history, which we call VAST pilots: Pilot 1: from Ancient Greek tragedies to modern theatrical plays, Pilot 2: from 17th century works of natural philosophy to exhibits in science museums, and Pilot 3: from traditional fairy tales across Europe to contemporary narrative.[2] In this respect, a key objective of the VAST project – that is reported in this paper – is the construction of a knowledge graph where different and potentially conflicting interpretations about the considered European values are represented and coexist.

As a way to moral values, we take as a reference the definition of Schwartz's basic human values, defined as "trans-situational goals, varying in importance, that serve as guiding principles in the life of a person or group." His refined theory of values [3] features 19 values. These are clustered in four universal value dimensions: Self-enhancement, Self-Transcendence, Openness to Change and Conservation. Schwartz's framework is widely adopted, e.g. in the European Social Survey [4] and the making of EU policy [5], as well as in various fields including automatic values extraction from natural language [6, 7]. A particularly relevant recent development is ValueNet [8], an ontology relating Schwartz's values to established semantic frameworks like FrameNet and DBpedia. However, while the proposed model and integration are an important cornerstone towards formalisation of values, the authors stop short of the actual population of a semantic knowledge graph. Arguably as a consequence, they do not address the challenge related to different perceptions of values that is the focus of our work presented here.

The knowledge graph proposed here features four major aspects of user interactions with values present in a historical artifact: i) the participant's individual characteristics such as demographics or expressed beliefs (*who*), ii) the description of the artifact and experience that define the interaction (*how*), iii) the collection of the expressed perceived values (*what*), and iv) the spatio-temporal context of the interaction (*where* and *when*).

Next in this paper we review previous research, both in terms of representation of relevant knowledge, and collaborative knowledge graph creation. In Section 3.1 we present our conceptual model and knowledge graph engineering process. This is followed by examples of concrete knowledge externalisation activities and details about our technical implementation. We conclude with

2. Background

The two relevant fields of research are works about formalisms and models for representation of knowledge in the humanities, as well as collaborative methodologies for knowledge graph engineering. An important role in the crossroads of the two is the already mentioned ValueNet [8]. While it does provide an important ontological foundation for the representation of values, it does not yet address the issue of perceptions of values, which could be divergent across space and time. Thus, we see it as complementary to the work presented here.

2.1. Representation of Humanities Knowledge

When looking for ways to represent historical interpretations, another natural candidate could be the *CIDOC Conceptual Reference Model* (CRM). It features elements to allow for the representation, integration, mediation, and interchange of documents and scientific activities [9]. Yet, formal models like ValueNet and CIDOC-CRM inherently assume an absolutist epistemology, where knowledge is universally established, and the challenge is to merely represent it [10]. While CIDOC-CRM does allow for "subjective opinions and inferences", it falls short of addressing cultural interpretations and differences that might lead to possible conflicting ontological constructs [11]. Such an epistemological perspective is too restrictive when working with intangible concepts and the critical approach commonly used in the humanities [12, 13]. This typical pluralism of interpretations highlights the need for a less absolutist and more culturally-aware representation of social constructs such as values. An established epistemological foundation of such a formal representation is social constructivism, where knowledge is understood to stem from shared interpretation and understanding. This leads to the idea that ontological entities are grounded in a particular socio-cultural context. A constructivist approach is taken e.g. in *PhiloSurfical* [14] in the construction of their ontology of philosophy. The authors argue that they cater for different interpretations by modelling the ontology at a high level of generality, thus making their representations "as re-usable as possible, especially among annotators having different philosophical views".

Beyond the aforementioned limitations, for the formal representation of values, we note that CIDOC-CRM supports the *appellation* relationship that can be employed for a simplistic to representation between *concepts* (as used in VAST) and their corresponding *terms*. The authors of the *PhiloSurfical* model [14] introduce the notion of *idea-appelation* to denote philosophical concepts. As a further contribution, the *Europeana Data Model* has been recently proposed to represent cultural-heritage resources by adopting the Simple Knowledge Organization System (SKOS) and thus also addressing possible definitions of of concepts [15]. Similarly to CIDOC-CRM, *SKOS* itself is a schema, but it hosts a number of thesauri that contain entities that could be useful as *terms* in our approach. Noteworthy are the GESIS Thesaurus for the Social Sciences[1] and the UNESCO Thesaurus[2].

2.2. Collaborative Knowledge Graph Engineering

When multiple contributors are involved in the specification of an knowledge graph, a clearly defined process becomes a necessity. The *NeON methodology framework* to collaborative ontology engineering provides variations of a process to be used in different circumstances, depending on planned project duration or availability of pre-existing knowledge resources [16]. The simplest among these is a flow of four subsequent steps to align contributors: initiation, design, implementation, and maintenance. Other versions also include reuse, merge and re-engineer phases, or an iterative repetition of the process. Particularly relevant here is the NeON iterative-incremental approach that consists of iteratively repeating the stages of the simple process. Special attention is paid to ensure that no backtracking takes place so that the process converges.

[1] https://www.gesis.org/en/services/research/thesauri-und-klassifikationen/social-science-thesaurus.

[2] https://skos.um.es/unescothes.

As an alternative, the *UPON Lite framework* consists of 6 steps identified by their objectives: terminology, glossary, then taxonomy, predication and parthood performed in parallel and, finally, ontology [17]. These steps can be considered as a detailed expansion of the design and implementation phases in NeON.

The idea of involving non-expert users in ontology engineering has been discussed in the literature by proposing to employ crowdsourcing-based solutions. Daga et al [18] consider how citizen are contributing to cultural heritage archives. However, none of the considered examples are addressing the need to organise diverse knowledge in a single knowledge graph. Highly structured software for ontology construction and crowdsourcing from social networks has been used as a validation mechanism [19]. Games with a purpose (GWAP) have been also proposed as a common solution to employ crowdsourcing contributions for the creation of knowledge graphs. These have been categorised into four types: specification of term relatedness, verification of relationship correctness, specification of relation type, and verification of domain relevance [20]. An example of specification of term relatedness is Free Association [21], in which given a word, the user is asked to find others that are related. A recent example for relationship verification is Indomilando [22] where players are asked to choose between different predefined relationships in the form of multiple-choice questions) as a way to verify that an undisclosed one of them is correct. A typical model for specification of relationship type is a game called SpotTheLink [23] where players are given two concepts and asked to identify a relationship that could connect them.

3. Approach

We describe our participatory knowledge graph design providing two overarching perspectives, one static describing the meta-model and one dynamic describing the process of knowledge externalisation.

3.1. Knowledge Graph Meta-model

Starting from the tangible artifacts and the users interacting and interpreting them, we develop a representation of the activities and the elicited values. This is reflected in the four principle domains of Figure 1 and discussed below.

Artifacts. The knowledge graph that we present here features a unified representation of tangible resources, from historical documents and pieces of art to questionnaires and activity recordings. A common feature of the representation of all these is their association to a specific spatio-temporal context. Whereas to the knowledge graph artifacts are simply stored objects, some of them might also have a role in the knowledge graph creation as structural representations of knowledge, including annotations, questions and answers or digital games provoking particular behaviours, as elaborated in Section 4. Correspondingly, to be able to facilitate a focused reflection or discussion of a particular part of an artifact, the knowledge graph features also a representation of a *segment* of an artifact, be it text, video or a more structured digital tool.

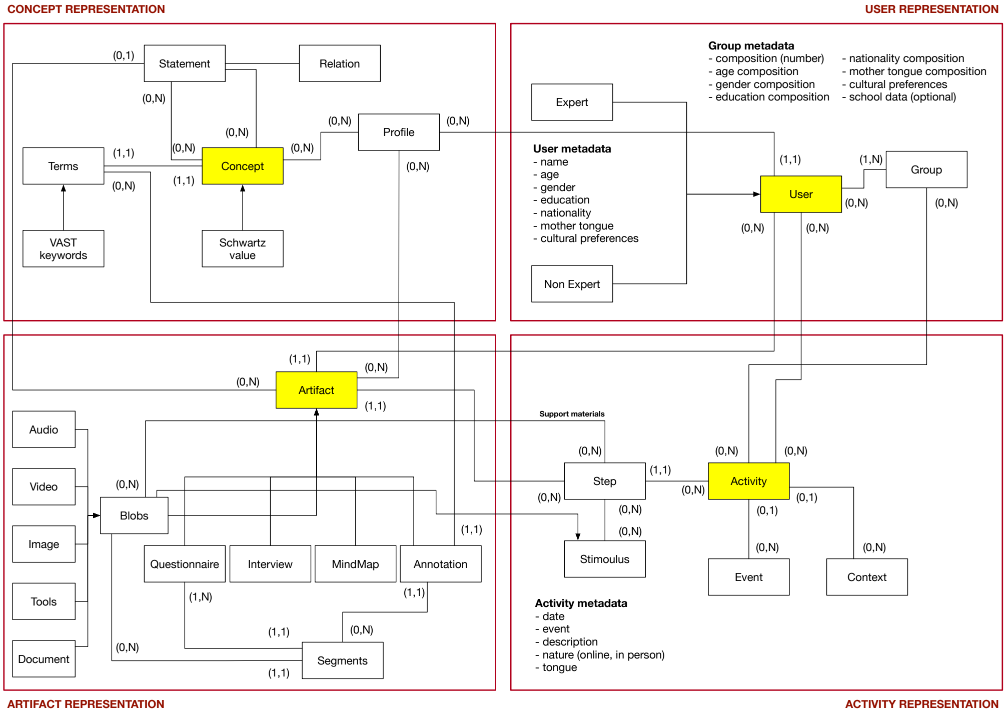

Figure 1: The VAST knowledge graph meta-model showing the four principal domains.

Users. To capture their perceptions of values over space and time, our knowledge representation features two types of users as represented in Figure 1, who could contribute differently: (i) *Expert* users include humanities researchers that study and interpret the historical intentions behind an artifact, social science researchers that study current perceptions of it and value communicators that reenact its contemporary significance. This separation of responsibility needs not be strict, but the contribution of each user has a clearly defined spatio-temporal context, which for expert users is their domain of expertise, as further explained in the Section 3.2. (ii) *Non-expert* users are the general audience participants that engage in VAST activities and share their perceptions and views. Based on an experience (*activity*) related to a historical artifact. This could happen in a museum or theatre visit, in a school activity or online, which defines their corresponding spatio-temporal context.

Activities. Users engage with artifacts in the context of particular activities that also need to be represented in the knowledge graph. This is not only due to the need to store the time (technically – date) and space (geolocation) of when and where these take place, but also their structure. In order to capture the potential influences that might shape the user experience, its steps (i.e. method) and stimuli (i.e. artifacts used as supplementary materials) are being represented. This allows the knowledge graph to relate responses to stimuli from different steps within a single activity, even when the related data is collected in anonymised form. Finally, the context of the activity is represented through the recordings of activity sessions (i.e. *events*), involved organisations, and individual and/or group demographics. Providing an example for a complete activity is out of the scope of this paper. Instead, here we take a data-centric approach, focusing on representation and collection of data. Accordingly, in Section 4 we provide examples of different types of used artifacts and the corresponding digitisation of concepts.

Concepts. To represent values, the VAST knowledge graph features two main entities: concepts and terms. *Concepts*, a core part of which are the Schwartz values, are abstractions that could feature also other commonly established psychological constructs beyond those studied by Schwartz. By adopting standard psychological instruments, such as the European Social Survey [4] (a type of *artifact* for our knowledge graph) in *activities*, these concepts are associated to user profiles. As already indicated, values – and concepts in general – are implicitly represented in communication across different media (e.g., text, visual arts, drama, oral narration). As communication changes in different social contexts, so does the perception of underlying values. As a consequence, we consider *terms* and *concepts* to be omnipresent across space and time, but we consider relationships between them to be related to a specific context, thus pinned to specific time and space. *Terms* represent explicit categories, each defined by a label and associated to different artifacts. The vocabulary of terms includes the values of the EU and an expert-led expansion of these, termed VAST keywords, but potentially also other terms introduced by users. Terms are associated to annotations in artifacts by users (expert and non-expert, as detailed in Section 3.2). Concepts and terms are variously interconnected by relationships. In particular, a binary relationship among a pair of terms/concepts can be specified to denote a semantic relation holding between them. Three particular relationship semantics have been identified as relevant: has-broader-term (semantically equivalent to P127 in CIDOC CRM), shows-features-of (P130), and is-in-conflict-with (no CIDOC CRM equivalent).

3.2. Knowledge Graph Engineering

Our approach to knowledge graph engineering is based on an iterative-incremental design approach inspired by the *NeON* methodology [16] and the *UPON Lite* methodology. NeON in a sentence. UPON Lite has been defined to "*shift responsibility for ontology building toward a community of end users through a social, highly participative approach supported by an easy-to-use method and tools* [17]".. In the first stages of the process, experts collaboratively create the core of the knowledge graph. For each VAST pilot, a team of domain experts composed of researchers and value communicators works to produce vocabulary (i.e. *terms* and *concepts*) and relationships, relevant to the pilot. In this respect, the following sequence of steps have been defined.

Vocabulary definition. Starting from EU values, experts from the three pilots define an extended set of *terms* that they consider relevant to their respective domain, both spatio-temporal and in terms of activities. A single vocabulary of *terms* and *concepts* is defined containing the whole set provided by all pilots, so that a *term* can be used where found appropriate regardless of the activity in which it emerged.

Schema definition. As explained in the previous subsection and shown in Figure 1, a general schema has been defined beforehand. Beyond that, the decisions of how to represent individual *artifacts* and *activities*, as well as how to connect these *terms* are taken by the core team of domain experts representing each of the three pilots.

Schema population. The content generation is reduced to micro-tasks doable within single isolated sessions, in which users are assumed not to have prior awareness and need to be introduced to the entire context. Thus, these sessions (i.e. *activities*) are much closer to field studies in museums, theatre, schools or particular online settings. They represent a highly structured way of schema population akin to citizen curation [18] and games with a purpose [20] designed by domain experts. Particular examples of activities including such micro-tasks with non-experts have been presented elsewhere [24, 25, 26] and an illustration of the corresponding integration of the data in the knowledge graph follows.

4. Digitisation Samples

Our approach allows for the integration of a wide range of user contributions into the knowledge graph as a way to formalise contributor knowledge. To illustrate this, here we discuss five distinct examples that illustrate the full range of structured and unstructured contributions. These are text annotation, questionnaire responses, dedicated social network discussions about values, mind map creation and digital game interactions.

Text Annotation. An established content analysis technique is used to annotate values in classical texts with dedicated annotation tool and methodology [27]. In particular, we work on English translations of the considered textual artifacts and users are asked to identify segments in artifacts and tag them with (one or more) appropriate terms from the vocabulary. In contrast to other activities described below, text annotation is a lengthy process that requires longer periods of engagement, but is also accessible to experts. Users are first asked to annotate explicit references to the considered values in the artifact. Then, they are asked to annotate implicit or indirect references to values, according to their subjective interpretation.

Questionnaires. Three models for mapping answers of common question types to the knowledge graph are supported, namely multiple choice questions, scale questions and open-answer questions. Multiple choice questions allow for formulation of questions aimed specifically at a particular set of alternative relationships, be it between *terms*, *concepts*, *profiles*, *artifacts*, etc. An example of such a question regarding the relationship between a segment of an artifact and a term is provided in Figure 2. These questions can typically be formulated to enquire about relationships proposed by experts and allow probing for their perceived validity across various user groups. Scale questions allow to examine the strength of relationships, for example using Likert-scale questions of standardised questionnaires from psychology [4]. This allows adding a quantitative measure to the relationships provided by an individual. To this end, beyond the information about the relationship supported by the user, also knowledge graph statements containing the question scale and the response rating are stored. Finally, open questions allow for unconstrained free-text responses. We propose two distinct ways to integrate these in the knowledge graph. One is to pose questions that provide two of the elements of the triple representing a relationship in the knowledge graph. Then users are asked about the third as illustrated in Figure 2. With this approach, a statement for the response is included in the knowledge graph only for those responses that can be automatically mapped (e.g. via exact matching to possible appellations) to preexisting objects. Answers that are not reduced to preexisting statements in this way are simply included as new artifacts for future qualitative analysis. Such addition as artifacts is the second (and non-automatically interpretable with the current functionality) way to integrate answers to any form of open questions into the knowledge graph. This is necessary because questionnaires have been designed beforehand without considering digitisation constraints.

Social Network Activity. Using the specific functionalities of social networks, such as polls, images and replies, questionnaire items can be represented online to reach wider audience, as shown in Figure 2. Posting such questions with specific hashtags or in dedicated groups provides a community context remotely similar to the one of a bespoke activity, albeit with limited demographics data.

Mind maps. Another widespread technique to collect opinions, that is both an intuitive representation of complex knowledge and naturally adapts itself to a knowledge graph, is network-based visualisations such as mind maps, and the similar concept maps and topic maps. Such data is being collected in a range of activities within the VAST project (see Figure 3), both on paper to be curated and digitised by educators, and in a dedicated digital tool that allows data to be directly imported in the knowledge base. Typically in such activities, students are asked to start from a concept or term present in the knowledge graph and – in a group or individually – expand the map with their own views. The nodes that are being introduced in the process are either mapped to terms with coinciding labels, or introduced as new terms [24]. A particularity of this type of activities is that more often than not, relationships are unlabeled. In such cases, we employ a general relatedness (shows-features-of) relationship.

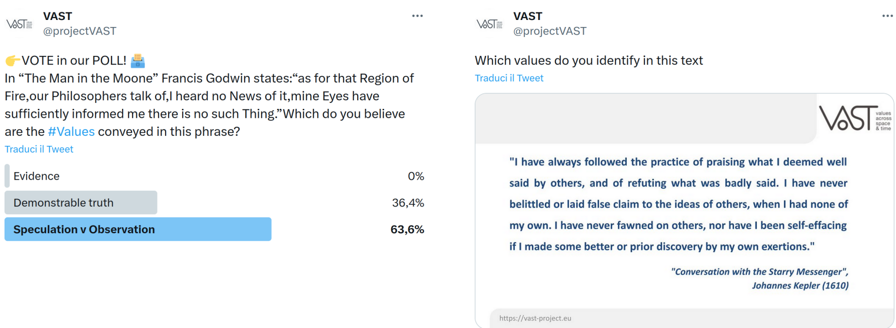

Figure 2: Two examples of responses collected from a social network: a multiple choice question on the left and open answer question to the right.

Figure 3: Examples of mind maps from paper-based and digital (right) activities with students.

Figure 4: The two games that have been integrated in the VAST platform. On the left is The trial of Antigone, and on the right is the Little Red Riding Hood.

Digital Game Experiences. Yet another very broad category of activities are bespoke digital games that provide highly structured experiences and are designed to generate specific representations of player perceptions based on their choices within the game. Two examples that have been developed within the VAST project and shown in Figure 4 are games where players are asked to role-play the Little Red Riding Hood and The trial of Antigone [25, 26] and value preferences are being derived from the player choices within the games.

5. Technical Implementation

The VAST platform is built on top of a GraphDB instance which ensures a RDF-compliant storage and a SparQL querying endpoint. This graph database is used as a data warehouse, synchronised via API endpoints and synchronisation scripts with the various components. These components include a text annotation tool, a survey management tool, and a set of digitisation tools for different activities. Beyond the specific bespoke integration with the digital mind maps and games, also a generic digitisation tool which allows for the import of generic, including non-digital activities. The resulting integration, including content produced from the above-mentioned digitisation types, is accessible via the VAST platform[3].

6. Conclusion and Future Work

In this paper we described an approach to the construction of a knowledge graph of moral values – a particular type of abstract concepts. The necessity to represent both past and present interpretations of the values conveyed in historical artifacts has led us to consider a broad range of contributors. We employ an iterative-incremental process. We also demonstrate a range of mechanisms that engage users to contribute to our knowledge graph.

The produced participatory knowledge graph is made available via the VAST platform. Besides supporting both qualitative and quantitative research (e.g. providing baselines for values mining and natural language processing), this collection of structured datasets can be used for citizen-informed curation, and design of future museum activities.

An important aspect of knowledge graph creation is interoperability with other relevant semantic resources, such as CIDOC-CRM and ValueNet. While not discussed here, this has been envisioned by trying to reuse semantics as much as possible, e.g. relationship types from CIDOC-CRM and shared concept entities with ValueNet. Yet, the actual integration remains to be detailed.

Furthermore, answer matching in the described approaches could be expanded by identifying *concepts* and *terms* with their corresponding appellations from established thesauri. Beyond this, answers to open questions could be further analysed for further semantic interpretation, e.g. identifying responses to generic questions who, how [28], where [29], etc.

Although it is beyond the scope of this paper, an important aspect of the knowledge graph design process are the activities that involve contributors. Whether it is questionnaires or mind map compilation done within conventional museum workshops or self-paced activities in social networks or games, enriching text content with multimedia is expected to contribute to engagement and retention. One way to make this more accessible, that has recently emerged, is through the semi-automatic generation of contextual illustrations. While text-to-image generator models are not commonly grounded in pre-existing text, preliminary tests have shown very positive results [30].

Acknowledgments

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101004949. This document reflects only the author's view and the European Commission is not responsible for any use that may be made of the information it contains.

[3] https://platform.vast-project.eu/

References

[1] K. Giouvanopoulou, X. Ziouvelou, G. Petasis, et al., Exploring Values and Value Transformation: A Multi-Perspective Approach, Open J. of Social Sciences 11 (2023) 313–338. doi:10.4236/jss.2023.113023.

[2] S. Castano, A. Ferrara, G. Giannini, et al., A Computational History Approach to Interpretation and Analysis of Moral European Values: the VAST Research Project, in: HistoInformatics 2021, 2021, pp. 1–9. URL: https://ceur-ws.org/Vol-2981/paper7.pdf.

[3] S. H. Schwartz, J. Cieciuch, M. Vecchione, et al., Refining the theory of basic individual values, J. of personality and social psychology 103 (2012). doi:10.1037/a0029393.

[4] S. H. Schwartz, A repository of schwartz value scales with instructions and an introduction, Online Readings in Psychology and Culture 2 (2021). doi:10.9707/2307-0919.1173.

[5] M. Scharfbillig, L. Smillie, D. Mair, et al., Values and identities-a policymaker's guide, Publications Office of the EU (2021). doi:10.2760/349527.

[6] A. Ferrara, S. Picascia, E. Rocchetti, Augustine of Hippo at SemEval-2023 Task 4, in: 17th Int'l W'shop on Semantic Evaluation (SemEval-2023), ACL, Toronto, Canada, 2023, pp. 1044–1053. doi:10.18653/v1/2023.semeval-1.143.

[7] G. Papadopoulos, M. Kokol, M. Dagioglou, et al., Andronicus of Rhodes at SemEval-2023 Task 4, in: 17th Int'l W'shop on Semantic Evaluation (SemEval-2023), ACL, Toronto, Canada, 2023, pp. 542–548. doi:10.18653/v1/2023.semeval-1.75.

[8] S. De Giorgis, A. Gangemi, R. Damiano, Basic human values and moral foundations theory in valuenet ontology, in: Knowledge Engineering and Knowledge Management, Springer, Cham, 2022, pp. 3–18. doi:10.1007/978-3-031-17105-5_1.

[9] M. Theodoridou, G. Bruseker, M. Daskalaki, et al., Methodological tips for mappings to CIDOC CRM, in: CAA2016: Oceans of Data: 44th Conf. on Computer Applications and Quantitative Methods in Archaeology, Archaeopress Publishing, 2018, pp. 89–102.

[10] B. Smith, W. Ceusters, Ontological realism: A methodology for coordinated evolution of scientific ontologies, Applied Ontology 5 (2010). doi:10.3233/AO-2010-0079.

[11] K. N. Vavliakis, G. T. Karagiannis, P. A. Mitkas, Semantic web in cultural heritage after 2020, 2021. URL: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.492.4352, last accessed 22 August 2023.

[12] L. Da Sylva, Towards linked data: Some consequences for researchers in the social sciences and humanities, Proc. of the Assoc. for Information Science and Technology 55 (2018). doi:10.1002/pra2.2018.14505501011.

[13] L. Jansen, Ontologies for the digital humanities: Learning from the life sciences?, in: Joint Ontology W'shops 2019, Episode V: The Styrian Autumn of Ontology, Graz, Austria, 2019. URL: http://ceur-ws.org/Vol-2518/paper-WODHSA5.pdf.

[14] M. Pasin, E. Motta, Ontological requirements for annotation and navigation of philosophical resources, Synthese 182 (2011). doi:10.1007/s11229-009-9660-3.

[15] V. Charles, A. Isaac, Enhancing the Europeana Data Model (EDM), 2015. URL: https://pro.europeana.eu/files/Europeana_Professional/Publications/EDM_WhitePaper_17062015.pdf, last accessed 15 August 2023.

[16] M. C. Suárez-Figueroa, A. Gómez-Pérez, M. Fernández-López, The NeOn Methodology for Ontology Engineering, Springer, Berlin, Heidelberg, 2012, pp. 9–34. doi:10.1007/

978-3-642-24794-1_2.
[17] A. De Nicola, M. Missikoff, A lightweight methodology for rapid ontology engineering, Communication of the ACM 59 (2016). doi:10.1145/2818359.
[18] E. Daga, L. Asprino, R. Damiano, et al., Integrating citizen experiences in cultural heritage archives: Requirements, state of the art, and challenges, J. on Computing and Cultural Heritage 15 (2022). doi:10.1145/3477599.
[19] V. Iyer, L. M. Sanagavarapu, Y. Raghu Reddy, A framework for syntactic and semantic quality evaluation of ontologies, in: Secure Knowledge Management In The AI Era, Springer, Cham, 2022, pp. 73–93. doi:10.1007/978-3-030-97532-6_5.
[20] G. Wohlgenannt, M. Sabou, F. Hanika, Crowd-based ontology engineering with the uComp Protégé plugin, Semantic Web 7 (2016). doi:10.3233/SW-150181.
[21] D. Vickrey, A. Bronzan, W. Choi, A. Kumar, et al., Online word games for semantic data collection, in: Empirical Methods in NLP, ACL, Honolulu, Hawaii, 2008, pp. 533–542. URL: https://aclanthology.org/D08-1056.pdf.
[22] I. Celino, G. Re Calegari, A. Fiano, Refining Linked Data with Games with a Purpose, Data Intelligence 2 (2020). doi:10.1162/dint_a_00056.
[23] K. Siorpaes, M. Hepp, Games with a purpose for the semantic web, IEEE Intelligent Systems 23 (2008). doi:10.1109/MIS.2008.45.
[24] E. Fani, A. Ferrara, C. Gagliardi, et al., Participatory citizen curation in the VAST project framework, in: ExICE - Extended Intelligence for Cultural Engagement, Bologna, Italy, 2023. doi:10.5281/zenodo.7973656.
[25] E. Aristodemou, S. Montanelli, M. Ruskov, Fairy tales as a swipe card game: Engaging digital audiences in value-centred discourses using storytelling, in: ExICE - Extended Intelligence for Cultural Engagement, Bologna, Italy, 2023. doi:10.5281/zenodo.7973654.
[26] L. Arnellou, M. Dagioglou, D. Katsamori, et al., Digitising intangible cultural heritage: The use case of Antigone, in: ExICE - Extended Intelligence for Cultural Engagement, Bologna, Italy, 2023. doi:10.5281/zenodo.7973658.
[27] G. Petasis, M. Ruskov, A. Gradou, et al., The VAST collaborative multimodal annotation platform: Annotating values, in: Information Systems and Technologies WorldCIST, Springer, Pisa, Italy, 2023. Publication pending.
[28] M. Ruskov, Who and How: Using sentence-level NLP to evaluate idea completeness, in: AI in Education, 2023, pp. 284–289. doi:10.1007/978-3-031-36336-8_44.
[29] S. Gupta, K. Nishu, Mapping local news coverage: Precise location extraction in textual news content using fine-tuned BERT based language model, in: 4th W'shop on NLP and Computational Social Science, ACL, Online, 2020, pp. 155–162. doi:10.18653/v1/2020.nlpcss-1.17.
[30] M. Ruskov, Grimm in Wonderland: Prompt engineering with Midjourney to illustrate fairytales, in: 19th Conf. on Information and Research Science Connecting to Digital and Library Science, 3365, 2023, pp. 180–191. URL: https://ceur-ws.org/Vol-3365/paper6.pdf.