

Detecting the Semantic Shift of Values in Cultural Heritage Document Collections (short paper)

Alfio Ferrara¹, Stefano Montanelli¹ and Martin Ruskov¹

¹Università degli Studi di Milano
Department of Computer Science
Via Celoria, 18 - 20133 Milano, Italy

Abstract

The paper presents the main features and goals of the EU H2020 VAST (*Values Across Space and Time*) project about the transformation of moral values across space and time, with particular emphasis on the core European Values that represent the essential pillars of the EU society. In particular, we discuss the preliminary results obtained by analysing a selected collection of historical documents by employing machine learning techniques. The aim is to classify document annotations and their relationships with values to discover possible shifts in the value meaning when different temporal contexts are considered.

Keywords

semantic shift detection, natural language processing, computational humanities

1. Introduction

The rapid development and diffusion of artificial intelligence techniques and data science approaches enable research in the field of humanities and social sciences to become more and more computational. Various studies are being appearing to exploit artificial intelligence techniques for heritage representation and processing, like for example the analysis of literary texts and historical documents, or the extraction of knowledge from spontaneous contributions provided by people involved in artistic/cultural experiences such as museum visits and theatrical plays [1, 2]. As a consequence, we are witnessing a shift from *digital humanities* to the so-called *computational humanities* research, where the role of artificial intelligence, data science and cutting edge digital technologies is fundamental to achieve research advances and results [3].

In this paper, we present the ongoing experience of VAST (*Values Across Space and Time*), an EU H2020 project providing a concrete example of computational humanities and, more specifically, computational history research (<https://www.vast-project.eu/>). VAST aims to study the transformation of moral values across space and time, with particular emphasis on the

1st Italian Workshop on Artificial Intelligence for Cultural Heritage (AI4CH22), co-located with the 21st International Conference of the Italian Association for Artificial Intelligence (AIXIA 2022). 28 November 2022, Udine, Italy.

✉ alfio.ferrara@unimi.it (A. Ferrara); stefano.montanelli@unimi.it (S. Montanelli); martin.ruskov@unimi.it (M. Ruskov)

🌐 <https://islab.di.unimi.it/team/alfio.ferrara@unimi.it> (A. Ferrara);

<https://islab.di.unimi.it/team/stefano.montanelli@unimi.it> (S. Montanelli);

<https://islab.di.unimi.it/team/martin.ruskov@unimi.it> (M. Ruskov)

🆔 0000-0002-4991-4984 (A. Ferrara); 0000-0002-6594-6644 (S. Montanelli); 0000-0001-5337-0636 (M. Ruskov)

© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

core European Values, such as *freedom, democracy, equality, rule of law, tolerance, dialogue, and dignity* [4]. Across time, from antiquity to modernity, a value represents a message that is communicated through different mediums (e.g., text, visual art, drama, oral narration) and this message can change when the context and the society where citizens live change. As a first goal, VAST aims at representing values and associated messages as they are extracted from documents of the past, such as for example literature texts. Beyond this, the VAST project will study how moral values are communicated and perceived today, by collecting, digitising, and analysing narratives and experiences of both communicators of moral values, like for example artists, directors, culture and creative industry institutions, museum curators, storytellers, educators, and the respective audiences, like spectators, museum visitors, students, and pupils.

In the following, we focus on the VAST approach developed in the project to extract knowledge from a selected collection of annotated documents. In particular, we explored the use of Natural Language Processing (NLP) techniques (i.e., word2vec) to classify the document annotations and their relationships with values to discover possible shifts in the value meaning when different temporal contexts are considered, namely the so-called project pilots. As a contribution, we provide some preliminary results on the VAST dataset to show the possible employment of machine learning techniques to automatically recognise semantic shift of values within a given document corpus.

The paper is organized as follows. Section 2 provides an overview of the VAST project. In Section 3, the VAST approach to semantic shift detection is presented. Preliminary results on the analysis of semantic shift on values are discussed in Section 4. Related work and concluding remarks are finally given in Section 5 and 6, respectively.

2. The VAST project overview

The purpose of VAST is to enhance the metadata of existing digital resources according to moral values in order to track values in space and time and to study how these values are appropriated in different cultural and societal contexts. The project is structured in three pilots, each of them concerned to a specific historical period and characterized by a general context, a narrative type, communication mediums and tangible and intangible assets [5]. The selection of three specific historical periods allow us to navigate the vastness of history making more comprehensible our study on the transformation of values. The space, as well as the time, is a core aspect of the project: VAST mainly focuses on Western Europe.

Pilot A: Ancient Greek Drama is related to values in ancient Greek tragedies and how they are perceived by contemporary theatrical plays and general audiences. The goal is to analyze how the values of the antiquity, that are recognized to be discussed in specific tragedies (e.g., *Lysistrata*, Comedy, 411 BC), are revisited in the present through modern artistic reproductions, such as acting, music, and voice.

Pilot B: Scientific Revolution Texts is related to values in texts of 17th century about natural philosophy and how they are perceived by experts in science museums and museum visitors like students and pupils. The texts considered in the project are mostly about imaginary travel stories or fictional communities of ideal perfection in which the new intellectual achievements were embedded in an imaginary narrative context (e.g., *The Man in the Moone*, Francis Godwin,

1638).

Pilot C: European Folktales is related to values in folktales throughout the History of Europe and how they are perceived by storytelling experts in fairytale museums and museum visitors. Though fictitious, folktales are important simulations of the reality. Moreover, the variability of tales makes them the ideal case study for cross-cultural comparisons on social dynamics, including cooperation, competition, or decision making. The pilot is mainly focused on archetypical stories (e.g., the *Grimms' Fairy Tales*, 1812) and it includes texts from several European countries (i.e., Portugal, Italy, Slovenia, Greece, Cyprus).

3. Semantic shift detection in VAST

The VAST approach to semantic shift of values is based on the use of word embedding techniques to analyse and compare the value interpretations across the three project pilots. The use of embeddings for semantic shift detection is getting more and more attention in the literature, by leveraging the idea that semantically-related words are close to each other in a given embedding space (see Section 5). In the literature, both context-free and contextualised embedding models are proposed. In VAST, we exploit a context-free solution (i.e., Word2Vec), mainly due to the limited extension of the considered document collection and to the availability of the documents as a static corpus without the incremental insertion of new texts throughout time. In our approach, a fine-tuning step is also performed to obtain pilot-specific models that are exploited to enforce the effective comparison of value descriptions across the pilots.

The VAST approach relies on a collection $\mathcal{D} = \{D_A \cup D_B \cup D_C\}$ composed of documents from Pilot A, Pilot B, and Pilot C, respectively. We work only with consolidated English translations of the original texts¹. The approach is articulated as follows (see Figure 1):

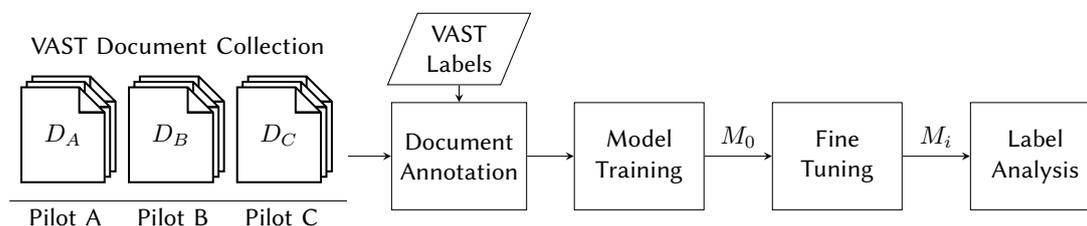


Figure 1: The VAST approach to semantic shift of values

Document annotation. This stage has the goal to associate the VAST documents in \mathcal{D} with labels that are descriptive of the values to recognise. A reference *vocabulary* has been defined in the project and it is composed of around 100 labels about values. A team of scholar experts is involved in the project to perform the annotation task. Each scholar is focused on a single pilot of her/his expertise and she/he has to read the full set of pilot documents and to highlight specific text snippets by associating the vocabulary labels that she/he considers appropriate

¹The full list of documents in the dataset is available at https://contents.islab.di.unimi.it/vastdocs/vast_collection_fulltext.zip.

according to an annotation methodology defined in the project. These expert annotations are included in the documents by inserting the labels used for annotation at the beginning and at the end of the annotated snippets.

Model training. This stage has the goal to train a word embedding model for representing the annotated documents of the collection \mathcal{D} [6]. The document collection is submitted to a lemmatisation process and the word2vec algorithm is then employed [7]. As a result, a word2vec model M_0 is created for the collection \mathcal{D} about all the three project pilots.

Fine-tuning. This stage has the goal to update the model M_0 into three models, each one specific of a pilot and related documents. The goal is to obtain pilot models that are able to capture the language peculiarities of the pilot time-periods. The result of fine-tuning over M_0 is the creation of three pilot models M_A , M_B , and M_C , each one trained by considering the documents in D_A , D_B , and D_C , respectively.

Label analysis. This stage has the goal to exploit the pilot models and to support the cross-pilot comparison of vectors related to target words, namely the values considered in VAST (e.g., justice). Two different analysis based on graph similarity and embedding clustering will be presented in Section 4.

Example. A summary view of pilot documents considered in VAST is provided in Table 1. As an example of document annotation, consider the Brothers Grimm’s version of the Snow-White fairy-tale and the following excerpt (pilot C): “*she realized that the huntsman had deceived her, and that Snow-White was still alive*”. A VAST expert of pilot C annotated this snippet with the label *deceptiveness vs honesty* belonging to the VAST vocabulary.

Table 1

Summary view of VAST pilot documents

Pilot	Docs	Words	Annotations
Pilot A: Greek Tragedy	20	57 578	1788
Pilot B: Scientific Revolution	18	74 291	2098
Pilot C: Fairy-Tales	12	55 623	1692

4. Analysis of semantic shift of values

By relying on the embedding models M_A , M_B , and M_C , in the following, we describe two different analysis where we compare the labels about the VAST values in the three project pilots with the aim to observe possible changes/shifts.

4.1. Similarity Graph of pilot labels

In this analysis, we build a similarity graph for each pilot where the nodes represent the labels of the VAST vocabularies and the edges denote similarities between pairs of nodes according to the cosine similarity calculated by considering the word embeddings (a threshold $\theta = 0.45$ is applied in our experiments to filter poorly-relevant similarity values). The similarity graphs allow to

immediately analyse the shift of a value in the three pilots by observing how the neighborhood changes in the corresponding graphs for a given label. An example of the resulting similarity graphs over the pilots is shown in Figure 2 where 10 labels per pilot are considered.

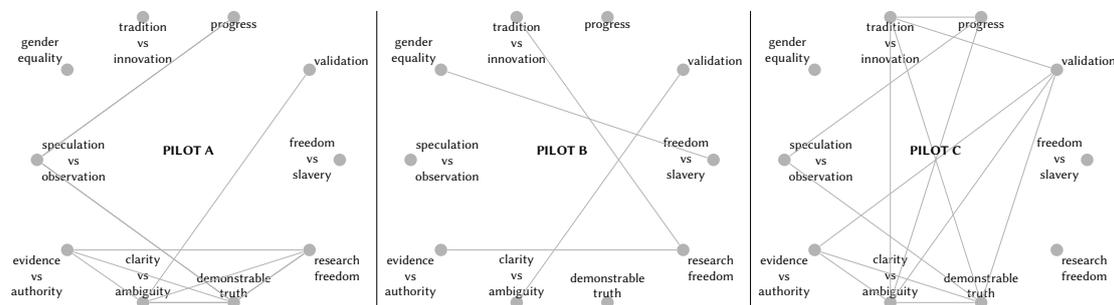


Figure 2: An example of similarity graph over a subset of labels in the VAST vocabulary

In the example, we note that the similarity between validation and clarity vs ambiguity persists in all the three pilots, meaning that such a relationship emerges in the whole dataset. Furthermore, we note that the similarity between gender equality and freedom vs slavery emerges only in pilot B about the Scientific Revolution texts. On the contrary, the similarity between progress and demonstrable truth via speculation vs observation are valid only in Pilots A and C.

4.2. Clustering of pilot labels

In this analysis, we build clusters of similar labels by relying on the similarity links over the labels calculated in the three pilots. In our experiment, the clique percolation method is exploited over the pilot-oriented similarity graphs and three sets of cluster labels are defined (i.e., one cluster-set per pilot). The similarity clusters allow to analyse the shift of a value in the three pilots by observing the overlaps and the differences on the obtained clusters in relation to a given label/value. An example of the resulting clusters over the pilots is shown in Figure 3 where we focus on the free thinking label/value.

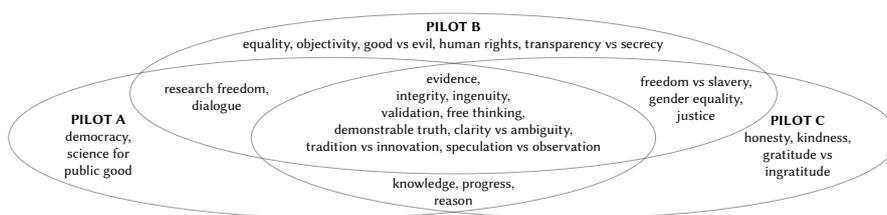


Figure 3: An example of cluster overlap across the three pilots for the label free thinking

The intersection for the three pilots shows the shared theme of the clusters: centred around e.g. validation, free thinking, and integrity. In Pilot A, the label democracy emerges and it is coherent

with the historical period of this pilot (i.e., the Ancient Greece), whereas freedom vs slavery and gender equality are not as prominent as in the other pilots. In Pilot B, the labels objectivity and human rights emerge and this is coherent with the Scientific Revolution period. Finally, honesty and kindness emerge in Pilot C, as typical moral values of the considered folktales.

5. Related work

The proposed VAST approach to shift detection of value meanings is closely related to the more general issue of *semantic shift detection*. In this context, a recent review of approaches is provided in [6] where the authors distinguish between word- and sense-level changes. Typically, word-level approaches focus on detecting changes on a single word meaning that is assumed to be the dominant one, whereas sense-based approaches focus on recognizing changes by considering also the so-called minor meanings. The use of a single, shared embedding, like the one used in VAST, is framed as a word-level approach and it allows to compare the embeddings of different pilots (i.e., sub-corpora) since they are aligned within the same vector space. In [8], a solution based on word2vec and Cosine Similarity is proposed. As a difference with our VAST approach, they consider the time in which the documents are added to the corpus, so they split the corpus in sub-corpora and they train the model on the first sub-corpus by fine-tuning it on the following periods, resulting in a model for each period.

The use of word embedding implies that the document analysis is focused and constrained by the considered corpus used for training. Thus, the results of semantic shift detection only represents a snapshot that depends on the actual semantics of the given vocabulary/corpus [9]. As a consequence, shifts might appear for “somewhat different from what a historical linguist would expect to see” [10]. This requires the capability to interpret the recognised shifts and to classify the word changes according to possible categories like i) words of strongly context-dependent meaning, ii) words frequently used in a very specific context in a particular time bin, and iii) words undergoing syntactic changes, not semantic ones.

As a final remark, solutions for semantic shift detection based on contextual word embedding are also being appearing in the literature (e.g., [11, 12]).

6. Concluding remarks

In this paper, we presented the preliminary results of the VAST project about the shift of values across three different project pilots based on a selected document collection. The obtained results provide interesting suggestions for possible improvements and further investigations.

Ongoing and future work are about i) the enrichment of the document collection with nowadays textual sources collected from non-expert users involved in VAST activities (e.g., museum visitors, theatrical actors/curators), and ii) the creation of an ontological knowledge base about the project pilots derived from the similarity graphs and clusters obtained in our experiment.

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.

References

- [1] H. El-Hajj, M. Valleriani, Cidoc2vec: Extracting information from atomized cidoc-crm humanities knowledge graphs, *Information* 12 (2021). doi:10.3390/info12120503.
- [2] E. Daga, L. Asprino, R. Damiano, M. Daquino, B. D. Agudo, A. Gangemi, T. Kuflik, A. Lieto, M. Maguire, A. M. Marras, D. M. Pandiani, P. Mulholland, S. Peroni, S. Pescarin, A. Wecker, Integrating citizen experiences in cultural heritage archives: Requirements, state of the art, and challenges, *J. Comput. Cultural Heritage* 15 (2022). doi:10.1145/3477599.
- [3] G. Michael, *Agent-Based Modeling and Historical Simulation*, DHQ: Digital Humanities Quarterly 8 (2014).
- [4] The EU values, The EU values, 2020. URL: <https://ec.europa.eu/component-library/eu/about/eu-values/>, last accessed 5 May 2022.
- [5] S. Castano, A. Ferrara, G. Giannini, S. Montanelli, F. Periti, A Computational History Approach to Interpretation and Analysis of Moral European Values: the VAST Research Project, in: *Proc. of the 6th JCDL Int. Workshop on Comp. History (HistoInformatics 2021)*, 2021.
- [6] N. Tahmasebi, L. Borin, A. Jatowt, *Survey of computational approaches to lexical semantic change detection*, Language Science Press, Berlin, 2021, pp. 1–91. doi:10.5281/zenodo.5040241.
- [7] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, 2013. doi:10.48550/ARXIV.1310.4546.
- [8] Y. Kim, Y.-I. Chiu, K. Hanaki, D. Hegde, S. Petrov, Temporal analysis of language through neural language models, in: *Proc. of the ACL 2014 Workshop on Language Technologies and Computational Social Science*, arXiv, 2014, pp. 61–65. doi:10.48550/ARXIV.1405.3515.
- [9] P. Shoemark, F. F. Liza, D. Nguyen, S. Hale, B. McGillivray, Room to Glo: A systematic comparison of semantic change detection approaches with word embeddings, in: *Proc. of the 2019 Conf. on Empirical Methods in Natural Language Processing and the 9th Int. Joint Conf. on Natural Language Processing (EMNLP-IJCNLP)*, Assoc. for Comp. Linguistics, Hong Kong, China, 2019, pp. 66–76.
- [10] A. Kutuzov, E. Velldal, L. Øvrelid, Contextualized embeddings for semantic change detection: Lessons learned, *Northern European J. of Language Technology* 8 (2022).
- [11] D. Schlechtweg, B. McGillivray, S. Hengchen, H. Dubossarsky, N. Tahmasebi, SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection, in: *Proc. of the 14th Workshop on Semantic Evaluation*, Barcelona (online), 2020, p. 1–23.
- [12] F. Periti, A. Ferrara, S. Montanelli, M. Ruskov, What Is Done Is Done: an Incremental

Approach to Semantic Shift Detection., in: Proc. of the Int. Workshop on Computational Approaches to Historical Language Change (LChange), 2022, pp. 33–43.