A complementary account to emotion extraction and classification in cultural heritage based on the Plutchik's theory

Andrea Bolioli‡
Alessio Bosca‡
andrea.bolioli@h-farm.com
alessio.bosca@h-farm.com
H-FARM Innovation, Torino
Torino, Italy

Rossana Damiano**
Antonio Lieto*
Manuel Striani*
rossana.damiano@unito.it
antonio.lieto@unito.it
manuel.striani@unito.it
Dipartimento di Informatica, Università di Torino
Torino, Italy

ABSTRACT

The paper presents a combined approach to knowledge-based emotion attribution and classification of cultural items employed in the H2020 project SPICE. In particular, we show a preliminary experimentation conducted on a selection of items contributed by the GAM Museum in Turin (Galleria di Arte Moderna), pointing out how different language-based approaches to emotion categorization (used in the systems Sophia and DEGARI respectively) can be powerfully combined to cope with both coverage and extended affective attributions. Interestingly, both approaches are based on an ontology of the Plutchik's theory of emotions.

KEYWORDS

Affective Content Aggregation, Commonsense Reasoning, Description Logics

ACM Reference Format:
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 *Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '22 Adjunct), July 4–7, 2022, Barcelona, Spain.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3511047.3537659

1 INTRODUCTION

For centuries, aesthetics has assigned to emotions a primary role in the experience of art; only in recent years, however, this intuition has been confirmed by experiments in neurophysiology, which have demonstrated how correlates of emotions, such as brain response and face expressions, are affected by art ([12, 27]). In addition to their role in the way people relate to artworks ([23]), from paintings and musical works to movies and novels, emotions provide a universal language through which people communicate their experience, well beyond words. Rooted in evolution, emotion are characterized by an universal basis ([7]), despite the differences in their expression across languages and the cultures. In this sense, emotions can provide a way for connecting people who belong to different groups, intended as culture, age, education, and different sensory characteristics. The expression of emotions through language, in particular, lies at the basis of several models of emotions, including Shaver's ([24]) and Plutchik's ([20]), and has prompted the creation of a number of resources for sentiment analysis ([4, 22, 26]). The application of these resources to art is straightforward: for example, WikiArt Emotions ([17]) is a dataset of 4,105 artworks from WikiArt annotated for the emotions evoked in the observer. The artworks were annotated via crowdsourcing for one or more of twenty emotion categories, in English language. Experiments such as WikiArt Emotions have paved the way to the extraction of emotions from text and tags to create affective art recommenders, like ArsEmotica ([1, 18]) or DEGARI ([15], [16]), able to classify and group artistic items well beyond the standard 6 basic emotions of Ekman's theory ([7]), embracing richer, finer-grained models. A recent experiment on emotions evoked by art was performed in the Art Emotions Map project ([25] under review): 1,300 people were asked to describe how 1,500 paintings make them feel by choosing from different words. The results revealed 25 different emotions that people linked to the artworks they saw. The authors plotted these feelings on an interactive map, grouping artworks that triggered specific emotions.

In the context of the project SPICE [3, 6], which aims at supporting citizens in creating and sharing their own interpretation of artworks by attaching personal responses and affective annotations to artworks, our work has been focused on developing knowledge-based and reasoning technologies that leverage the role of emotions in the tasks of interpreting and reflecting on museums exhibits. In particular, we have developed two different complementary strategies to equip museum exhibits with emotional labels from user-generated comments. Emotional labels have been derived from an online campaign where users had to annotate their emotional feelings on a selection of artworks chosen by the curators of the Gallery of Modern Art located in Turin (Galleria di Arte Moderna, GAM), Italy. People were asked to answer the question: "How does this artwork make you feel?". We analyzed the answers using two complementary strategies that rely on the same emotion model (the Plutchik's theory) as instantiated in the Plutchik's ontology that

*All the authors contributed equally to this research.

UMAP '22 Adjunct, July 4–7, 2022, Barcelona, Spain

ACM ISBN 978-1-4503-9232-7/22/07…$15.00
https://doi.org/10.1145/3511047.3537659

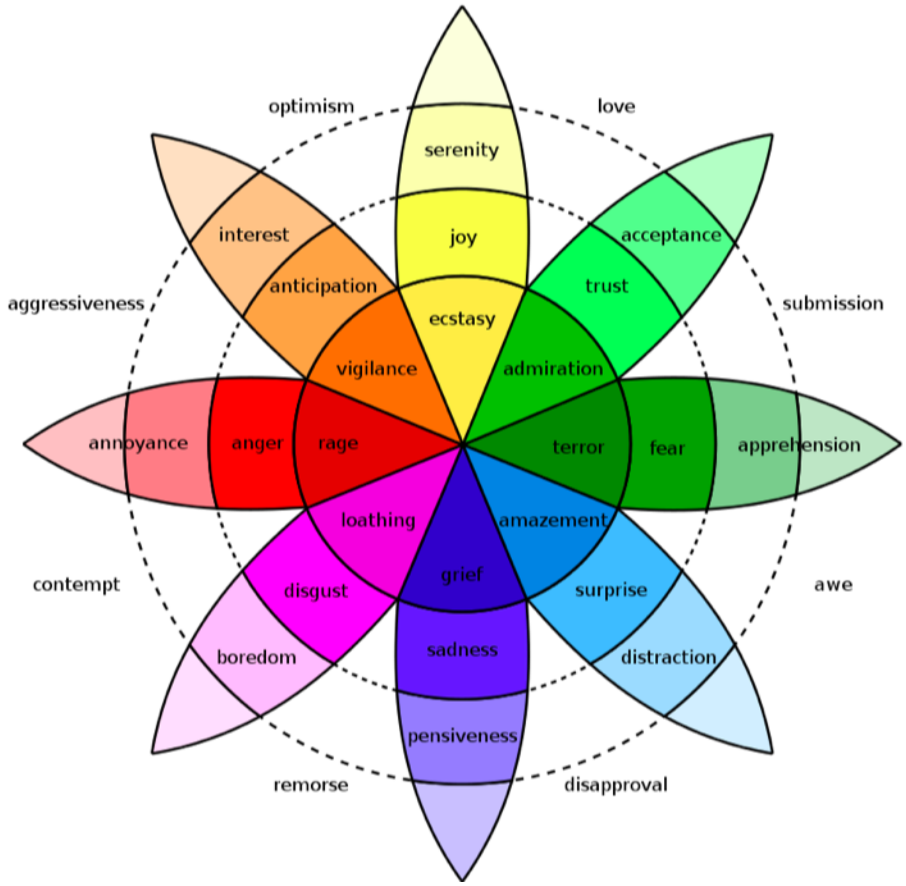

Figure 1: The Wheel of Emotion of the Plutchik Model

will be described below, but work in complementary ways to yield a fine-grained, comprehensive account of the emotions evoked by the artworks.

2 THE PLUTCHIK'S ONTOLOGICAL MODEL

The reference theory for the two systems is encoded in an ontology of emotional categories based on Plutchik's psychological model of emotions [19]. The ontology structures emotional categories in a taxonomy, which currently includes 32 emotional concepts. The design of the taxonomic structure of emotional categories, of the disjunction axioms and of the object and data properties mirrors the main features of Plutchik's circumplex model. As mentioned before, such model can be represented as a wheel of emotions (see Figure 1) and encodes the following elements:

- Basic or primary emotions: *Joy*, *Trust*, *Fear*, *Surprise*, *Sadness*, *Disgust*, *Anger*, *Anticipation*; in the color wheel, this is represented by differently colored sectors.
- Opposites: basic emotions can be conceptualized in terms of polar opposites: *Joy* vs *Sadness*, *Anger* vs *Fear*, *Trust* vs *Disgust*, *Surprise* vs *Anticipation*.
- Intensity: each emotion can exist in varying degrees of intensity; in the wheel, this is represented by the vertical dimension.
- Similarity: emotions vary in their degree of similarity to one another; in the wheel, this is represented by the radial dimension.
- Complex emotions: a complex emotion is a composition of two basic emotions; the pair of basic emotions involved in the composition is called a *dyad*. Looking at the Plutchik wheel, the eight emotions in the blank spaces are compositions of similar basic emotions, called *primary dyads*. Pairs of less similar emotions are called *secondary dyads* (if the radial distance between them is 2) or *tertiary dyads* (if the distance is 3), while opposites cannot be combined.

Within this ontology, the class Emotion is the root for all the emotional concepts. The Emotions hierarchy includes all the 32 emotional categories as distinct labels. In particular, the Emotion class has two disjoint subclasses: BasicEmotion and ComplexEmotion. Basic emotions of the Plutchik model are direct sub-classes of BasicEmotion. Each of them is specialized again into two subclasses representing the same emotion with weaker or stronger intensity (e.g. the basic emotion *Joy* has *Ecstasy* and *Serenity* as sub-classes). Therefore, we have 24 emotional concepts subsumed by the BasicEmotion concept. Instead, the class CompositeEmotion has 24 subclasses, corresponding to the primary (*Love, Submission, Awe, Disapproval, Remorse, Contempt, Aggressiveness e Optimism*), secondary (*Hope, Guilt, Curiosity, Despair, Unbelief, Envy, Cynicism e Pride*) and tertiary (*Anxiety, Delight, Sentimentality, Shame, Outrage, Pessimism, Morbidness, Dominance*) dyads. Other relations in the Plutchik model have been expressed in the ontology by means of object properties: the hasOpposite property encodes the notion of polar opposites; the hasSibling property encodes the notion of similarity and the isComposedOf property encodes the notion of composition of basic emotions. In the following we introduce the SOPHIA Engine and the DEGARI systems, both relying on the model of emotion developed by Plutchik, and show the main differences concerning the obtained affective classifications.

3 SOPHIA ENGINE

Sophia Analytics is a platform for text and data mining combining structured data with more fine-grained information, extracted from natural language contents. In the context of SPICE project it is used in order to realize and expose an annotation service for the semantic enrichment of textual contents, targeting user generated contents as well as descriptions of museum artifacts.

The process of semantic annotation is realized by a Natural Language Processing Pipeline that includes different analysis modules, each one responsible for annotating the document with respect to a specific aspect: sentiment analysis, emotion detection, entity linking. The overall process is exposed by means of standard RESTful APIs and produces a JSON-LD document as output. The service is multilingual and supports English, Finnish, Hebrew, Italian and Spanish.

The analysis performed by Sophia Annotation Service makes it possible to focus on the visitors, their thoughts, cultural and social context, emotional inclinations so to enhance their role in the curatorial process, both as individuals and as part of a community (or more communities). It also allows for retrospective social studies on how the same type of content can produce different emotions and polarities and, also, how the same emotion or object interpretation is instead shared by people belonging to different communities.

The multilingual Emotion Detection component for the Art domain combines language specific domain knowledge (SPICE Emotion Lexicon and rules) with state-of-the-art AI models that allow for tailoring the system to the domain, jargon and style of final users. For each input text, the emotions detected by the 2 components (rule-based and deep learning based) are combined in a single result.

The component references emotions from the Plutchik Emotion ontology. In the first classification experiment described in this paper (called SPICE Art&Emotions experiment) the emotions supported are a subset of the complete model and consist of Anger, Anticipation, Disapproval, Disgust, Fear, Interest, Joy, Love, Sadness, Serenity, Surprise, Trust. i.e., 8 basic emotions + 2 positive emotions (Interest and Serenity) + 2 complex emotions (Disapproval and Love). This emotions tagset derives from the union of several emotion lexicons for English, Italian and other languages, as described below. Therefore, it is not meant to represent a complete model of emotions but a linguistic resource that can evolve over time.

3.1 SPICE Emotion Lexicon and Rule Based Emotion Detection

For the creation of the multilingual lexicon "SPICE Emotion Lexicon", we started with a subset of the Italian Emotion Lexicon developed by CELI (9,321 entries) as the source dataset ([2]). This lexicon was then integrated with words taken from a subset (6,468 entries) of the NRC Word-Emotion Association Lexicon (Mohammad S. M., & Turney, P. D., 2013). The Lexicon was also integrated with a list of Italian emotive words (555 entries) contained in the ItEM lexicon (Passaro, L. et al., 2015). The final output consist of a multilingual aligned lexicon for English (1,865 entries), Italian (2,483 entries) and Spanish (1,795) and singular lexicons for Hebrew (1,003 entries) and Finnish (5,836 entries).

The lexicon comprises as well a set of 122 emojis associated with emotions, which seem to be the most used and widespread on social media. The emojis are taken from Emojipedia. The association between the emoji and the emotion has been made both arbitrarily and on the basis of previous studies on emotion detection in the emoji field (Wolny, W. (2016); Shoeb, A.A. et al (2019); Arva, H. et al (2018)).

A rule-based system identifies patterns as a sequence of words (enriched with morphological features, as lemma or part of speech). The rules are automatically generated combining the entries of the emotion lexicon with language specific rules (used for handling conjunctions, modifiers as adverbs, negations) manually defined by linguists.

The rule engine used for applying these rules to the sequence of analyzed terms is Drools, a rule management system with a forward and backward chaining inference-based rules engine, also known as a production rule system using an enhanced implementation of the Rete algorithm. More details on the Rete algorithm implementation used in Drools can be found in: Proctor, Mark, et al. "Drools documentation." JBoss 5.05 (2008): 2008.

3.2 Deep Learning based approach

A Language Model (LM) assigns probabilities to a sequence of words and is a crucial component in NLP applications such as machine-translation and information extraction. In the last years, the Deep Learning era has brought new neural LMs (as Bert or GPT-3) that have outperformed the traditional statistical ones in many NLP tasks. Deep Learning Neural Language Models are pretrained on very large corpora of textual data (typically extracted from the web) on unsupervised tasks as predicting the next word in a text or filling the blank. There are several benefits in using a pretrained model, but the most important one is the possibility of fine-tuning it on a specific task with a (relatively) small amount of domain-related data. Thanks to the capability of abstracting and generalizing contents, this type of LM is suitable for dealing with contents coming from users with different language skills (e.g., native speakers, non-native speakers, kids, tourists, etc) and is an effective solution for harmonizing linguistic differences between the users' groups.

Another important benefit comes from Multilingual Neural Language Models in which the tokens from different languages share the same embedding space, thus the experience (annotated data)

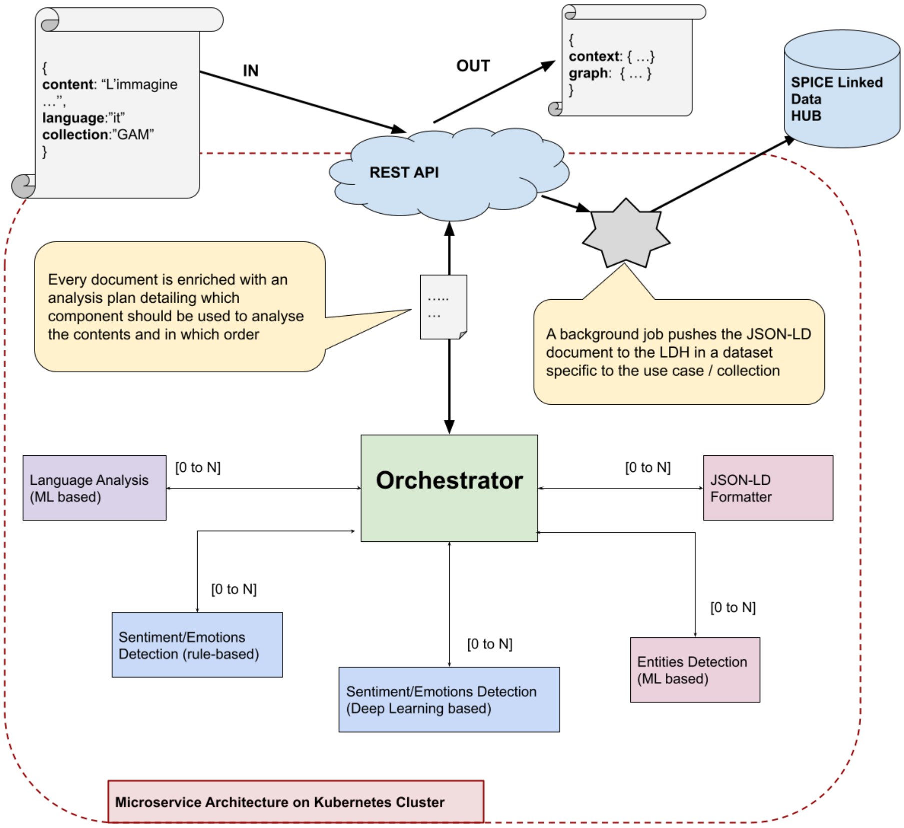

Figure 2: Sophia Analysis Pipeline

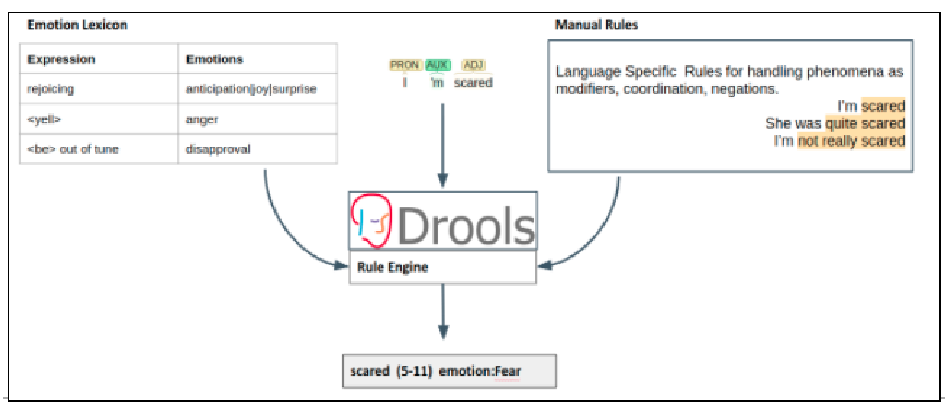

Figure 3: Emotion Detection Rule Based System

learned in one language will be exploited as well in the other languages, leveraging the transfer learning capabilities of the model.

An AI model for Emotion / Sentiment detection in the Art domain based on a pre-trained Deep Learning LM has been trained leveraging the data collected in SPICE use cases and some data from the GoEmotions public dataset (in order to handle emotions that were under-represented in the SPICE dataset). The pretrained LM we adopted is bert-multilingual-base-cased16 from HuggingFace repository. The Bert model (pre-trained on a large amount of data from over 100 languages) has been fine-tuned using annotated data. The iterative model creation process was performed through Sophia Analytics platform; the model created was finally exposed as a microservice in Sophia Analysis Pipeline. As we will continue to work on this ML approach in the coming months, a comparison of the results of the two approaches (rule-based vs. deep learning based) will be presented at the end of the project.

Figure 4: Model Creation Process

4 DEGARI

The core component of DEGARI relies on a probabilistic extension of a typicality-based Description Logic called $\mathbf{T}^{\mathsf{CL}}$, (Typicality-based Compositional Logic, introduced in [14]). This framework allows one to describe and reason upon an ontology with commonsense (i.e. *prototypical*) descriptions of emotional concepts, as well as to dynamically generate novel prototypical concepts in a knowledge base as the result of a human-like recombination of the existing ones [8, 13].

The logic $\mathbf{T}^{\mathsf{CL}}$, that we recall here for self-containedness, is the result of the integration of two main features: (i) an extension of a nonmonotonic Description Logic of typicality $\mathcal{ALC} + \mathbf{T_R}$ introduced in [9, 10] with a distributed semantics; (ii) a well-established heuristics inspired by cognitive semantics for concept combination and generation ([11]), in order to formalize a dominance effect between the concepts to be combined: for every combination, it distinguishes a HEAD, representing the stronger element of the combination, and a MODIFIER. The basic idea is to extend an initial knowledge base (ontology) with a prototypical description of a novel concept, obtained by the combination of two existing ones, namely a HEAD concept and a MODIFIER concept. In this logic, typical properties can be directly specified by means of a *typicality operator* $\mathbf{T}$ enriching the underlying Description Logic, and a knowledge base can contain inclusions of the form $p\ ::\ \mathbf{T}(C) \sqsubseteq D$ to represent that typical Cs are also Ds, where p is a real number between 0.5 and 1, representing the probability of finding elements of C being also D. From a semantic point of view, it considers models equipped by a preference relation among domain elements, where $x < y$ means that x is more normal than y, and that the typical members of a concept C are the minimal elements of C with respect to this relation. An element x is a *typical instance* of a given concept C if x belongs to the extension of the concept C, written $x \in C^{\mathcal{I}}$, and there is no element in $C^{\mathcal{I}}$ more normal than x. $\mathbf{T}^{\mathsf{CL}}$ also considers the key notion of *scenario*. Intuitively, a scenario is a knowledge base obtained by considering all rigid properties as well as all ABox facts, but only a subset of typicality properties. To this aim, it considers an extension of the Description Logic $\mathcal{ALC} + \mathbf{T_R}$ based on the distribution semantics known as DISPONTE ([21]). The idea is to assume that each typicality inclusion is independent from each other in order to define a probability distribution over scenarios: roughly speaking, a scenario is obtained by choosing, for each typicality inclusion, whether it is considered as true of false. Reasoning can then be restricted to either all or some scenarios. $\mathbf{T}^{\mathsf{CL}}$ equips each scenario with a probability, easily obtained as the product, for each typicality inclusion, of the probability p in case the inclusion is involved, $(1 - p)$ otherwise. It immediately follows that the probability of a scenario introduces a probability distribution over scenarios, that is to say the sum of the probabilities of all scenarios is 1.

The bridge from the definition of the emotions in the ontology and the annotations associated with an artwork is provided by an emotion lexicon. Emotional concepts are described by using the NRC Emotion Intensity Lexicon [22] (one of the lexica used also by SOPHIA). Such lexicon provides a list of English words, each with real-values representing intensity scores for the eight basic emotions of Plutchik's theory. The lexicon contains close to 10,000 words, including terms already known to be associated with emotions as well as terms that co-occur in Twitter posts that convey emotions. The intensity scores were obtained via crowd-sourcing, using best-worst scaling annotation scheme. In this work, we considered the most frequent terms available in such lexicon (and associated to the basic emotions of the Plutchik wheel) as typical features of such emotions. In this way, once the prototypes

Total GAM artworks	24	
Total DEGARI artworks classification	12	(50%)
Total SOPHIA artworks classification	24	(100%)
% of items where DEGARI classification extended SOPHIA classification	29%	

Table 1: The Figure shows the aggregate statistics on the 24 selected GAM artworks. SOPHIA classifies all 24 GAM artworks (100%) while DEGARI 12 (50% of the total), extending SOPHIA's coverage with compound emotions. It is worth-noticing that DEGARI is only able to classify complex emotions, while Sophia classifies both basic and (a subset of) complex emotions.

of the basic emotional concepts were formed, the $\mathbf{T}^{\mathrm{CL}}$ reasoning framework was used to generate the compound emotions.

In the context of our system, $\mathbf{T}^{\mathrm{CL}}$ allows us to provide a formal, explainable framework for combining prototypical descriptions of concepts. It is adopted to automatically build the prototypical representations of the compound emotions according to the Plutchik's theory. The prototypes of basic emotions are formalized by means of a $\mathbf{T}^{\mathrm{CL}}$ knowledge base, whose TBox contains both *rigid* inclusions of the form

$$BasicEmotion \sqsubseteq Concept,$$

to express essential desiderata but also constraints, e.g. $Joy \sqsubseteq PositiveEmotion$ as well as *prototypical* properties:

$$p \ :: \ \mathbf{T}(BasicEmotion) \sqsubseteq TypicalConcept,$$

representing typical concepts of a given emotion, where $p \in (0.5, 1]$, expressing the frequency of such a concept in items belonging to that emotion: for instance, $0.72 \ :: \ \mathbf{T}(Surprise) \sqsubseteq Delight$ is used to express that the typical feature of being surprised contains/refers to the emotional concept *Delight* with a frequency/probability/degree of belief of the 72%.

Once the association of lexical features to the emotional concepts in the Plutchik's ontology is obtained and the compound emotions are generated via the logic $\mathbf{T}^{\mathrm{CL}}$, the system is able to reclassify the artworks in the novel emotional categories. Intuitively, an item belongs to the new generated emotion if its metadata (name, description, title, user-generated annotations) contain all the rigid properties as well as at least the 30% of the typical properties of such a derived emotion. The 30% threshold was empirically determined: i.e., it is the percentage that provides the better trade-off between overcategorization and missed categorizations [5].

Overall, once the association of lexical features to the emotional concepts in the ontology is obtained and the hybrid emotion concepts have been generated via the logic $\mathbf{T}^{\mathrm{CL}}$, the system is able to reclassify the artworks through the textual descriptions associated to that cultural items.

More specifically, the current version of the system, available as a web service, accepts JSON files containing a textual description of the artworks (e.g. from user comments or from the museum catalogues) and performs an automatic information extraction step generating a lemmatized version of the JSON descriptions of the artworks and a frequentist-based extraction of the typical terms associated to each artwork in its textual descriptions (the assumption is that the most frequently used terms to describe an item are also the ones that are more typically associated to it). The frequencies are computed as the proportion of each term with respect to the set of all terms characterizing the item, in order to compare the lemmatized version of the artwork description with the prototypes of the compound emotions generated. These two tasks are performed by using standard libraries like Natural Language Toolkit [1] and TreeTagger [2]. Once this pre-processing step is automatically done, the final representation of the artwork is compared with the representations of the typical compound values obtained with $\mathbf{T}^{\mathrm{CL}}$.

5 DISCUSSION AND FUTURE WORKS

We have compared the results of the two systems on a subset of selected items provided by the GAM museum curators. We collected answers, comments and tags from users in English and Italian in order to trigger the affective classification using the two approaches described above. How many different emotions can visual art evoke in viewers? How are these emotions verbalized? The question we asked ("How does this artwork make you feel? Write your feelings, emotions, thoughts") was answered in different ways, even in front of the same work of art. For example, looking at The Siren by G. A. Sartorio one answer was "Love, romantic, calm, a bit sensual. Does it suppose to be a sad ending?", the second one was "Serene, curious, happy", the third one was "I feel anxiety", etc. A precise textual analysis of the answers will be presented in a forthcoming paper, together with the collected dataset.

In this paper we present the overall results, which are shown in Table 1, where we have grouped the emotions extracted from all the answers relating to each picture. Importantly, Sophia shows a better coverage of the itemset while DEGARI (which focus exclusively on complex emotions) is able to perform a fine-grained emotional classification. To this end, Tables 2 and 3 provide both a detailed and synthetic description of the way in which the emotional nuances detected by DEGARI extend the basic ones detected by SOPHIA. In particular Table 2 shows, for the subset of artworks where DEGARI is able to label the user comments with complex emotions, the difference with the affective classification reported by Sophia. In the column reporting the SOPHIA results, in bold, are highlighted the complex emotions (DYADS) extracted by SOPHIA. Table 3, finally, reports the overlap of the two affective systems for the subset of artworks considered in Table 2.

Looking at Table 2, we see that 12 emotion categories have been included for "Aracne", 11 for Asphissia, 10 for The Siren (from Anger to Trust). This result doesn't mean that the automatic system

[1] https://www.nltk.org/
[2]https://www.cis.uni-muenchen.de/ schmid/tools/TreeTagger/

GAM Artefact	SOPHIA emotions	DEGARI emotions
 Asphissia	Anger Anticipation **Disapproval** Disgust Fear Interest Joy **Love** Sadness Surprise Trust	Aggressiveness Anxiety Cynism Delight Disapproval Guilt Hope Love Morbidness Optimism Pride
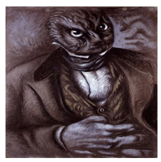 Self-portrait of Owl	Anticipation **Disapproval** Disgust Fear Interest Sadness Surprise	Awe Curiosity Delight Disapproval Outrage Unbelief
 Daphne	Anger Anticipation Interest Joy **Love** Sadness Serenity Surprise	Delight Guilt Love Morbidness Optimism Pride Disapproval
 The Siren	Anger Anticipation Fear Interest Joy **Love** Sadness Serenity Surprise Trust	Anxiety Awe Shame Submission
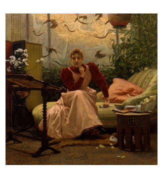 Aracne	Anger Anticipation **Disapproval** Disgust Fear Interest Joy **Love** Sadness Serenity Surprise Trust	Awe Curiosity Delight Disapproval Outrage Unbelief
 The mirror of life	Anger Anticipation Disgust Interest Joy Serenity Trust	Delight Guilt Love Morbidness Optimism Pride

Table 2: Simple and complex emotions extracted by SOPHIA compared with the complex emotions extracted by DEGARI (that - as mentioned - only classifies this category of emotions). In bold are the complex emotions extracted by SOPHIA. As described along the paper, SOPHIA has a better coverage of the overall extracted emotions, while DEGARI is more nuanced in assigning the complex emotions of the Plutchik's wheel.

GAM item	DEGARI emotion classification	SOPHIA emotion classification	DEGARI ∩ SOPHIA
Asphissia, Angelo Morbelli	11	11	14%
Self-portrait of Owl by Alberto Savino	6	7	16%
Daphne by Felice Casorati	7	8	14%
The Siren by Giulio Aristide Sartorio	4	10	0%
Aracne by Carlo Stratta	6	12	16%
The mirror of life by Giuseppe Pellizza da Volpedo	6	7	0%

Table 3: The Figure shows, for the 6 selected GAM artworks (the operas are listed in Table 2), the complex emotions extracted by DEGARI extending the overall emotions (basic + complex) extracted by SOPHIA. The last column shows the overlap percentage between the emotions extracted for both emotional classification systems.

faces overcategorization issues: the manual annotation confirmed that users expressed different emotions in front of the same work.

Interestingly, the results provided by the two complementary approaches (SOPHIA and DEGARI) are cumulative in nature. They, in fact, contribute to enrich the same knowledge graph that is queried to retrieve emotional concepts associated, by means of users answers, comments and tags, to the set of cultural items.

Overall, the results of these experiments suggested a combined use of the two systems to enrich the affective knowledge. In particular, as SOPHIA has a wide coverage in detecting basic emotions, DEGARI can use the output of SOPHIA (basic emotions) in order to identify fine-grained (complex) emotions. This combination is technically obtained by automatically updating knowledge graphs of cultural items available in a Linked Data Format [6] More importantly, such enriched knowledge can be used to feed personalized recommendations to citizens exploiting cultural items in museum.

We plan to extend the evaluation carried out in this preliminary work to artworks from the collections of the other museum partners of the SPICE project, i.e., the Hecht Museum in Haifa, the IMMA Museum in Dublin, the Design Museum in Helsinki and the Museum of National Science in Madrid. We also plan to extend the analysis to other languages, i.e. Finnish, Hebrew and Spanish.

FOUNDING

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] F. Bertola and V. Patti. 2016. Ontology-based affective models to organize artworks in the social semantic web. *Information Processing & Management* 52, 1 (2016), 139–162. https://doi.org/10.1016/j.ipm.2015.10.003
[2] Andrea Bolioli, Federica Salamino, and Veronica Porzionato. 2013. Social Media Monitoring in Real Life with Blogmeter Platform. *ESSEM@ AI* IA* 1096 (2013), 156–163.
[3] L.E. Bruni, E. Daga, R. Damiano, L. Diaz, T. Kuflik, A. Lieto, A. Gangemi, P. Mulholland, S. Peroni, S. Pescarin, et al. 2020. Towards advanced interfaces for citizen curation. In *AVI²CH@ AVI*.
[4] E. Cambria, Y. Li, F.Z. Xing, S. Poria, and K. Kwok. 2020. SenticNet 6: Ensemble application of symbolic and subsymbolic AI for sentiment analysis. In *ACM CIKM '20*. ACM, 105–114. https://doi.org/10.1145/3340531.3412003
[5] Eleonora Chiodino, Davide Di Luccio, Antonio Lieto, Alberto Messina, Gian Luca Pozzato, and Davide Rubinetti. 2020. A Knowledge-Based System for the Dynamic Generation and Classification of Novel Contents in Multimedia Broadcasting. In *Proceedings of ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain, 2020 (Frontiers in Artificial Intelligence and Applications, Vol. 325)*. IOS Press, 680–687. https://doi.org/10.3233/FAIA200154
[6] E. Daga, L. Asprino, R. Damiano, B. Diaz Agudo, A. Gangemi, T. Kuflik, A. Lieto, et al. 2022. Integrating citizen experiences in cultural heritage archives: requirements, state of the art, and challenges. *ACM J. on Computing and Cultural Heritage (JOCCH)* (2022).
[7] P. Ekman and W.V. Friesen. 1971. Constants across cultures in the face and emotion. *Journal of personality and social psychology* 17, 2 (1971), 124.
[8] Laura Giordano, Valentina Gliozzi, Antonio Lieto, Nicola Olivetti, and Gian Luca Pozzato. 2020. Reasoning about Typicality and Probabilities in Preferential Description Logics. *Applications and Practices in Ontology Design, Extraction, and Reasoning, arXiv preprint arXiv:2004.09507* (2020).
[9] L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. 2014. Rational Closure in SHIQ. In *Informal Proc. of Description Logics, Vienna, Austria, July 17-20, 2014 (CEUR Workshop Proceedings, Vol. 1193)*. CEUR-WS.org, 543–555. http://ceur-ws.org/Vol-1193/paper_20.pdf
[10] L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. 2015. Semantic characterization of rational closure: From propositional logic to description logics. *Artificial Intelligence* 226 (2015), 1–33. https://doi.org/10.1016/j.artint.2015.05.001
[11] J. A. Hampton. 1987. Inheritance of attributes in natural concept conjunctions. *Memory & Cognition* 15, 1 (1987), 55–71.
[12] H. Leder, G. Gerger, D. Brieber, and N. Schwarz. 2014. What makes an art expert? Emotion and evaluation in art appreciation. *Cognition and Emotion* 28, 6 (2014), 1137–1147.
[13] Antonio Lieto. 2021. *Cognitive design for artificial minds*. Routledge.
[14] A. Lieto and G.L. Pozzato. 2020. A description logic framework for commonsense conceptual combination integrating typicality, probabilities and cognitive heuristics. *Journal of Experimental and Theoretical Artificial Intelligence* 32, 5 (2020), 769–804. https://doi.org/10.1080/0952813X.2019.1672799
[15] A. Lieto, G.L. Pozzato, S. Zoia, V. Patti, and R. Damiano. 2021. A commonsense reasoning framework for explanatory emotion attribution, generation and reclassification. *Knowledge Based Systems* 227 (2021), 107166. https://doi.org/10.1016/j.knosys.2021.107166
[16] Antonio Lieto, Gian Luca Pozzato, Manuel Striani, Stefano Zoia, and Rossana Damiano. 2022. Degari 2.0: A Diversity-Seeking, Knowledge-Based, Explainalable, and Affective Art Recommender for Social Inclusion. *Submitted* (2022).
[17] S. Mohammad and S. Kiritchenko. 2018. WikiArt Emotions: An Annotated Dataset of Emotions Evoked by Art. In *Proc. of (LREC 2018)*. Miyazaki, Japan. https://www.aclweb.org/anthology/L18-1197
[18] V. Patti, F. Bertola, and A. Lieto. 2015. Arsemotica for arsmeteo. org: Emotion-driven exploration of online art collections. In *Proc. of FLAIRS 2015*. AAAI Press, 288–293. http://www.aaai.org/Library/FLAIRS/flairs15contents.php
[19] R. Plutchik. 1980. A general psychoevolutionary theory of emotion. In *Theories of emotion*. Elsevier, 3–33.
[20] R. Plutchik. 2001. The nature of emotions. *American scientist* 89, 4 (2001), 344–350.
[21] F. Riguzzi, E. Bellodi, E. Lamma, and R. Zese. 2015. Reasoning with Probabilistic Ontologies. In *Proc. of IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*. AAAI Press, 4310–4316. http://ijcai.org/Abstract/15/613
[22] M. Saif. 2018. Word Affect Intensities. In *Proc. of LREC 2018, Miyazaki, Japan, May 7-12, 2018*. http://www.lrec-conf.org/proceedings/lrec2018/summaries/329.html
[23] I. Schindler, G. Hosoya, W. Menninghaus, U. Beermann, V. Wagner, M. Eid, and K.R. Scherer. 2017. Measuring aesthetic emotions: A review of the literature and a new assessment tool. *PloS one* 12, 6 (2017), e0178899.

[24] P. Shaver, J. Schwartz, D. Kirson, and C. O'connor. 1987. Emotion knowledge: further exploration of a prototype approach. *Journal of Personality and Social Psychology* 52, 6 (1987), 1061.
[25] D. Stamkou, E.and Keltner, R. Corona, E. Aksoy, and A.S. Cowen. 2022. Emotional Palette: A Computational Mapping of 25 Emotional Experiences Evoked by Visual Art. (2022).
[26] Y. Susanto, A.G. Livingstone, B.C. Ng, and E. Cambria. 2020. The hourglass model revisited. *IEEE Intelligent Systems* 35, 5 (2020), 96–102.
[27] Noah NN Van Dongen, Jan W Van Strien, and Katinka Dijkstra. 2016. Implicit emotion regulation in the context of viewing artworks: ERP evidence in response to pleasant and unpleasant pictures. *Brain and Cognition* 107 (2016), 48–54.