Annotating Homeric Emotions by a Domain-Specific Language

Federico Boschetti[1], Laura Chilla[2], Maria Konstantinidou[3] and John Pavlopoulos[4]

[1]CNR-ILC
[2]Ca' Foscari University of Venice, Italy
[3]Laboratory of Papyrology and Palaeography, Department of Greek Philology, Democritus University of Thrace, Greece
[4]Athens University of Economics and Business, Greece

Abstract

In this study, we introduce a novel approach to understanding the emotive content present in ancient literature, specifically focusing on the first Book of Homer's *Iliad*. Recognizing the challenges inherent in interpreting emotions from ancient texts, we developed a Domain-Specific Language (DSL) tailored for this purpose. This DSL not only allows for the annotation of basic sentiments such as positive, negative, neutral, or mixed but also facilitates the identification and categorization of specific emotions. To ensure the relevance and applicability of our annotations, we mapped the extracted emotions to some authoritative domain ontologies. This mapping process aids in bridging the gap between ancient emotional expressions and contemporary understanding. Our preliminary results, which we discuss in detail, highlight the potential of our approach in offering deeper insights into the emotional landscape of ancient texts. We believe that our methodology can serve as a foundation for future studies aiming to decode emotions in historical literature.

Keywords

sentiment analysis, digital philology, collaborative annotation, domain-specific languages

1. Introduction

This study follows the investigations of Pavlopoulos [1] for the annotation of sentiment and emotions in the first Book of Homer's *Iliad*, translated in modern Greek. In this second step, the ancient Greek text is analysed and the main focus is on the expressiveness of the annotation system to capture multiple aspects [2] of the textual units under observation.

Like in the previous work, annotators are asked to indicate both the sentiment (i.e. positive, negative, neutral or mixed) and the specific emotions (from an open set of possibilities). But they can annotate at any level of granularity (from a single word to several verses), both on the paradigmatic (i.e. words outside context) and the syntagmatic (i.e. textual units in context) axes, from the perspective of different experiencers (e.g. the character and the ancient audience), towards different participants to the scene.

CLiC-it 2023: 9th Italian Conference on Computational Linguistics, Nov 30 — Dec 02, 2023, Venice, Italy
federico.boschetti@ilc.cnr.it (F. Boschetti); lauradchilla@gmail.com (L. Chilla); mkonst@helit.duth.gr (M. Konstantinidou); annis.pavlo@gmail.com (J. Pavlopoulos)
https://www.ilc.cnr.it/people/federico-boschetti-2/ (F. Boschetti); https://helit.duth.gr/author/mkonst/ (M. Konstantinidou); https://ipavlopoulos.github.io/ (J. Pavlopoulos)
0000-0002-7810-7735 (F. Boschetti); 0009-0009-0160-4467 (L. Chilla); 0000-0002-8744-1444 (M. Konstantinidou); 0000-0001-9188-7425 (J. Pavlopoulos)

2. Related work

Sprugnoli et al. [3] apply sentiment analysis to the *Odes* of Horace at sentence level. One of the main interests is to test *LatinAffectus* [4] the lexical resource for Latin in which words are associated to their polarity: positive, neutral, negative, or mixed. Concerning the ancient Greek Literature, Yeruva et al. [5] studied the inter-annotation agreement of human annotators and machines on an English translation of Aeschylus's tragedies. Luraghi and Sausa [6] study the construal of emotions in Homeric verbs.

Annotating emotions mentioned in ancient literary texts is a complex task because there are no native speakers (as pointed out by Sprugnoli about LatinAffectus), there is not a perfect match between emotions defined in different languages and cultures, and there is abundant secondary literature to take into account (such as exegetical commentaries, lexica, translations).

Sini et al. [7] demonstrate that different languages (in their case: Neo-Latin) structure the lexicon of the emotions in similar but not identical ways.

Kleinginna and Kleinginna [8] and Williams et al. [9] discuss multiple definitions for the category of emotions and suggest how to merge or harmonise different tables of them.

Studies on the emotions in the ancient world are necessary to keep the correct historical distance between the contemporary audience or the annotators and the text or the ancient audience: for instance, Braund and Most [10] on ancient anger (text) or Becker [11] on Stoic emotions

(ancient, but late, audience).

3. Method

A group of 15 volunteers (students and scholars) are asked to annotate the first book of the Iliad on the Euporia [12] web platform developed by CoPhiLab.[1] and they are provided with detailed guidelines, based on the syntax and semantics of the Domain-Specific Language (DSL) created for this task. A DSL is a formal language, usually defined by a Context-Free Grammar, which is compact and familiar to the user because it is optimized for limited purposes. Each participant annotates the same portion of the text, resulting in the first book being annotated in its entirety 15 times. As shown in Fig. 3, an annotation is constituted by a reference to the Homeric text and a sequence of one or more structured sentences, ending with a punctuation mark. If we consider the first line of the bottom box, 1.11*ētímasen* (which means: [he, i.e. Agamemnon] dishonored, disdained [Chryses]) is the reference to the text, followed by two sentences: a lexical observation and a situational observation. The structure of a sentence extends the syntax of the scripting language *turtle*: it is constituted by subject(s), predicate(s) and object(s), each of which can be preceded by hashtags to categorise them; the object(s) can be followed by a recipient (or cause) of the emotion. In our example, the first sentence is the lexical observation: `#lex this expresses #quite_neg shame, humiliation.` In this case the annotator asserts that *ētímasen* expresses two quite negative emotions: shame and humiliation. The second sentence is the situational observation: `#character Chryses feels #neg humiliation.` As pointed out by Luraghi and Sausa [6], an emotion is a complex experience that involves an experiencer (the subject), an expertum (the emotion) and a stimulus (the cause of the emotion, that can be mentioned or not). The classifiers for experiencers are: `#character` (an Iliadic hero), `#narrator` (that can be `the_poet`, or an internal narrator), `#recipient` (that can be `the_ancient_audience` or `the_contemporary_audience`), and `#annotator` (that is always expressed by "`I`"). In this way, multiple perspectives can be captured: e.g. different characters may experience different emotions in the same scene. The polarity of the emotions are graded on a scale of seven degrees from `#very_neg` to `#very_pos`.

The annotations are parsed with a compiler compiler[2] and converted in XML or json to facilitate data analysis.

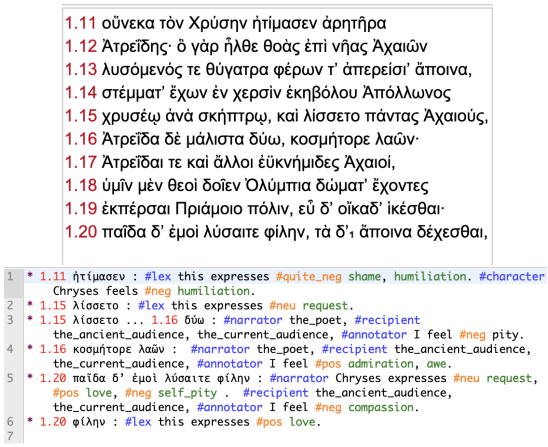

Figure 1: The user interface of Euporia

3.1. Mapping annotations to existing ontologies

x‘ Because the annotators are free to add new emotion terms to their list, we structure these terms *a posteriori* by mapping them to the following ontologies: the Emotion Ontology (*MFOEM*), the Semanticscience Integrated Ontology (*SIO*), Visualized Emotion Ontology (*VEO*), and the National Cancer Institute Thesaurus ontology (*NCIT*).
MFOEM [13] applies a cognitive perspective and builds upon MF (Mental Functioning) and BFO (Basic Format Ontology) and considers affective phenomena, e. g. emotions, moods, appraisals as well as subjective feelings.
SIO [14] uses a simpler approach but assigns positive and negative polarities to the emotions.
VEO [15] builds on MFOEM, but aims at the visual representation of emotions.
NCIT [16] applies a clinical perspective and distinguishes between emotions and feelings.

We take into account also the Time Event Ontology, *TEO*, [17] to shape temporal aspects of the annotations, such as the phases of a complex event that involves multiple emotions.
To better clarify the differences between the different ontologies, we shall now describe how *Anger* is classified in each:

- MFOEM, anger < emotion < affective process < mental process < bodily process < process < occurrent < entity;
- SIO, anger < disgust < hostility < negative emotion < emotion < behaviour < process < entity;
- VEO, anger < emotion < affective process < mental process < behaviour < action < bodily process < process < occurrent < continuant <

[1] https://cophilab.ilc.cnr.it/ A prototype of the platform, which is an app for eXist-db, is available at https://github.com/CoPhi/euporia
[2] We use the compiler compiler ANTLR: https://www.antlr.org/

entity < consequence < aspect;
- NCIT, anger < emotion < mental process < neurologic process < organismal process < biological process.

We decided to use MFOEM as the main reference ontology, because it is more suitable for our purposes.In fact, both MFOEM and our ontologies have a cognitive perspective. We then proceeded to map the list of emotions extracted from the annotations to MFOEM, SIO, VEO and NCIT.

Only 7 of these emotions - *anger, fear, hope, joy, sadness, satisfaction* - were included in all the ontologies of emotions, but other 32 terms can be mapped to one or another of them. Besides 7 unfound terms, all the remaining terms were classifiable through MultiWordNet[3], but as hyponyms of the following synsets: *feeling* (6), *speech act* (3), *cognitive state, state of mind* (3), *trait* (2), *emotion* (2), *human action* (2), *feeling, cognitive state, state of mind* (1), *sentiment* (1), *emotion, feeling* (1), *communication* (1), *state* (1), *human action, feeling* (1). Furthermore, for each emotion the following attributes have been instantiated:

- time, indicating when the emotion is perceived, with values *present, future*;
- agent, meaning who perceives the emotion, with the values *oneself, external*. There is also one instance - *submission* - of oneself toward external;
- valence, with the values, *positive, negative, ambiguous*;
- consequence, i.e. when the result of the emotion will take place and what type of result will be, with the values *unpleasant, pleasant, expected, unexpected, actual, future*. In one instance - *suspense* - it was not possible to determine the effect of the emotion.

For example, *empathy* is classified as somebody's reaction for an actual consequence to an event happened toward another agent Therefore, it is perceived in the present by oneself with a negative and it has an actual, unpleasant consequence.

By comparing the list of emotions in Pavlopoulos et al. [1] and the list extracted from the current annotations, 2 terms are missing: *guilt, loneliness* and 10 are new entries: *scorn, threat, acknowledgement, warlike, sadness, betrayal, contempt, disrespect, emotion, rage*.

4. Current results

The most frequently annotated emotion was anger (Figure 4), with 97 occurrences. The emotions of respect, aggression, and fear followed with less than 50 occurrences each. On the other hand, the most infrequent annotations

[3]https://multiwordnet.fbk.eu

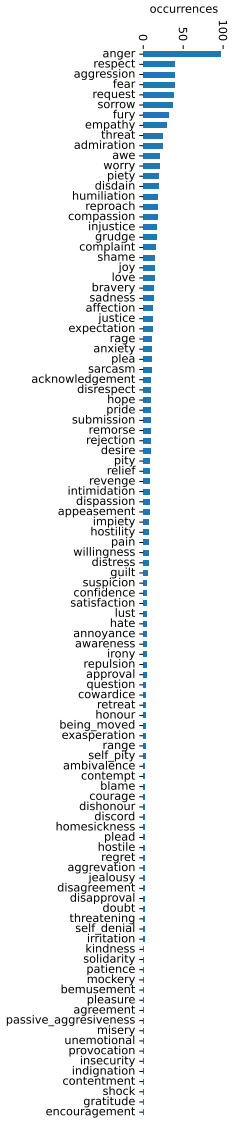

Figure 2: Frequency of annotated emotions

regarded encouragement, gratitude, and shock. In 62% of the annotations of anger, the most frequently annotated emotion, the polarity was negative (very negative in 5, quite negative in one) while in the rest it was neutral. As can be seen in Figure 4, the number of polarity-carrying emotions per verse varies. One of the highest peak was observed in verse 474, where Apollo is satisfied by the song sang to him by the Achaeans. All annotations were positive.

Currently only 4 annotators out of 15 have completed their tasks.

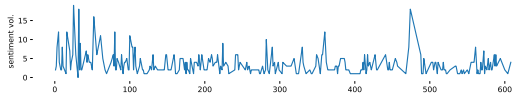

Figure 3: Number of annotations per verse, not distinguishing regarding polarity.

The DSL context-free grammar, the reference text of the first book of the *Iliad* encoded in XML-TEI, the updated annotations and the script to convert the DSL in XML are available at https://github.com/CoPhi/emohomer.

5. Conclusions

This study developed a domain-specific language for the annotation of emotive content of the first Book of Homer's *Iliad*. We mapped the list of emotions we extracted from the annotations and we discussed the results. Next steps comprise the study of polarity for more emotions, an exploration of verses with contradicting polarity, and sentiment analysis based on the subject's role (annotator, character, or narrator). Also, we plan to employ more annotators, in order to measure inter-annotator agreement, and study verses which provoke consistent and diverse emotions to the different annotators.

Acknowledgments

We are grateful to Giota Nikolopoulou and to the students and scholars that are annotating the emotional aspects related to the first book of the *Iliad*.

References

[1] J. Pavlopoulos, A. Xenos, D. Picca, Sentiment Analysis of Homeric Text: The 1st Book of Iliad, in: Proceedings of the Thirteenth Language Resources and Evaluation Conference, European Language Resources Association, Marseille, France, 2022, pp. 7071–7077. URL: https://aclanthology.org/2022.lrec-1.765.

[2] M. Pontiki, D. Galanis, J. Pavlopoulos, H. Papageorgiou, I. Androutsopoulos, S. Manandhar, SemEval-2014 task 4: Aspect based sentiment analysis, in: Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014), Association for Computational Linguistics, Dublin, Ireland, 2014, pp. 27–35. URL: https://aclanthology.org/S14-2004. doi:10.3115/v1/S14-2004.

[3] R. Sprugnoli, F. Mambrini, M. C. Passarotti, G. Moretti, Sentiment Analysis of Latin Poetry: First Experiments on the Odes of Horace, in: Proceedings of the Eighth Italian Conference on Computational Linguistics (CLiC-it 2021). Milan, Italy, January 26-28, 2022, Accademia University Press, Torino, 2021, pp. 314–320. URL: https://api.semanticscholar.org/CorpusID:245125493.

[4] R. Sprugnoli, M. Passarotti, D. Corbetta, A. Peverelli, Odi et Amo. creating, evaluating and extending sentiment lexicons for Latin., in: Proceedings of the Twelfth Language Resources and Evaluation Conference, European Language Resources Association, Marseille, France, 2020, pp. 3078–3086. URL: https://aclanthology.org/2020.lrec-1.376.

[5] V. K. Yeruva, M. ChandraShekar, Y. Lee, J. Rydberg-Cox, V. Blanton, N. A. Oyler, Interpretation of Sentiment Analysis in Aeschylus's Greek Tragedy, in: Proceedings of the The 4th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature, International Committee on Computational Linguistics, Online, 2020, pp. 138–146. URL: https://aclanthology.org/2020.latechclfl-1.17.

[6] S. Luraghi, E. Sausa, Hate and anger, love and desire: The construal of emotions in Homeric Greek, in: D. Haug (Ed.), Historical Linguistics 2013, John Benjamins, Amsterdam/Philadelphia, 2015, pp. 233–255. URL: https://api.semanticscholar.org/CorpusID:148416336.

[7] B. Sini, C. Tinti, D. Galati, Semantic Structure of Emotional Lexicon, in: A. C. Michalos (Ed.), Encyclopedia of Quality of Life and Well-Being Research, Springer Netherlands, Dordrecht, 2014, pp. 5821–5828. URL: https://doi.org/10.1007/978-94-007-0753-5_4197. doi:10.1007/978-94-007-0753-5_4197.

[8] P. R. Kleinginna, A. M. Kleinginna, A categorized list of emotion definitions, with suggestions for a consensual definition, Motivation and Emotion 5 (1981) 345–379. URL: http://link.springer.com/10.1007/BF00992553. doi:10.1007/BF00992553.

[9] L. Williams, M. Arribas-Ayllon, A. Artemiou, I. Spasić, Comparing the Utility of Different Classification Schemes for Emotive Language Analysis, Journal of Classification 36 (2019) 619–648. URL: http://link.springer.com/10.1007/s00357-019-9307-0. doi:10.1007/s00357-019-9307-0.

[10] S. Braund, G. W. Most, Ancient anger: Perspectives from Homer to Galen, volume 32, Cambridge University Press, 2004.

[11] L. C. Becker, Stoic emotion, in: S. K. Strange, J. Zupko (Eds.), Stoicism: Traditions and Transformations, Cambridge University Press, 2004,

p. 250–276. doi:10.1017/CBO9780511498374.014.

[12] G. Mugelli, F. Boschetti, A. Bellandi, R. Del Gratta, A. F. Khan, A. Taddei, Annotating ritual in ancient greek tragedy: a bottom-up approach in action, Digital Humanities Quarterly 15 (2021) 1–11.

[13] A. D. Spear, W. Ceusters, B. Smith, Functions in Basic Formal Ontology, Applied Ontology 11 (2016) 103–128. URL: https://www.medra.org/servlet/aliasResolver?alias=iospress&doi=10.3233/AO-160164. doi:10.3233/AO-160164.

[14] M. Dumontier, C. J. Baker, J. Baran, A. Callahan, L. Chepelev, J. Cruz-Toledo, N. R. Del Rio, G. Duck, L. I. Furlong, N. Keath, D. Klassen, J. P. McCusker, N. Queralt-Rosinach, M. Samwald, N. Villanueva-Rosales, M. D. Wilkinson, R. Hoehndorf, The Semanticscience Integrated Ontology (SIO) for biomedical research and knowledge discovery, Journal of Biomedical Semantics 5 (2014) 14. URL: https://doi.org/10.1186/2041-1480-5-14. doi:10.1186/2041-1480-5-14.

[15] R. Lin, M. T. Amith, C. Liang, R. Duan, Y. Chen, C. Tao, Visualized Emotion Ontology: a model for representing visual cues of emotions, BMC Medical Informatics and DecisionMaking 18 Suppl. 2 (2017) 102–113.

[16] OBO Foundry, NCI Thesaurus OBO Edition, 2022. URL: https://obofoundry.org/ontology/ncit.html, https://obofoundry.org/ontology/ncit.html Accessed = 2023-07-12.

[17] F. Li, J. Du, Y. He, H.-Y. Song, M. Madkour, G. Rao, Y. Xiang, Y. Luo, H. W. Chen, S. Liu, L. Wang, H. Liu, H. Xu, C. Tao, Time event ontology (TEO): to support semantic representation and reasoning of complex temporal relations of clinical events, Journal of the American Medical Informatics Association 27 (2020) 1046–1056. URL: https://doi.org/10.1093/jamia/ocaa058. doi:10.1093/jamia/ocaa058.