Ancient Greek's New Technological Muse: Extracting Topoi in the Anacreontea with LLMs

Rafael O. Nunes[1], **João G. Zandoná** [1], **Júlia V. Maia**[1], **Andre Spritzer**[1], **Dennis G. Balreira**[1], **Carla M.D.S. Freitas**[1]

[1]Federal University of Rio Grande do Sul, Porto Alegre, RS, Brazil

{ronunes,jgzandona,dgbalreira,carla}@inf.ufrgs.br

vieiramaiajulia@gmail.com

Abstract. *Natural Language Processing, along with Large Language Models (LLMs), holds significant potential in the domain of literature, leveraging its computational capabilities to analyze and comprehend human language. These techniques prove to be particularly useful in a specific part of Greek literature called Anacreaontea, a collection of poems emulating the style of the 6th-century BCE Greek poet Anacreon. This paper presents an LLM approach to automatically classify Anacreontea poems in their respective topoi. Our methodology explores two well-established autoregressive language models (LLama 2 and Mistral) and investigates the use of contextual prompting in this scenario. We also provide an annotated corpus with 21 fragments of the Anacreontea with topos for Greek and Portuguese text.*

1. Introduction

Natural Language Processing (NLP) can play an essential role in literature because its methods can lead to valuable insight into modalities. Methods such as Text Analysis [Silva et al. 2024, Plaisant et al. 2006], Translation [Ghazvininejad et al. 2018, Genzel et al. 2010], Named Entity Recognition [Santos 2024, Silva and Moro 2024] and Authorship Attribution [Ramezani 2021, da Rocha Bartolomei and Drummond 2020] can be especially useful to researchers in linguistics and history.

This valuable insight also extends to the Anacreontea, a collection of poems that imitate the style of the ancient Greek poet Anacreon. This is an essential part of Greek literature as these texts represent the continuation of the Anacreon tradition and its connection to the broader context of Hellenistic literature. Scholars and literary critics have categorized the different Anacreontea poems by specific *topos* (topics) and approaches. These categorizations help to understand the main themes and motifs in the Anacreontea and the different ways poets approached Anacreon's style.

In this study, we classify anacreontic poems in their respective *topoi* using autoregressive language models with a context configuration according to the categorization assigned by Antunes [Antunes 2016]. Our goal is to address the scarcity of studies on ancient Greek literature, especially regarding the Anacreontea corpus. To the best of our knowledge, NLP research has not yet addressed this topic.

The contributions of this work are as follows:

- A study on the extraction of Anacreontea topoi using contextual prompts to Large Language Models.[1]
- A corpus with 21 fragments of the Anacreontea annotated with topoi based on the original Greek and its Portuguese translation[2].

2. Related Work

Most of the research on the application of NLP to the domain of ancient Greek literature focuses on sentiment analysis [Yeruva et al. 2020a, Yeruva et al. 2020b, Pavlopoulos et al. 2022]. This is possibly due to the potential of how computational exploration of cultural and linguistic context of ancient literature can help with the challenge of bridging the ancient Greek language with modern understandings [Pavlopoulos et al. 2022]. Works in this field include investigating the human-machine collaboration in sentiment analysis for classical Greek tragedy [Yeruva et al. 2020a], proposing the use of sentiment estimators as mechanical annotators for distant reading of Homeric texts [Pavlopoulos et al. 2022], and evaluating the consistency of annotation in the application of human-in-the-loop methodology in literary text [Yeruva et al. 2020b].

Regarding LLMs for the modern Greek language, one of the most prominent is GREEK-BERT [Koutsikakis et al. 2020], which outperforms two multilingual Transformer-based models in two of the proposed benchmarks. The pre-training corpora include texts from Wikipedia, the European Parliament Proceedings Parallel Corpus, and OSCAR, a cleansed version of Common Crawl. Researchers have also introduced a pilot study focused on the automatic linguistic preprocessing of ancient and Byzantine Greek [Singh et al. 2021] along with a distilled version of the GREEK-BERT model, called DistilGREEK-BERT [Karavangeli et al. 2023].

According to the recent work analyzed, ancient Greek literature still presents an interesting challenge, which we will address using NLP techniques.

3. Anacreontea Corpus

The Anacreontea corpus is a set of texts attributed to the ancient Greek poet Anacreon that were actually written by different authors over many centuries [Sens 2014]. Even though the Anacreontea are not the original works of Anacreon himself, they hold significant importance in the realm of poetry and contribute to our understanding of Anacreon's image in antiquity [Antunes 2016, Antunes 2013]. Furthermore, they serve as sources of inspiration for other poets [Antunes 2016].

We obtained each topos[3] and approaches from Antunes [Antunes 2016], along with a selection of poems mapped to each approach and topos. Table 1 shows the names, the number of annotated fragments, and statistics about the number of tokens to each topos. We only had two approaches with one fragment, which are personification (*personificação*) and tribute (*homenagem*). Additionally, we used this source's Portuguese translations. However, we could not find a work containing all poems and their analyses. To address this issue, we referred to Antunes's other work, encompassing all translations of Anacreontea [Antunes 2013], to retrieve specific texts that were missing. We also find

[1]The code is available at https://github.com/RafaelOleques/anacreontea

[2]The corpus is available at https://huggingface.co/datasets/ronunes/anacreontea

[3]The Anacreontea thematic.

Table 1. Name of each topos and statistics. The following coding is used for the topoi: (I) Bittersweet eros (*Eros doceamargo*), (II) Madness, sobriety, and the many voices of the Anacreontea (*Loucura, sobriedade e as muitas vozes das Anacreônticas*), (III) Poíesis, (IV) Wealth (*Riqueza*) and (V) Old age (*Velhice*).

		Portuguese tokens			Greek tokens		
Code	**Fragments**	**Mean**	**Min**	**Max**	**Mean**	**Min**	**Max**
I	6	51.67 ± 22.81	27	82	40.67 ± 16.82	19	60
II	4	73.00 ± 19.44	49	95	47.50 ± 11.41	35	64
III	6	119.50 ± 50.54	57	209	93.17 ± 47.11	40	180
IV	2	70.00 ± 8.00	62	78	60.00 ± 2.00	58	62
V	3	101.17 ± 46.29	59	180	71.50 ± 34.71	40	123
Total	21	91.00 ± 45.62	27	209	67.62 ± 37.70	19	180

three more poems not mentioned in Antunes's latter work [Antunes 2016] but classified in the former [Antunes 2013].

We represent the corpus with the attributes *fragment*, *topos*, *topos_english*, *portuguese_text*, *greek_text*, and *source*, with the source containing the DOI of the paper with the Portuguese and Greek texts of the poem.

4. Extracting Anacreontea Topoi

This section outlines the proposed method for extracting the topoi from the Anacreontea. We introduce the LLM models used in this work, briefly describing their architecture and data used for instruction tuning. We then describe the prompt used with these components and our design choices.

4.1. LLM Models

Mistral-7B-Instruct-v0.2 [Jiang et al. 2023] is an enhanced fine-tuned version of Mistral-7B-Instruct-v0.1, both released under the open-source Apache 2.0 license. Mistral uses a 7-billion parameter model with a transformer-based architecture, leveraging Sliding Window Attention (SWA) and stacking 32 transformer blocks. The SWA mechanism uses a window size of 4096, allowing the model to work with a context length of 8096 tokens. Fine-tuning of the model was performed using publicly available instruction datasets from the Hugging Face repository.

LLama2 7B Chat [Touvron et al. 2023] is an open-source auto-regressive language model released by Meta in July of 2023. The model uses a transformer architecture and was pre-trained on publicly available data sources adding up to 2 trillion tokens, considering a context size of 4096 tokens. The tuned version uses supervised fine-tuning (SFT) with over 100,000 examples and reinforcement learning with human feedback (RLHF) with over 1,000,000 human preferences.

4.2. Prompt

We designed our prompt by incorporating four components: persona, instruction, context, text, and question. We aimed to provide the LLM with a more comprehensive understanding of the problem and domain knowledge, enhancing its performance. We provide a detailed description of the prompt components below.

The structure of the prompt follows the format {persona-instruction-context-text-question}[4]. We defined a domain expert **persona** to help the model follow the classification [Salewski et al. 2023, White et al. 2023]. Next, we provided an **instruction** to directly guide the model to the task [White et al. 2023]. The instruction pointed out that the model needed (i) to choose the topos that better represent the text subject and (ii) to use topos in the same way given in the context. Part (i) of the instruction aimed to make the model return only one classification whereas part (ii) was meant to ensure that it uses the correct classes.

The **context** could be provided in two ways: label names or definitions. Our aim with this was to analyze how knowledge of the domain could help an LLM understand a task. We based our question on the specific topos domain, which is uncommon in general texts (e.g., Bittersweet eros). Moreover, topoi that carry many meanings (e.g., Madness, sobriety, and the many voices of the Anacreontea) can be difficult for the models to understand without the context and thus provide the right classification [White et al. 2023]. All the definitions were extracted from Antunes's study [Antunes 2016]. To avoid data loss, we removed direct references to specific texts in the definitions.

The **text** begins with "*Anacreôntica:*", followed by the text of the poem in triple quotes. We used triple quotes to provide the model with a better context for the poem's lines.

In the final part of the prompt, a **question** asks the model to inform the most significant topos of the poem. This directly follows the poem's text because a previous study showed that, within the prompt, information near the question can positively interfere with the result [Gupta et al. 2023]. The question is accompanied by a question-answer template. The answer begins with "*Vamos pensar passo a passo*" (*Let's think this through step by step.*) in order to instigate the model to use Chain of Thoughts (CoT) [Wei et al. 2022].

5. Experimental Evaluation

This section describes our experimental evaluation, beginning with our hardware setup and our choice of programming language and libraries. We then detail the hyperparameters used in the models and the metrics employed to evaluate our results.

5.1. Setup

We ran the LLM models on a computer with an Nvidia GeForce RTX 4070 GPU and 12 GB of RAM. We chose Python 3.7.6 as our programming language because of its extensive support for Machine Learning and Natural Language Processing libraries. We used the *bitsandbytes* library[5] for quantization.

5.2. Hyperparameters

To load Mistral and LLama2 in our system, we used the AutoModelForCausalLM class[6] from HuggingFace, setting *load_in_4bit = True* to use 4-bit quantization,

[4]An example of a complete prompt can be found in the readme on our GitHub repository: https://github.com/RafaelOleques/anacreontea/

[5]https://github.com/TimDettmers/bitsandbytes

[6]https://huggingface.co/docs/transformers/model_doc/auto

Table 2. Results of Mistral 7B using definitions as contextual information.

Category	Precision	Recall	F1-Score	Support
Bittersweet eros	1.00	0.33	0.50	6
Madness, sobriety and the many voices of the Anacreontea	0.33	0.50	0.40	4
Poíesis	0.83	0.83	0.83	6
Wealth	0.00	0.00	0.00	2
Old age	0.75	1.00	0.86	3
Accuracy			0.57	21
Macro avg	0.49	0.44	0.43	21
Weighted avg	0.69	0.57	0.58	21

bnb_4bit_compute_dtype = torch.bfloat16 to accelerate the computation and *torch_dtype = torch.float16* to save memory. To generate the results, we used the standard values of the *generate* method of *AutoModelForCausalLM*.

5.3. Metrics

We used scikit-learn [Pedregosa et al. 2011] to compute the metrics. The metrics used were **precision**, **recall**, **F1-Score**, and **accuracy**. With the exception of accuracy, all metrics featured macro and weighted averages, and results were obtained for each topos. In Section 6, we thoroughly discuss precision and recall in order to explain the results and the impact of the definitions. Our focus is nonetheless on the F1-Score of each class because our aim is to balance precision and recall. Furthermore, our study focuses on understanding the use of definitions in the final results for each topos.

6. Results and Discussion

Before we delve into a detailed analysis of the topoi, we provide a bird's eye view of the key results.

As shown in Tables 2 and 3, for the Mistral model, the contextual reference could enhance the Macro F1-Score by 25% and weighted F1-Score by 20%, with other metrics following the same pattern. Considerable differences arise when comparing models with and without definition information. Those incorporating definitions introduce one out-of-context class (*Joy and Drink*), while non-contextual prompts lead to the addition of three classes (*Immortality*, *Wealth and Love Bittersweet eros*, and *Wealth and Bittersweet eros*).

This trend persists with LLama2, as shown in Table 4 for models with definitions and Table 5 for those without. The former reveals new classes like *Madness, sobriety, and the many voices of the Anacreontea and Poíesis*, while the latter introduces *Irony*, *None*, *Sovereignty*, and *Old age and Madness*. However, Mistral had better results in all metrics and classes, so our analysis focuses on its results.

6.1. Analysis of specific topos

Insight can be derived from analyzing a specific topos. This examination underscores the significance of utilizing definitions to comprehend domain classes.

Table 3. Results of Mistral 7B using only label names.

Category	Precision	Recall	F1-Score	Support
Bittersweet eros	0.33	0.50	0.40	6
Madness, sobriety and the many voices of the Anacreontea	0.25	0.25	0.25	4
Poíesis	1.00	0.17	0.29	6
Wealth	0.00	0.00	0.00	2
Old age	1.00	0.33	0.50	3
Accuracy			0.29	21
Macro avg	0.32	0.16	0.18	21
Weighted avg	0.57	0.29	0.31	21

Table 4. Results of LLama2 7B using definitions as contextual information.

Category	Precision	Recall	F1-Score	Support
Eros sweet-bitte	0.50	0.17	0.25	6
Madness, sobriety and the many voices of the Anacreontea	0.33	0.75	0.46	4
Poíesis	0.67	0.67	0.67	6
Wealth	0.00	0.00	0.00	2
Old age	0.33	0.33	0.33	3
Accuracy			0.43	21
Macro avg	0.31	0.32	0.29	21
Weighted avg	0.44	0.43	0.40	21

Table 5. Results of LLama2 7B using only label names.

Category	Precision	Recall	F1-Score	Support
Bittersweet eros	0.00	0.00	0.00	6
Madness, sobriety and the many voices of the Anacreontea	0.29	0.50	0.36	4
Poíesis	0.33	0.17	0.22	6
Wealth	0.00	0.00	0.00	2
Old age	0.00	0.00	0.00	3
Accuracy			0.14	21
Macro avg	0.07	0.07	0.07	21
Weighted avg	0.15	0.14	0.13	21

Bittersweet eros shows a notable improvement in precision, rising from 33% to 100%, even if expressing a degree of recall of 50% to 33%. One contributing factor to precision improvement was the specificity of the class name, which may have been misconstrued as a generic love poem. For instance, Fragment 1 delves into the lyrical self's affection for Anacreon despite Anacreon's advanced age. A poem in the primary topos encompasses the theme of old age (Old age) but also intertwines themes of love and references to Eros. When the model lacks precise connections to the nuanced meaning of topos Bittersweet eros, it runs the risk of misidentifying it as a love poem that articulates a certain level of ambiguity or confusion, exemplified in this fragment where the author expresses, "Eu, tolo, então a aceitei" [7].

Regarding the model that does not use definitions, the most probable reason for the lower precision is that Anacreontea has love associated with many topoi. Love as a recurring theme in the corpus is illustrated by the many fragments that mention gods that have love in their domains, such as Eros and Aphrodite, or talk about the madness of love, such as that driven by wine (e.g., Dionysus) [Antunes 2015]. Since poems cited below in other topoi classified as Bittersweet eros correpond to this subject, this seems to be a plausible explanation.

Madness, sobriety, and the many voices of the Anacreontea had an improvement in both precision and recall. Unlike other topoi that focus on specific subjects, it exhibits a rich tapestry of themes, including madness, the interplay between wine and intense emotions, sobriety, and diverse voices. This diversity can be challenging to identify themes with precision. One example is Fragment 12, which portrays mythological scenes of madness. Madness is repeatedly exhibited not only by the characters, but also by the speaker himself, who expresses a desire to go mad. Using this definition, the fragment could be identified in the right topos, but in contrast, using only label names, it was identified as Bittersweet eros, probably due to the expression of love and madness by Attis to Cybele.

Poíesis obtained a large improvement by using definitions, with the F1-Score increasing from 29% to 83%. In the case of using only the label name, the model could only identify Fragment 17, which contains an explicit description of the sculpting of a statue with several mythical references and discussions about art. In contrast, poems such as Fragment 60(a), which express the composition of music, are harder to classify because they talk about the subject of the piece of music/art. This poem, in particular, concerns a musical composition on the unrequited love of Apollo, which in principle would express Bittersweet eros as classified by the model that does not use definitions. However, this is not actually author's focus; his concern is actually music itself and its composition.

Wealth was the most difficult class to classify, achieving 0% of F1-Score. This was probably because of the scarcity of examples (two) and the possible intersection with other topoi. For instance, Fragment 36 discusses death and the importance of enjoying life, which could fit into the topos Old age — as was the model's prediction — however, the poem does not refer to the advanced age of the speaker, but to contempt for money and the enjoyment of life.

Old age also obtained an improvement in the F1-Score through the use of defi-

[7] "I, fool, then accepted it" (free translation to English)

nitions, increasing from 50% to 86%. In Table 3, we can see that it obtained 100% of precision, but this occurred because only one fragment was classified as Old age, and it was corrected. Instead, Table 2 shows a more balanced distribution between precision and recall. As discussed in the other classes, this topos also suffers from having examples, such as Fragment 39, that are classified as Bittersweet eros in the prompt that does not use definitions.

Out-of-domain topoi were detected in predictions. Even if there are possible reasonable *topoi*, they are not present in Antunes's original classification [Antunes 2016], e.g., *Immortality* for Fragment 36 and *Joy and Drink* for Fragment 42. The out-of-domain topoi could be manually assigned to specific topoi, which could be compound topoi, such as *Wealth and Bittersweet eros*. Still, we prefer not to do this to avoid a positive result bias due to modifications by the study's authors.

6.2. Qualitative Analysis

We asked an undergraduate student specializing in linguistics and literature to analyze instances where the model-generated topoi diverged from the originals. The detailed analysis is available in Appendix A. Our goal with this analysis was to explore the model's capability to extract novel topoi from fragments, as the original work [Antunes 2016] was not aimed at making exhaustive analyses.

We observed that the model identified topoi that were not originally annotated in six of the nine fragments with distinct classifications from the golden corpus. This underscores the ability of LLMs to comprehend subtle variations and highlights the efficacy of contextual prompts in facilitating a nuanced understanding of different topoi.

7. Conclusion

In this study, we proposed an LLM-based approach to automatically extract topoi from the Anacreontea corpus. We performed experiments using the LLama2 and Mistral models using context and no-context prompts to assess similarity of results with and without additional domain-specific knowledge. We found that Mistral achieved better performance and that specific knowledge of the definition of the topoi significantly increased all metrics considered.

Lack of data was a limitation of this study. We obtained classes with very few examples because we performed a manual extraction of topoi from a work that was not meant for conducting exhaustive analyses of all individual fragments, but rather of the topoi and the approach itself. As such, it is difficult to analyze all classes with precision. This is especially true of the class level of "approaches", in which we found only one example of each.

In future studies, we plan to extend this work by searching for other classifications of the fragments and increasing the number of fragments in our corpus. We also intend to test other techniques to improve the context, such as In-context Learning and Retrieval-Augmented Generation.

References

Antunes, C. L. B. (2013). As anacreônticas e a imagem de anacreonte na antiguidade. *Letras Clássicas*, 17(1):109–149.

Antunes, C. L. B. (2015). Byron, tradutor anacreôntico: Anacreôntica 23 e "to his lyre". *Translatio*, (10):90–96.

Antunes, C. L. B. (2016). Entre amores ébrios e fazer artístico: Tópoi e poíesis nas'anacreônticas'— between drunken love and artistic praxis: Tópoi and poíesis in the'anacreontea'. *Estudos Linguísticos e Literários*, (55):374–385.

da Rocha Bartolomei, B. and Drummond, I. N. (2020). Authorship attribution of brazilian literary texts through machine learning techniques. In *Brazilian Conference on Intelligent Systems*, pages 389–402. Springer.

Genzel, D., Uszkoreit, J., and Och, F. (2010). "poetic" statistical machine translation: Rhyme and meter. In Li, H. and Màrquez, L., editors, *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 158–166, Cambridge, MA. Association for Computational Linguistics.

Ghazvininejad, M., Choi, Y., and Knight, K. (2018). Neural poetry translation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 67–71.

Gupta, S., Gardner, M., and Singh, S. (2023). Coverage-based example selection for in-context learning. In Bouamor, H., Pino, J., and Bali, K., editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13924–13950, Singapore. Association for Computational Linguistics.

Jiang, A. Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D. S., de las Casas, D., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., Lavaud, L. R., Lachaux, M.-A., Stock, P., Scao, T. L., Lavril, T., Wang, T., Lacroix, T., and Sayed, W. E. (2023). Mistral 7b. *arXiv preprint arXiv:2310.06825*.

Karavangeli, E. A., Pantazi, D.-A., and Iliakis, M. (2023). Distilgreek-bert: A distilled version of the greek-bert model.

Koutsikakis, J., Chalkidis, I., Malakasiotis, P., and Androutsopoulos, I. (2020). Greek-bert: The greeks visiting sesame street. In *11th Hellenic conference on artificial intelligence*, pages 110–117.

Pavlopoulos, J., Xenos, A., and Picca, D. (2022). Sentiment analysis of homeric text: The 1st book of iliad. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 7071–7077.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Plaisant, C., Rose, J., Yu, B., Auvil, L., Kirschenbaum, M. G., Smith, M. N., Clement, T., and Lord, G. (2006). Exploring erotics in emily dickinson's correspondence with text mining and visual interfaces. In *Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries*, pages 141–150.

Ramezani, R. (2021). A language-independent authorship attribution approach for author identification of text documents. *Expert Systems with Applications*, 180:115139.

Salewski, L., Alaniz, S., Rio-Torto, I., Schulz, E., and Akata, Z. (2023). In-context impersonation reveals large language models' strengths and biases. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Santos, D. (2024). Literary similarity of novels in Portuguese. In Gamallo, P., Claro, D., Teixeira, A., Real, L., Garcia, M., Oliveira, H. G., and Amaro, R., editors, *Proceedings of the 16th International Conference on Computational Processing of Portuguese*, pages 539–543, Santiago de Compostela, Galicia/Spain. Association for Computational Lingustics.

Sens, A. (2014). Dialect in the anacreontea. *Imitate Anacreon*, pages 97–112.

Silva, M. O., de Melo-Gomes, L., and Moro, M. M. (2024). From words to gender: Quantitative analysis of body part descriptions within literature in portuguese. *Information Processing & Management*, 61(3):103647.

Silva, M. O. and Moro, M. M. (2024). Evaluating pre-training strategies for literary named entity recognition in Portuguese. In Gamallo, P., Claro, D., Teixeira, A., Real, L., Garcia, M., Oliveira, H. G., and Amaro, R., editors, *Proceedings of the 16th International Conference on Computational Processing of Portuguese*, pages 384–393, Santiago de Compostela, Galicia/Spain. Association for Computational Lingustics.

Singh, P., Rutten, G., and Lefever, E. (2021). A pilot study for bert language modelling and morphological analysis for ancient and medieval greek. In *5th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature, co-located with EMNLP 2021*, pages 128–137. Association for Computational Linguistics.

Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C. C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P. S., Lachaux, M.-A., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E. M., Subramanian, R., Tan, X. E., Tang, B., Taylor, R., Williams, A., Kuan, J. X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., and Scialom, T. (2023). Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q. V., Zhou, D., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.

White, J., Fu, Q., Hays, S., Sandborn, M., Olea, C., Gilbert, H., Elnashar, A., Spencer-Smith, J., and Schmidt, D. C. (2023). A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382*.

Yeruva, V. K., Chandrashekar, M., Lee, Y., Rydberg-Cox, J., Blanton, V., and Oyler, N. A. (2020a). 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*, pages 138–146.

Yeruva, V. K., Chandrashekar, M., Lee, Y., Rydberg-Cox, J., Blanton, V., and Oyler, N. A. (2020b). Interpretation of sentiment analysis with human-in-the-loop. In *2020 IEEE International Conference on Big Data (Big Data)*, pages 3099–3108. IEEE.

A. Manual Analysis

This appendix presents a manual analysis conducted by an undergraduate student specializing in linguistics and literature. The student has already completed the courses on classical studies, specifically focusing on ancient Greek and ancient Greek literature, thereby contributing with expertise to the qualitative analyses.

Below, we present the fragment number, the original topos in Antunes [Antunes 2016], the inferred topos, and the analysis made. It is important to highlight that the goal of the analysis was to understand if some of the predictions could contribute to the finding of extra topoi in the text, since the original work did not focus on analyzing all the topoi of each fragment.

Fragment: 2
Original topos: Madness, sobriety, and the many voices of the Anacreontea
Predicted topos: Poíesis
Analysis: Yes. The emphasis on "doing poetry" is already described by the motto: "Give me Homer's lyre // Without the murderous string."[8] while the construction of the nomos in "Bring me the laws mixed in them"[9] is a predicate of this act of poetry-making. Furthermore, it is worth reiterating that the motto traditionally describes the total meaning of the poem, so the option to "poetry-making" is quite clear. It seems that the original topic occurs in relation to other works, the predicted one analyzes the poem in isolation.

Fragment: 8
Original topos: Wealth
Predicted topos: Madness, sobriety, and the many voices of the Anacreontea
Analysis: Yes. The choice of topos as "Loucura, sobriedade e as muitas vozes das Anacreôntica" highlights the poet's tone and conclusions regarding life, which is a hallmark of his corpus. See the third comment.

Fragment: 60b
Original topos: Bittersweet eros
Predicted topos: Madness, sobriety, and the many voices of the Anacreontea
Analysis: Yes. Once again, the poet's perception and his moral code seem to affirm the topos predicted by the "Anacreontic voices." By inserting himself into the poem, the poet seems to favor his own view rather than a purely erotic speech: "themes that generate division and pain are not welcome in Anacreonte's poetry"[10] (RAGUSA, 2021, p. 212). As attested by fragment 60b in its rejection of martial aspects (e.g., "avoiding the star // Whose light shines scarlet.") in favor of the Venusians.

[8] *Dá-me a lira de Homero// Sem a corda do assassínio*
[9] *Traz-me as leis mescladas nelas*
[10] *na poesia de Anacreonte não são bem vindos temas que gerem cizânia e dor*

Fragment: 4
Original topos: Poíesis
Predicted topos: Wealth
Analysis: Yes. Apparently, the program classified fragment 4 as "Wealth" due to the description of "Hephaestus' Cup"[11], which, being interpreted for its usefulness, in this case, has its ornaments rejected as being useless for drinking. As such, the term's classification seems rather to first recover the tone of carpe diem and the rejection of luxuries, and this only in a secondary sense, in the act of poetry-making.

Fragment: 35
Original topos: Bittersweet eros
Predicted topos: Madness, sobriety, and the many voices of the Anacreontea
Analysis: No, it's wrong.

Fragment: 59
Original topos: Bittersweet eros
Predicted topos: Wealth
Analysis: No. Although you might say, "No, definitely."[12] It is possible to make a distant inference from Bacchus and his relationship to the city, since this was a god marked by certain traits that are opposed to the polis or its values, as Jeanmarie attests in his work "Dionysos," or even Euripides, in the "The Bacchae."

Fragment: 36
Original topos: Wealth
Predicted topos: Old age
Analysis: No, it's wrong.

Fragment: 42
Original topos: Madness, sobriety, and the many voices of the Anacreontea
Predicted topos: Alegria e Beber
Analysis: Yes. The poem is set in the symposium scene, its description is quite precise and concerns the general content of the poem, Anacreon's common exhortation. However, this subject appears indirectly and circumstantially. In this case, the choice of topic seems to centralize the common theme in the poet's body of work, and not so much in the specific facets presented.

Fragment: 29
Original topos: Bittersweet eros
Predicted topos: Madness, sobriety, and the many voices of the Anacreontea
Analysis: Yes. Unlike the previous poem, this one has its morality highlighted, with its judgment being a piece of anacreontic perception.

---

[11] *Taça de Hefesto*
[12] *Não, definitivamente.*