HoLM: Analyzing Linguistic Unexpectedness in Homeric Poetry

John Pavlopoulos[1,2], Ryan Sandell[3], Maria Konstantinidou[4], Chiara Bozzone[3]

[1] Department of Informatics, Athens University of Economics and Business, Greece
[2] Archimedes/Athena RC
[3] Ludwig-Maximilians-Universität München, Germany
[4] Department of Greek Philology, Democritus University of Thrace, Greece

annis@aueb.gr, {ryan.sandell,chiara.bozzone}@lrz.uni-muenchen.de, mkonst@helit.duth.gr

Abstract

The authorship of the Homeric poems has been a matter of debate for centuries. Computational approaches such as language modeling exist that can aid experts in making crucial headway. We observe, however, that such work has, thus far, only been carried out at the level of lengthier excerpts, but not individual verses, the level at which most suspected interpolations occur. We address this weakness by presenting a corpus of Homeric verses, each complemented with a score quantifying linguistic unexpectedness based on Perplexity. We assess the nature of these scores by exploring their correlation with named entities, the frequency of character *n*-grams, and (inverse) word frequency, revealing robust correlations with the latter two. This apparent bias can be partly overcome by simply dividing scores for unexpectedness by the maximum term frequency per verse.

Keywords: Homeric Poetry, Ancient Greek, Language Modeling

1. Introduction

The authorship of the *Iliad* and the *Odyssey* has been a matter of debate since at least the third century BCE (Nesselrath, 2011): during the Hellenistic period, scholars active at Alexandria (e.g., Zenodotus, Aristophanes of Byzantium) worked to establish a text for the *Iliad* and *Odyssey* on the basis of variant texts already in circulation (which are fragmentarily known today through surviving papyri). The technique of composition of the poems is oral and traditional (Parry, 1971; Lord, 1960); see further the work of Foley (2007) for the impact of oral-formulaic theory on the interpretation of the Homeric epics, and the work of Bozzone (2014, 2024) on their natural-language character and deviations therefrom. However, we do not know precisely how or when the poems came to be written down, and to what an extent they were altered during their written transmission (Turner, 2011). One popular hypothesis is that the poems were first put into writing through dictation by one or more singers in the context of a Panathenaean festival during the 6th century BCE (Jensen, 2011; Ready, 2019), though some (West, 2011, 2013) specifically imagine (one or more) literate poets who were responsible for first committing the songs to writing. Scholars nowadays agree that the poems may have absorbed contributions from different performers, scribes, and editors over several centuries, before being largely stabilized by the work of scholars in Alexandria (Haslam, 2011). West (1998, v) considers it "evident" that the poems contain later interpolations, and most specialists of Homeric philology regard Book 10 of the *Iliad* (the *Doloneia*) as a large-scale interpolation (Danek, 1988, 2012). Nonetheless, there is considerable disagreement as to whether the entire Homeric corpus is the product of a single poet (Wachter, 2007), whether each poem overwhelmingly reflects the work of a single author (West, 2011, 2013), which only underwent relatively minor alterations during its transmission, or whether both poems should be seen as the product of a centuries-long process of textualization, in which individual contributions are either difficult or impossible to identify (Nagy, 2004, 2009). See also Bozzone and Sandell (2022, 21–24) for a similar brief introduction to these issues of textualization of the Homeric epics.

Computational approaches can provide experts with tools for reassessing and making crucial headway on these issues (Pavlopoulos and Konstantinidou, 2023; Bozzone and Sandell, 2022). Fasoi et al. (2021), for example, identified linguistically non-conforming (unexpected) passages by employing character-level statistical language modeling, which might point to potential interpolations, revisions, or excerpts of different provenance in the *Iliad* and *Odyssey*. The structural units used in their models were at the book-level, but excerpts of up to several hundred lines were proposed by Pavlopoulos and Konstantinidou (2023). However, although lengthy interpolations (e.g., the entirety of Book 10 of the *Iliad*) may exist (Bozzone and Sandell, 2022, 22), most of the suspected interpolations in the Homeric epics consist of just single lines. This situation poses a special challenge for quantitative analyses, especially if we consider that the language of Homer is, to a high degree, formulaic. i.e., composed of prefabricated sequences of words (*formulas*), such that many phrases, verses,

LREC-COLING 2024, pages 8166–8172
20-25 May, 2024.

	Indices of (near) identical verses
I1	(38, 452), (23, 377) (43, 457)
I2	(534, 545, 556, 630, 644, 710, 737, 759)
O1	(213, 230, 306, 345, 388, 412), (44, 80, 314)
O2	(129, 208, 309, 371), (25, 161, 229)

Table 1: Sample identical or near-identical verses in Books 1 and 2 of the *Iliad* (I) and *Odyssey* (O). In O1, in the first tuple, one letter of one article is changed across the verse occurrences. The first tuple of O1 and O2 refers to the same verse.

	Verse	**Line numbers**
I1	killan te zatheēn tenedoio te iphi anasseis	(38, 452)
I1	aideisthai th' hierēa kai aglaa dechthai apoina	(23, 377)
I1	hōs ephat' euchomenos, tou d' eklye Phoibos Apollōn	(43, 457)
I2	tō d' hama tessarakonta melainai nēes heponto	(534, 545, 556)
I2	*idem*	(630, 644, 710)
I2	*idem*	(737, 759)
O1	tēn d' au Tēlemachos pepnumenos antion ēuda	(213, 230, 306, 345)
O1	ton d' au Tēlemachos pepnumenos antion ēuda	(388, 412)
O1	ton d' ēmeibet' epeita thea glaukōpis Athēnē	(44, 80, 314)
O2	ton d' au Tēlemachos pepnumenos antion ēuda	(129, 208, 309)
O2	keklyte dē vyn meu, Ithakēsioi, hotti ken eipō	(25, 161, 229)

Table 2: Examples of identical verses in Books 1 and 2 of the *Iliad* and *Odyssey*.

and even entire blocks of verses are repeated in identical or near-identical form in the text (Parry, 1971; Bozzone, 2024). A few cases of such identical or near-identical verses are listed in Table 1, to provide some examples, and identical verses are also cited in Table 2.[1]

Our work employs the statistical testing seen in Pavlopoulos and Konstantinidou (2023), but instead of applying it to entire books or lengthy excerpts, we apply it to single verses of the Homeric poems. The outcome of this application is the identification of linguistically surprising verses, which constitutes a novel result in the literature. Using a variant of the measure of PERPLEXITY, we estimated the linguistic unexpectedness of each verse, and we investigated its correlation with the presence of named entities, the poem-level frequency of character 5-grams, and (inverse) word frequency. Our empirical analysis shows that a strong correlation exists between PERPLEXITY and the latter two, revealing a bias towards high PERPLEXITY where *hapax legomena* occur. We argue, however, that this bias can potentially be bypassed by combining the two into a single measure. The resource presented here, named HoLM, along with all scores computed, is publicly released at https://github.com/ipavlopoulos/holm.

2. Background and Related Work

Methods of **authorship attribution** (Love, 2002; Juola, 2006; Stamatatos, 2009; Kestemont et al., 2012; Manousakis and Stamatatos, 2018; Kabala, 2020) are related to our study, but our goals in this paper are orthogonal to questions of authorship. In principle, one could adopt some predefined set of (Ancient Greek) authors to which to attempt to attribute the authorship of a given Homeric text. Alternatively, one could compare the text of a known author with that of a given Homeric text, in order to verify whether the author of the two is the same (known-author verification). Direct attempts at attribution and verification of the Homeric epics in such a fashion confront, however, a serious obstacle: the near-total absence of any contemporary literature with which the Homeric epics can be directly compared. At best, one can employ unsupervised machine learning methods (e.g., clustering) to attempt to verify whether two documents (e.g., books of Homer) might have the same author, or whether all material traditionally attributed to Homer can consistently be distinguished from Ancient Greek texts of a similar genre (i.e., hexametric poetry). See the discussion of and some initial, tentative answers offered to these questions in Bozzone and Sandell (2022). The approach adopted in our study does not rely on any resources external to the Homeric poems.

3. Estimation of Unexpectedness

Given the verses V^b from book b of a poem (*Iliad* or *Odyssey*), we train a statistical character-level language model m_b on all the verses of the remaining books. We compute the negative log-likelihood of m_b for each character of $v \in V^b$, and then we average this score across the verse's characters, which we call as the average bits per character (BPC) score following the notation of Hwang and Sung (2017). Then, as suggested by Graves (2013) and in accord with the work of Pavlopoulos and Konstantinidou (2023), we define the equivalent Perplexity variant (Dror et al., 2020) for v as:

$$PPL(v, m^b) = 2^{|\overline{w}|*BPC(v,m^b)} \quad (1)$$

Iterating b over the books of each poem yields forty-eight language models (LMs) in total, twenty-four per poem, and we obtain one PPL score per verse of Homer. The verses, along with their PPL scores, constitute the core of the actionable dataset in HoLM that was developed. In the following section, this dataset is discussed in detail.[2]

[1] Transliteration of Greek follows the Library of Congress/ ALA-LC Romanization Conventions. All experiments and computational analyses were performed, of course, on the original Greek text. Note that the apostrophe shown in these transliterations represents a special type of apostrophe in the Greek, which always represents the elision of a vowel, and which was not removed during preprocessing of the text (see §4).

[2] We would like to point out here that the same ba-

4. The HoLM Dataset

Data Preparation and Preprocessing The HoLM dataset comprises one PPL score per Homeric verse. Fig. 1 depicts these scores as a time series, beginning from the first verse of each poem (left) and continuing until the last (right). Underlying this dataset are digitizations of modern textual editions of the *Iliad* and *Odyssey*, specifically, Monro and Allen (1920) for the *Iliad* and Murray (1919) for the *Odyssey*; these are the same textual editions that have been employed in other recent computational studies of the Homeric epics (Pavlopoulos and Konstantinidou, 2023).[3] The texts were preprocessed by dividing them into individual hexametric lines (12107 for the *Odyssey*, 15683 for the *Iliad*) as indicated by line-breaks in the textual editions, then removing all non-alphabetic characters (punctuation, numbers, etc.) from the text and converting all characters to lowercase.[4]

Exploratory analysis A given verse in each poem consists of approximately 42$\pm$1 characters (mean and st. deviation) and 7$\pm$1 words. A word is, on average, made up of 5.04 (*Iliad*) or 5.00 (*Odyssey*) characters. Books of the *Iliad* typically contain a larger number of verses. While our LMs were trained only on data from either the *Iliad* or *Odyssey*, one based on the totality of the Homeric corpus would exhibit bias towards the *Iliad*.

Recurring verses As is shown in Table 3, specific identical verses can and do appear multiple times across each poem. In the *Odyssey*, the occurrences of the five verses shown in the Table make up approximately 1% (109/12107) of the verses in the entire poem. The same number is considerably lower for the *Iliad*: the verses shown in Table 3 comprise only 0.33% (51/15683) of the poem. The application of agglomerative clustering (Ward's method) to each book further reveals that many near duplicates also exist. That is, slightly altered variations of the same verse may appear; for example, verses I.1.84 and I.1.215 differ only in the gender of a pronoun, resulting in a difference of just a single character (/tēn/ $\sim$ /ton/).

	Verse	#
I	kai min phōnēsas epea pteroenta prosēuda	15
I	hoi d' hote dē schedon ēsan ep' allēloisin iontes	10
I	ton d' apameibomenos prosephē podas ōkys achilleus	9
I	hōs eipōn otryne menos kai thymon hekastou	9
I	atreidē kydiste anax andrōn agamemnon	8
O	ton d' au tēlemachos pepnymenos antion ēuda	30
O	ton d' apameibomenos prosephē polymētis odysseus	25
O	ēmos d' ērigeneia phanē rododaktylos ēōs	20
O	tin d' apameibomenos prosephē polymētis odysseus	19
O	hōs ephat' autar ego min ameibomenos proseeipon	15

Table 3: Most frequently occurring verses in the *Iliad* (I) and *Odyssey* (O) verses, with total number of occurrences (#) in that work.

I.5.887	hē ke zōs amenēnos ea chalkoio typēsi
I.11.385	toxota lōbētēr kera aglae parthenopipa
I.5.723	chalkea oktanēma sidēreō axoni amphis
I.2.363	hōs phrētē phrētrēphin arēgē phyla de phylois
I.13.589	thrōskōsin kyamoi melanochroes ē erebinthoi
O.11.320	anthēsai pykasai te genys euanythei lachnē
O.19.177	dōriees te trichaikes dioi te pelasgoi
O.11.415	ē gamō ē eranō ē eilapinē tethalyiē
O.5.368	hōs d' anemos zaēs ēiōn thēmōna tinaxē
O.12.453	autis arizēlōs eirēmena mythologēuein

Table 4: The five most linguistically unexpected verses per poem (ranked by PPL descending).

Perplexity Perplexity scores for each verse were calculated as described in §3 above, using 5-gram LMs trained on sets of 23 books of the *Iliad* or *Odyssey*, respectively. In the *Iliad* (Fig. 1(a)), several peaks exist, with the top ten being considerably higher than the rest. The single most linguistically unexpected verse, as measured by PPL score, is the 887th verse of the 5th book (Table 4) followed by I.11.385. In the *Odyssey* (Fig. 1(b)), six of the most unexpected verses among the top ten come from books 10 to 12 (O.11.320, O.11.415, O.12.453, O.12.238, O.11.301, O.10.279). This particular concentration of linguistically unusual properties in Books 10 to 12 of the *Odyssey* dovetails with findings of Bozzone and Sandell (2022).

sic methodology may be readily applied to any other (set of) texts for any language. Once a unit of analysis (e.g., sentence, verse, paragraph, etc.) is established, character-level language models can be trained on the corpus (while holding out a given, typically larger unit, to avoid introducing biases), on which a Perplexity score is computed. This score can be interpreted as a measure of the degree of unexpectedness of that unit within the corpus. For example, a Perplexity score can be obtained for each sentence in the corpus of the novels of Jane Austen, by training seven language models (i.e., given seven published novels).

[3] By necessity, variant readings of specific forms in a given verse cannot be accounted for here. Other modern textual editions (van Thiel, 1991; West, 1998), make different choices, but the degree of difference is unlikely to substantially affect the overall results of the model. Future work could compare the Perplexity scores obtained for models trained on different textual editions; we hypothesize that the resulting sets of scores would be strongly positively correlated.

[4] Note that, among philological questions surrounding the text of the Homeric epics, divisions between lines — which constitute the relevant unit of analysis for this study — are among the least controversial matters, since these are, in most cases, easy to independently verify on the basis of the meter.

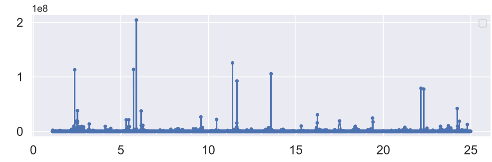

(a) Iliad

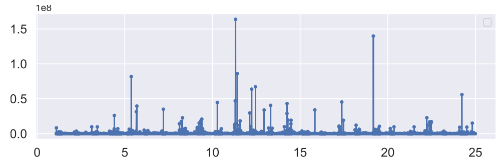

(b) Odyssey

Figure 1: PPL per verse across the books (horizontally) of the two Homeric poems.

5. Empirical analysis

As already discussed, the same verse may appear in an identical form up to nearly thirty times in a given poem (Table 3). At the same time, the occurrence of a particular lexical item in just a single line (i.e., if a given lexeme is a *hapax legomenon*) might result in an especially high PPL score. We therefore examined the extent to which an LM's predictions are affected by the frequencies of specific terms, character sequences, or the occurrence of particular named entities. In addition, by scoring each verse with two LMs, one per poem, we identified the verses that were surprising to the model trained on the source poem but not to that trained on the other poem.

5.1. Term frequency

Words We computed the number of verses in which each word of each poem is included, using its inverse to approximate the word's uniqueness in the poem.[5] The more verses a word is included in, the lower its inverse verse frequency (IVF) score will be. Words with a high IVF, for their part, might surprise an LM, yielding exceptionally high PPL. To assess such cases, we computed the minimum, maximum, and average IVF per verse, and measured their correlation with the poem's PPL scores. For both poems, Spearman's correlation coefficient was highest with the maximum IVF (0.669 for the *Iliad* and 0.702 for the *Odyssey*). After applying a base e logarithmic transformation to the IVF and PPL of each verse, Pearson's correlation coefficient between PPL and maximum IVF was similarly positive (0.581 for the *Iliad* and 0.618 for the *Odyssey*).[6] In effect, low-frequency lexemes, as measured by IVF, act as a good proxy for uncommon character transitions, which in turn directly affect the LM and PPL (see §3 above).

Character 5-grams We investigated the correlation between the token frequency of specific character 5-grams ($C5F$), composed of both word characters and whitespace (e.g., all the following are 5-grams [mēnin], [ēnin], [nin á]), within a given poem and the PPL score of a line. On the one hand, certain relatively short lexical items – mostly function words, such as *kaí* ('and') or *tís* ('some'), with accentual variants *kaì* and *tìs*, respectively – will cause certain character sequences to occur frequently in the corpus.[7] Thus, morphological and phonotactic factors may influence the frequency with which certain character sequences occur and in turn impact the extent to which an LM is surprised by a given line. We computed the $C5F$ in the *Iliad* and *Odyssey* separately.[8] Table 5 shows, as an example, the ten most frequent 5-grams in the *Iliad*; note, indeed, that many of these correspond to function words of two or three characters surrounded by whitespace. We then calculated correlation coefficients between the PPL of verses and the minimum, maximum, median, and mean $C5F$ in a given verse. Good negative correlations obtained between the minimum $C5F$ and PPL (better than -0.6), and even stronger correlations were found between median $C5F$ and PPL: Spearman's coefficient was -0.646 for the *Iliad* and -0.667 for the *Odyssey*; Pearson's coefficient (for base e log-transformed data) was -0.688 for the *Iliad* and -0.715 for the *Odyssey*. Thus, the lower the median frequency of all 5-grams in a verse, the greater the LM's surprise as measured by PPL.

IVF-C5F As intuitively should be expected, a reasonable negative correlation likewise holds between the maximum IVF and the minimum $C5F$ found in a verse. For the *Iliad*, we found a correlation of -0.602 (Spearman) and -0.660 (Pearson, after base e log-transformation). For the *Odyssey*, we found a correlation of -0.645 (Spearman) and -0.780 (Pearson, after base e log-transformation). Such a correlation is unsurprising, since a word form that occurs infrequently, if it contains a rare character sequence, will appreciably affect the total number of occurrences of that sequence.

[5] We trained scikit-learn's TfIdfVectorizer (Pedregosa et al., 2011), using default values, and the `idf_` attribute. Note that with default these default settings, words consisting of just a single character are excluded.

[6] Baayen (2001) treats the typically non-linear distribution of word frequencies in natural language texts.

[7] Sequences may also be (in)frequent due to the language's morphology and phonotactics (cf. Hayes and Wilson (2008) on the quantitative modeling of phonotactic preferences).

[8] `NLTK's ngram` module (Bird et al., 2009) was used.

Character 5-gram	Frequency
(, k, a, ì,)	2352
(, m, è, n,)	933
(n, , d, ’,)	864
(s, ,d, ’,)	639
(, d, ’, , e)	637
(, a, kh, a, i)	622
(, g, à, r,)	547
(,d, ’, , á)	533
(, d, ’, , a)	524
(a, l, l, ’,)	494

Table 5: Ten most frequent character 5-grams (incl. whitespace) in the *Iliad* (ranked descending), given here in transliteration.

5.2. Named Entities

The third and final baseline of testing consisted in the frequency of named entities, more specifically, personal names like *Achilleus*, per verse. Named entities were automatically identified using a Transformer-based recognizer trained in Ancient Greek (Yousef et al., 2022).[9] Most verses (in the *Iliad*, 65%; in the *Odyssey*, 73%) contain no personal names; no verse was automatically recognized as containing potentially more than four (only eight verses in the *Iliad*). Correlation (Pearson) between the number of personal names in a line and the line's PPL score proved to be negligible: -0.007 (*Iliad*) and -0.003 (*Odyssey*). We thus conclude that the presence of personal names in and of themselves does not materially impact the degree of linguistic unexpectedness of a verse.

5.3. Cross-poem modeling

We hypothesize that linguistically unexpected verses might be less unexpected for a model trained on verses of the other poem. To evaluate this hypothesis, we trained one LM per poem on all verses in that poem, and then used both models to score each verse per poem and to compute their difference in PPL that is described next. Given a source model S and another model O, the difference d for verse v is computed as: $d^v(S, O) = PPL(v, S) - PPL(v, O)$.[10] A verse with a large positive difference indicates that the verse is significantly more unexpected to a model trained on the entirety of the source poem than a model trained on the entirety of the other poem. For example, the famous opening verse of Odyssey (*andra moi ennepe mousa polutropon hos mala polla*) is much less surprising to the model trained on the *Odyssey* ($PPL = 305.6$) than to the model trained

[9]See huggingface.co/UGARIT/flair_grc_multi_ner.

[10]No verses were removed during training, thus giving an advantage to source models and increasing the difficulty that O would have in obtaining a lower PPL.

I6.490	all’ eis oikon iousa ta s’ autēs erga komize	5059.86
I1.485	nēa men hoi ge melainan ep’ ēpeiroio eryssan	4717.82
I13.821	hōs ara hoi eiponti epeptato dexios ornis	3963.98
O24.488	bē de kat‘ oulympoio karēnōn aixasa	4789.3
O22.124	hippourin deinon de lophos kathyperthen eneuen	3075.1
O11.270	tēn echen amphitryōnos uios menos aien ateirēs	2706.8

Table 6: Verses relatively surprising for a model trained on verses from the source poem, but less so for a model trained on verses from the other.

on the *Iliad* ($PPL = 13487.6$, thus the difference: $d = -13182.0$).[11] Table 6 presents the verses with the largest difference per poem. Overall, the proportion of verses with a positive cross-score is small for each poem (3.25% for the *Iliad*, 2.24% for the *Odyssey*), but the difference between *Iliad* and *Odyssey* in this regard is statistically significant (2-sample test for equality of proportions: χ^2 = 25.17, $p < 0.01$); this may point to an overall lower level of linguistic homogeneity in the *Iliad*.

6. Discussion and Conclusions

While neither the inverse verse frequency of a word form nor the token frequency of character 5-grams perfectly accounts for the behavior of the LM, both the frequency with which particular lexemes are employed as well as phonotactics (as partly captured by 5-gram frequency) must be factors in the raw natural language data that have a significant impact on the LM. The presence of a named entity in a verse, in contrast, has no significant impact; although a given named entity may itself be rare or contain a rare character sequence, others (e.g., *Zeus*) occur dozens of times, and precisely the presence of forms that occur frequently will tend to drive down the PPL score of a line. The HoLM dataset presented here is publicly available at https://github.com/ipavlopoulos/holm, and can be used by experts to detect verses that are unexpected to an LM trained on verses in one but not the other poem, or verses unexpected to the LM but which do not contain infrequent terms (i.e., by monitoring the $\frac{PPL}{IVF}$ ratio). Going forward, we intend to explore how PPL scores could be used for pedagogical purposes. First, we will hypothesize that verses with high PPL scores might also be verses that are more difficult for students of Ancient Greek to translate, and vice-versa. Second, for the study of Homeric style, we will hypothesize that PPL scores might aid us in distinguishing more and less “traditional” portions of the Homeric poems, with verses with lower scores being more “traditional” and formulaic, and verses with higher scores possibly showing more individual poetic style.

[11]Interestingly, the PPL for O1.1 for a model trained on all books of the *Odyssey* except Book 1 (84939.05) is much higher than the PPL of the full *Iliad* model.

Ethics Statement

The resource presented here is intended to provide experts with a tool for researching issues regarding the authorship of the Homeric poems. The broader impact of this resource is, however, not limited to use by experts on authorship analysis, but also extends to educational purposes. By identifying the verses most and least expected by a language model (trained on a given poem), instructors of Ancient Greek, and of Homer specifically, can develop more targeted teaching materials. By focusing on the most expected verses, for instance, instructors can acquaint students with some of the most typical examples Homeric language. Conversely, the least expected verses could be used to present students with special challenges, which could be exploited both for purposes of instruction and assessment. Furthermore, as was noted in §3, similar applications are in principle possible for any corpus of natural-language texts. Finally, the authors declare that their involvement in this research does not entail any conflicts of interest.

References

R. Harald Baayen. 2001. *Word Frequency Distributions*. Kluwer, Dordrecht.

Steven Bird, Edward Loper, and Ewan Klein. 2009. *Natural Language Processing with Python*. O'Reilly Media, Sebastopol, CA.

Chiara Bozzone. 2014. *Constructions: A New Approach to Formularity, Discourse, and Syntax in Homer*. Ph.D. thesis, University of California, Los Angeles.

Chiara Bozzone. 2024. *Homer's Living Language*. Cambridge University Press, Cambridge.

Chiara Bozzone and Ryan Sandell. 2022. One or many Homers? Using quantitative authorship analysis to study the Homeric question. In *Proceedings of the 32nd Annual UCLA Indo-European Conference: November 5th, 6th, and 7th, 2021*, pages 21–48. Helmut Buske Verlag.

Albio Cesare Cassio. 2002. Early editions of the Greek epics and Homeric textual criticism in the sixth and fifth centuries B.C. In Franco Montanari and Paola Ascheri, editors, *Omero tremila anni dopo*, pages 105–136. Edizioni di storia e letteratura, Roma.

Georg Danek. 1988. *Studien zur Dolonie*. Österreichische Akademie der Wissenschaften, Vienna.

Georg Danek. 2012. The Doloneia revisited. In Øvind Andersen and Dag T. T. Haug, editors, *Relative Chronology in Early Greek Epic Poetry*, pages 106–121. Cambridge University Press, Cambridge.

Rotem Dror, Lotem Peled-Cohen, Segev Shlomov, and Roi Reichart. 2020. Statistical significance testing for natural language processing. *Synthesis Lectures on Human Language Technologies*, 13(2):1–116.

Maria Fasoi, John Pavlopoulos, and Maria Konstantinidou. 2021. Computational authorship analysis of Homeric language. In *The 1st Digital Humanities Workshop*, Kiev, Ukraine.

John Miles Foley. 2007. "Reading" Homer through oral tradition'. *College Literature*, 34:1–28.

Alex Graves. 2013. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*.

Michael Haslam. 2011. Homeric papyri and the transmission of the text. In Ian Morris and Barry Powell, editors, *A New Companion to Homer*, pages 55–101. Brill, Leiden.

Bruce Hayes and Colin Wilson. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry*, 39:379–440.

Kyuyeon Hwang and Wonyong Sung. 2017. Character-level language modeling with hierarchical recurrent neural networks. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5720–5724. IEEE.

Minna Skafte Jensen. 2011. *Writing Homer: A Study Based on Modern Fieldwork*. The Royal Danish Academy of Sciences and Letters, Copenhagen.

Patrick Juola. 2006. Authorship attribution. *Foundations and Trends in Information Retrieval*, 1:233–334.

Jakub Kabala. 2020. Computational authorship attribution in medieval Latin corpora: the case of the Monk of Lido (ca. 1101–08) and Gallus Anonymous (ca. 1113–17). *Lang. Resources &. Evaluation*, 54(1):25–56.

Mike Kestemont, Kim Luyckx, Walter Daelemans, and Thomas Crombez. 2012. Cross-Genre Authorship Verification Using Unmasking. *English Studies*, 93(3):340–356.

Albert Bates Lord. 1960. *The Singer of Tales*. Harvard University Press, Cambridge, MA.

Harold Love. 2002. *Attributing Authorship: An Introduction*. Cambridge University Press, Cambridge, England, UK.

Nikos Manousakis and Efstathios Stamatatos. 2018. Devising rhesus: A strange 'collaboration' between Aeschylus and Euripides. *Digital Scholarship in the Humanities*, 33(2):347–361.

David B. Monro and Thomas W. Allen, editors. 1920. *Homeri Opera*, 3rd edition. Clarendon Press, Oxford.

A. T. Murray, editor. 1919. *Homer, Odyssey*. Harvard University Press, Cambridge, MA.

Gregory Nagy. 2004. *Homer's Text and Language*. University of Illinois Press.

Gregory Nagy. 2009. *Homer the Classic*. Center for Hellenic Studies, Washington, DC.

Heinz-Günter Nesselrath. 2011. Geschichte der Homerforschung. In Antonios Rengakos and Bernhard Zimmermann, editors, *Homer-Handbuch: Leben – Werk – Wirkung*, pages 175–198. Metzler, Stuttgart.

Milman Parry. 1971. *The Making of Homeric Verse: The Collected Papers of Milman Parry*. Oxford University Press, Oxford.

John Pavlopoulos and Maria Konstantinidou. 2023. Computational authorship analysis of the Homeric poems. *International Journal of Digital Humanities*, 5:45–64.

Fabian Pedregosa, Gael Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Jonathan Ready. 2019. *Orality, Textuality, and the Homeric Epics: An Interdisciplinary Study of Oral Texts, Dictated Texts, and Wild Texts*. Oxford University Press, Oxford.

Efstathios Stamatatos. 2009. A survey of modern authorship attribution methods. *Journal of the American Society for Information Science and Technology*, 60(3):538–556.

Frank Turner. 2011. The Homeric question. In Ian Morris and Barry Powell, editors, *A New Companion to Homer*, pages 123–145. Brill, Leiden.

Helmut van Thiel, editor. 1991. *Homeri Odyssea*. Hildesheim, Georg Olms.

Rudolf Wachter. 2007. Greek dialects and epic poetry: Did Homer have to be an Ionian? In Miltiade Hatzopoulos, editor, *Φωνης χαρακτ ρ εθνικ ς: Actes du Ve congrès internationale de dialectologie greque*, pages 317–328. Boccard, Paris.

Martin L. West, editor. 1998. *Homeri Ilias*. Teubner, Stuttgart.

Martin L. West. 2011. *The Making of the Iliad: Disquisition and Analytical Commentary*. Oxford University Press, Oxford.

Martin L. West. 2013. *The Making of the Odyssey*. Oxford University Press, Oxford.

Tariq Yousef, Chiara Palladino, and Stefan Jänicke. 2022. Transformer-based named entity recognition for Ancient Greek.