

Building UD Cairo for Old English in the Classroom

Lauren Levine, Junghyun Min, and Amir Zeldes

Georgetown University, Department of Linguistics
{le176, jm3743, amir.zeldes}@georgetown.edu

Abstract

In this paper we present a sample treebank for Old English based on the UD Cairo sentences, collected and annotated as part of a classroom curriculum in Historical Linguistics. To collect the data, a sample of 20 sentences illustrating a range of syntactic constructions in the world’s languages, we employ a combination of LLM prompting and searches in authentic Old English data. For annotation we assigned sentences to multiple students with limited prior exposure to UD, whose annotations we compare and adjudicate. Our results suggest that while current LLM outputs in Old English do not reflect authentic syntax, this can be mitigated by post-editing, and that although beginner annotators do not possess enough background to complete the task perfectly, taken together they can produce good results and learn from the experience. We also conduct preliminary parsing experiments using Modern English training data, and find that although performance on Old English is poor, parsing on annotated features (lemma, hyperlemma, gloss) leads to improved performance.

1 Introduction

Treebanking historical languages is a challenging task for multiple reasons – to name just a few: 1. annotators are never native speakers; 2. it is difficult to recruit annotators with sufficient background in the language, in theoretical syntax, and in corpus annotation; 3. guidelines and sample datasets illustrating annotation principles are often tailored to modern languages, with examples that would be implausible in context for ancient texts; 4. there is a scarcity of previous work for annotators to reliably refer to, making examples for contexts that are common in ancient texts difficult to find.

At the same time, tackling ancient language syntax through treebanking can be a rewarding and pedagogically valuable way of confronting students of ancient languages and historical linguistics with

authentic syntax in ways that simply looking at an original sentence and its translation cannot (Bamman and Crane, 2007; Mambrini, 2016; Keersmaekers et al., 2019). An additional challenge here is defining a realistic target for treebanking which can be completed as a homework assignment without becoming overwhelming.

In Universal Dependencies (UD, de Marneffe et al. 2021), the most popular treebanking formalism, there is no single template for starting a treebanking project. However, a set of basic sentences is sometimes suggested, known as the Cairo sentences and originally proposed at CICLing 2015 in Cairo.¹ The data has been translated and used as a UD starting point for a number of languages, including Ligurian (Lusito and Maillard, 2021), Luxembourgish (Plum et al., 2024) and Gujarati (Jobanputra et al., 2024). In this paper, we therefore propose and report on a classroom historical treebanking activity creating a UD Cairo-style corpus for Old English (Ald-Ænglisc or OE), including generating the text, treebanking proper, and adding supplementary annotations.

Although some work on treebanking OE and later stages of older English exists – notably The York–Toronto–Helsinki Parsed Corpus of Old English Prose (YCOE; Taylor, 2007), the Penn-Helsinki Parsed Corpus of Early Modern English (PPCEME; Kroch et al., 2004), and their follow-up work (Kroch, 2020; Kulick et al., 2023) – no such data exists in UD. We add to the literature by providing the first OE contribution to UD. We evaluate treebanking from students with and without previous exposure to UD, and run some parsing experiments to reveal how similar or distinct OE syntax is compared to Modern English (ME). We release all data publicly² and hope this work in-

¹<https://github.com/UniversalDependencies/cairo>

²<https://universaldependencies.org/treebanks/ang-cairo/index.html>

spires others to incorporate more treebanking into classroom activities.

2 LLM-assisted UD Cairo in Ald-Ænglisc

Translating the UD Cairo sentences to Old English is a non-trivial task, especially considering the goal of expressing the basic constructions of the language in the most natural possible way, without introducing translation effects conforming to the Modern English syntax or lexicon. As a first step, we considered simply asking an LLM, in this case GPT4o, to translate the sentences (see Appendix A for the full list of sentences with initial GPT outputs and final versions).

Inspecting GPT’s output revealed that while the model was surprisingly capable of translating the Cairo sentences into superficially correct Old English, many problems emerged, especially in morphology and the tendency to mimic the words and word order of the input one for one. For example, consider sentence (1):

- (1) a. The girl wrote a letter to her friend
- b. Sēo mæden wrāt ænne ærendgecwidd
 tō hyre frēonde.

The model translated ‘the girl’ into ‘seo mæden’ - a correct lexical choice with an incorrect article gender (feminine, when ‘mæden’ was in fact neuter in OE). The model also tended to insert indefinite articles (here ‘ænne’), which were not yet grammaticalized or expected in OE, where the Modern English original had them, resulting in superfluous ‘ones’ (e.g. ‘one letter’ for ‘a letter’, where just ‘letter’ would be correct). We also note that GPT recreates the input word order with the ‘to’ dative, rather than using the more usual ditransitive construction, (‘wrote her friend a letter’), which was used more in OE. Finally, the word for ‘letter’ is a creative invention, compounded from two attested OE stems, but which is not attested as a word. In the final data, we therefore corrected the gender error, removed the indefinite article, used the ditransitive word order, and used Bosworth Toller’s Anglo-Saxon Online Dictionary³ to select an attested word for ‘letter’, resulting in (2):

- (2) Pæt mæden wrat hyre freonde ærend-writ

At the same time, GPT was adept at proposing creative solutions to vocabulary gaps: for example,

³<https://bosworthtoller.com/>

Cairo sentence 6 refers to washing a car, to which GPT, unprompted, added: “(Note: ‘car’ is anachronistic, but Old English did have ‘cār,’ meaning ‘chariot’ or ‘cart.’)”, and selected OE *cræt* ‘cart’ for its own translation, which we adopted. In some cases, reprompting for additional translation ideas was also an effective strategy.

As a final source to ensure our chosen constructions corresponded to attested Old English syntax, we cross-referenced sentences containing target constructions in a Modern English New Testament translation with the OE Wessex Gospels, the oldest existing manuscript of the Gospels in OE. For example, GPT’s initial version of Cairo sentence 17 contains a counterfactual ‘she should have been doing’, which GPT rendered with *heo scolde þæt ... dyde*, roughly corresponding to a Modern English ‘she should that did’. However searching the gospels for ‘should have’, we quickly realized the most common equivalent in the Wessex Gospels is actually *hit gebyrede þæt heo dyde*, i.e. “it was appropriate that she did...”.

In summary, our approach to producing an ancient language version of the Cairo sentences is: 1. ask an LLM to translate the sentence; 2. check the result for morphological errors; 3. ensure chosen words are attested in a dictionary and replace/reprompt for alternatives as needed; 4. check a modern translation for equivalent modern construction instances, then ensure their original counterparts correspond to the selected constructions.

3 Annotation

3.1 UD trees and classroom annotation

Our primary objective was to design a classroom learning activity for an audience that might not be familiar with UD guidelines, rather than to create data for OE parsing. As part of a mid-level, mixed over/undergraduate course on Historical Linguistics, we formulated an assignment in the Old English module of the class in which students were asked to treebank four random OE Cairo sentences. We allowed students to access both UD corpora (using the Grew-match interface, [Guillaume 2021](#)) and Modern English parsers – we pointed them to online interfaces for UDPipe ([Straka, 2018](#)) and Stanza ([Qi et al., 2020](#)), but any parser was allowed.

Since the students mostly had only minimal instruction about the UD guidelines, they were instructed to look in existing UD data for similar constructions to the sentence they were annotat-

ing, and to play with the parsers by feeding them what they felt were modern English equivalents (the original Cairo sentences in English were not provided, though students may have found them online too). Use of AI agents was also allowed, though we used the Arborator graphical interface (Guibon et al., 2020), meaning students could not simply let LLMs solve the assignment for them directly.

Figure 1 plots students’ accuracy on the assignment in terms of unlabeled and labeled attachment scores (UAS and LAS). We note that some of the students in the class, who had background in Computational Linguistics and previous exposure to UD annotation, unsurprisingly did better on average in terms of accuracy, and therefore we report scores separately for those two annotator subsets.

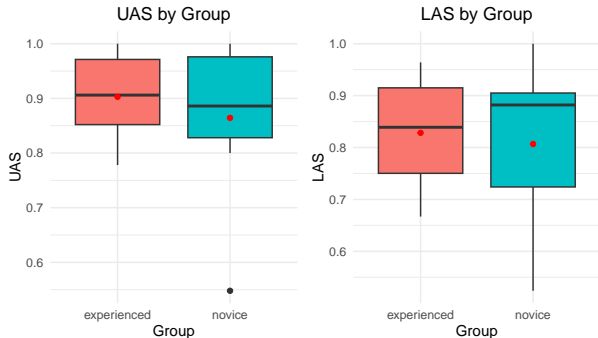


Figure 1: Unlabeled and labeled attachment scores (UAS & LAS) for experienced and novice annotators.

As Figure 1 shows, while the mean scores (indicated with red dots) are somewhat higher for the experienced group, novice students were able to achieve rather comparable scores by using corpus searches and automatic modern English parsers as references, and the median novice LAS score is even higher than the experienced one. At the same time, the range of scores is wider for novices, which we suspect is due to the heterogeneity of their background in general linguistics, with students who had more background in theoretical (but non-UD) syntax achieving better accuracy.

We report inter-annotator agreement (IAA) metrics for our gold reference annotations (accuracy, Cohen’s Kappa (Cohen, 1960), and F1-score) between two groups of experienced Modern English UD annotators in Table 1. We also include IAA accuracy on section 23 of the Penn Treebank (Marcus et al., 1993) measured by Berzak et al. (2016) as a point of reference. As our numbers are similar to Penn Treebank ones, we believe experience

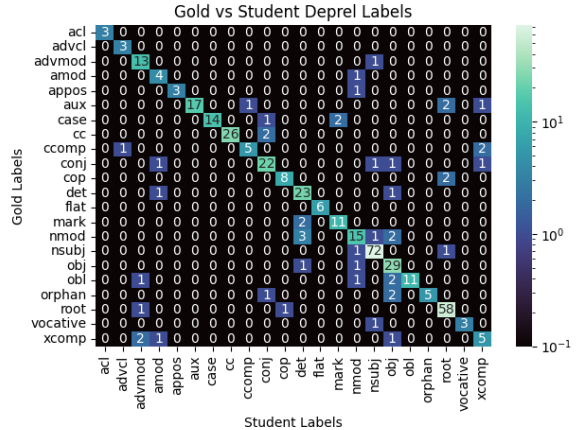


Figure 2: Confusion matrix for gold dependency relation (deprel) labels vs the student annotations. Deprel subtypes are collapsed, only deprel labels that occur > 5 times are included, and punct was excluded.

with ME UD annotation translates non-trivially to performance on OE UD annotation.

IAA score	UAS	LAS
UD Cairo OE		
acc.	94.15	88.30
κ	93.38	88.27
F1	96.99	93.79
Reference		
PTB-WSJ (acc.)	93.07	88.32

Table 1: Inter-annotator agreement metrics: accuracy, Cohen’s Kappa, and F1-Score. We present an IAA metric on PTB-WSJ by Berzak et al. (2016) for reference.

3.2 Disagreement Analysis

In Figure 2, we show a confusion matrix comparing our final gold labels for the Cairo UD sentences with the student annotations from our classroom activity. In line with our high inter-annotator agreement scores in the previous section, we see that there is strong overlap in the selection of dependency relation labels. However, there is still a smattering of disagreements between the gold labels and the student labels. The dependency relation that was mislabeled by students the most was nmod, which was mislabeled 6 times (3 times as det, 2 times as obj, and 1 time as nsubj). These confusions primarily stemmed from students not understanding when to apply the nmod:poss subtype when possessive case was involved in the sentence. The label that was most misapplied by students was obj, which was misapplied a total of 9 times

(2 times each on *nmod*, *obl*, and *orphan* relations, and 1 time each as on *conj*, *det*, and *xcomp* relations). In Figure 3, we show a dependency analysis for one of the UD Cairo sentences where a student both misapplied the *obj* label rather than utilizing the orphan relation, as well as incorrectly using *det* for a possessive pronoun. We see that the student uses a promotion strategy, rather than utilizing orphan. Likely the use of more complicated relations, such as orphan, could be learned sufficiently with more time devoted to annotator training.

3.3 Additional annotations

In addition to the classroom UD annotations that comprise part of speech tags and dependency labels and their attachments, we provide morphology, hyperlemma, and root annotations (a sample analysis is included in Appendix B).

Morphological features: We follow UD’s standard CoNLL-U format and provide annotations on case, number, gender, tense, person, verb-form, and mood. In our corpus, we count four cases (nominative, accusative, genitive, and dative), two numbers (singular and plural), three genders (masculine, neuter, and feminine), two tenses (past and present), three persons (1st, 2nd, and 3rd), three verb-forms (finite, infinitive, and participle), and two moods (indicative and imperative). However, the above list is not exhaustive for Old English morphology. In submitting our corpus to UD, we define Old English’s language-specific morphological feature space as outlined in Appendix C.

Gloss: We provide word-level glosses in Modern English, whose uninflected base forms are often equivalent to the hyperlemma described below.

Hyperlemmas: Hyperlemmas, first proposed in Dipper et al. (2004), align word forms and lemmas from different language branches and historical periods that correspond to each other. For example, a hyperlemma *have* would normalize German *haben*, English *have*, Old English *habban* and their inflected forms. We follow several annotated corpora of historic European languages (Meyer, 2011; Kučera and Stluka, 2014; Shokina and Mocken, 2016; Chiarcos et al., 2018) and include hyperlemmas in our annotations.

Roots and language of origin: Roots are another method of normalizing within language family and across historical periods. For example, German *trinken*, English *drink*, and Old English *drincan* can all be associated with their reconstructed Indo-European root **d^hrenǵ*. While it is more com-

Input	ELECTRA		mBERT	
	UAS	LAS	UAS	LAS
GUM 11 dev	95.35	93.17	94.21	91.95
Form	23.98	12.87	41.52	25.73
+Norm	35.09	20.47	41.52	27.49
Lemma	18.13	7.60	38.60	22.81
+Norm	32.16	16.96	39.18	23.39
Hyperlemma	61.99	49.12	70.76	56.73
+Norm	74.27	59.06	83.04	67.84
Gloss	81.87	75.44	85.38	78.36

Table 2: Unlabeled and labeled attachment scores (UAS & LAS) for parsers trained on ME data (GUM V11).

mon for Afro-Asiatic corpora to contain root information (e.g. Hajič et al., 2009), we provide root information in addition to hyperlemmas. If the root is of non-Indo-European origin, its language of origin is also specified.

4 Parsing experiments

In order to shed light on how much Old English dependency syntax analysis differs from that of Modern English, we conduct parsing experiments on our OE Cairo sentences using dependency parsers trained on ME. While it may be possible and potentially effective to use dependency parsers trained on other modern and historical Germanic languages, we limit the scope our parsing experiments to English. We train DiaParser (Dozat and Manning, 2017), a neural biaffine dependency parser, using English ELECTRA (Clark et al., 2020) and multilingual mBERT (Devlin et al., 2019) embeddings, in order to compare the impact of monolingual vs multilingual pre-trained inputs. For training data, we use version 11 of the GUM corpus (Zeldes, 2017), which is the largest available UD corpus of Modern English, containing 268k tokens across a variety of genres. We evaluate the ELECTRA and mBERT models on our OE Cairo sentences, using our various token annotations (Form, Lemma, Hyperlemma, and Gloss) as input. The results of these experiments are reported in Table 2, giving the UAS and the LAS for each model. Form refers to the unaltered text of the tokens, and Lemma refers to the uninflected form of the tokens. Hyperlemma and Gloss are as described in Section 3.3. +Norm indicates the addition of orthographic normalization to the annotation of the preceding row – replacing OE letters that no longer exist in

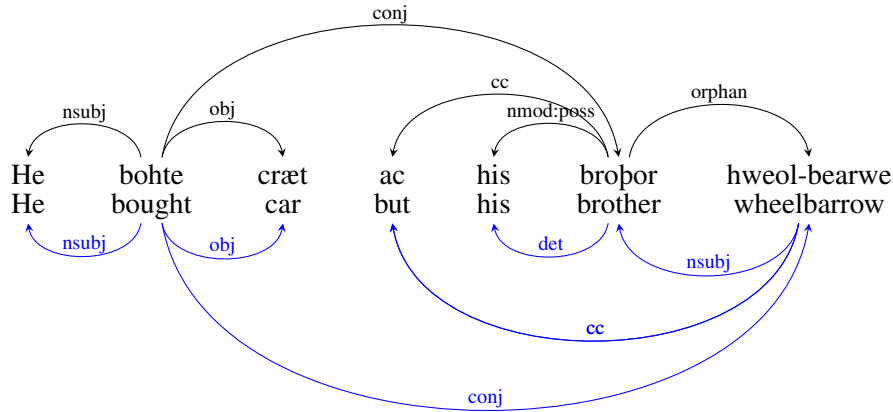


Figure 3: Dependency analysis of an ellipsis construction in OE: *He bought a car but his brother just a wheelbarrow*. ME gloss is included below the OE tokens. The gold analysis is shown on top in black, and a student analysis is shown on the bottom in blue.

ME ("th" for "þ", "e" for "æ", etc.).

Looking at Table 2, we see that although the ELECTRA model has higher performance on the Modern English GUM 11 dev data, mBERT uniformly performs better on the Old English data for all annotation categories. This demonstrates the advantage that multilingual pretrained embeddings provide when a parsing model is used on a different language than the one used during model training. We also note that for the categories of Form and Lemma, orthographic normalization leads to large improvements for the ELECTRA model with monolingual embeddings, but only minor improvement to the mBERT model with multilingual embeddings. However, both models show a large increase in LAS between Hyperlemma and Hyperlemma+Norm, with an increase of 9.94% for the ELECTRA model and an increase of 11.11% for the mBERT model.

For both models, performance increases from Form to Hyperlemma and again to Gloss, for both UAS and LAS, with the mBERT model giving the highest score of UAS 85.38% and LAS 78.36%. Using Hyperlemma and Gloss for the model input allows us to abstract away from the lexical differences between Old English and Modern English. Although parsing the Gloss leads to a large improvement in performance over parsing the Form (LAS Δ 62.65% for the ELECTRA model and LAS Δ 52.63% for the mBERT model), there is still a large drop in performance when compared to the ME dev data (LAS Δ 17.73% for ELECTRA and LAS Δ 13.59% for mBERT). Although some degradation can be expected in cross-corpus pars-

ing (Zeldes and Schneider, 2023), its magnitude indicates that there are notable syntactic differences between OE and ME for dependency parsing, even after controlling for lexical differences. Still, the scores for parsing on glossed tokens of Old English are high enough to suggest that leveraging Modern English models to parse glossed Old English texts could be an effective method to bootstrap the creation of Old English syntactic dependency datasets in the future.

5 Conclusion

In this paper, we presented an Old English version of the UD Cairo sentences, annotated for syntactic dependencies, morphological features, and historically relevant features of hyperlemma and root. In using LLMs for the initial generation of the OE version of the data, we demonstrate how current AI tools can be used in conjunction with quality assurance procedures in the workflow of data creation to produce more reliable results. By including the treebanking task as part of a course in Historical Linguistics, we show that a brief introduction to UD in the classroom, along with providing reference resources, can result in relatively good performance, even for a language with which most of the students have no familiarity. Additionally, through preliminary parsing experiments, we demonstrate that dependency parsers trained on Modern English data may be of use for the creation of new syntactic dependency datasets for Old English. We hope that the sample Old English treebank presented here can serve as a guide for future classroom annotation efforts.

Acknowledgments

We would like to thank the students in the course Corpus Approaches to Historical Linguistics at Georgetown University for participating in the course and the creation of the dataset. The final public repository will include all annotation versions of all students who agreed to release their work and signed an open license agreement. Credit is given by name to all students who wished to be named, as well as anonymously to students who wish to remain anonymous.

References

- David Bamman and Gregory Crane. 2007. [The Latin dependency treebank in a cultural heritage digital library](#). In *Proceedings of the Workshop on Language Technology for Cultural Heritage Data (LaTeCH 2007)*, pages 33–40, Prague, Czech Republic. Association for Computational Linguistics.
- Yevgeni Berzak, Yan Huang, Andrei Barbu, Anna Korhonen, and Boris Katz. 2016. [Anchoring and agreement in syntactic annotations](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2215–2224, Austin, Texas. Association for Computational Linguistics.
- Christian Chiarcos, Benjamin Kosmehl, Christian Fäth, and Maria Sukhareva. 2018. Analyzing middle high german syntax with rdf and sparql. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Kevin Clark, Thang Luong, Quoc V. Le, and Christopher Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#). In *ICLR*.
- Jacob Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and Psychological Measurement*, 20(1):37–46.
- Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. [Universal Dependencies](#). *Computational Linguistics*, 47(2):255–308.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Stefanie Dipper, Lukas Faulstich, Ulf Leser, and Anke Lüdeling. 2004. Challenges in modelling a richly annotated diachronic corpus of german. In *Workshop on XML-based richly annotated corpora*, pages 21–29.
- Timothy Dozat and Christopher D. Manning. 2017. [Deep biaffine attention for neural dependency parsing](#). In *International Conference on Learning Representations*.
- Gaël Guibon, Marine Courtin, Kim Gerdes, and Bruno Guillaume. 2020. [When collaborative treebank curation meets graph grammars](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 5291–5300, Marseille, France. European Language Resources Association.
- Bruno Guillaume. 2021. [Graph matching and graph rewriting: GREW tools for corpus exploration, maintenance and conversion](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 168–175, Online. Association for Computational Linguistics.
- Jan Hajič, Otakar Smrž, Petr Zemánek, Petr Pajas, Jan Šnidauf, Emanuel Beška, Jakub Kracmar, and Kamila Hassanová. 2009. [Prague arabic dependency treebank 1.0](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Mayank Jobanputra, Maitrey Mehta, and Çağrı Çöltekin. 2024. [A Universal Dependencies treebank for Gujarati](#). In *Proceedings of the Joint Workshop on Multiword Expressions and Universal Dependencies (MWE-UD) @ LREC-COLING 2024*, pages 56–62, Torino, Italia. ELRA and ICCL.
- Alek Keersmaekers, Wouter Mercelis, Colin Swaelens, and Toon Van Hal. 2019. [Creating, enriching and valorizing treebanks of Ancient Greek](#). In *Proceedings of the 18th International Workshop on Treebanks and Linguistic Theories (TLT, SyntaxFest 2019)*, pages 109–117, Paris, France. Association for Computational Linguistics.
- Anthony Kroch. 2020. [Penn parsed corpora of historical english](#).
- Anthony Kroch, Beatrice Santorini, and Lauren Delfs. 2004. [Penn-helsinki parsed corpus of early modern english](#).
- Seth Kulick, Neville Ryant, and Beatrice Santorini. 2023. [Parsing “early English books online” for linguistic search](#). In *Proceedings of the Society for Computation in Linguistics 2023*, pages 222–242, Amherst, MA. Association for Computational Linguistics.
- Karel Kučera and Martin Stluka. 2014. [Data processing and lemmatization in digitized 19th-century czech texts](#). In *Proceedings of the First International Conference on Digital Access to Textual Cultural Heritage, DATeCH ’14*, page 193–196, New York, NY, USA. Association for Computing Machinery.

Stefano Lusito and Jean Maillard. 2021. [A Universal Dependencies corpus for Ligurian](#). In *Proceedings of the Fifth Workshop on Universal Dependencies (UDW, SyntaxFest 2021)*, pages 121–128, Sofia, Bulgaria. Association for Computational Linguistics.

Francesco Mambrini. 2016. [The Ancient Greek Dependency Treebank: Linguistic annotation in a teaching environment](#). In M. Romanello and G. Bodard, editors, *Digital Classics Outside the Echo-Chamber*. Ubiquity Press, London.

Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. [Building a large annotated corpus of English: The Penn Treebank](#). *Computational Linguistics*, 19(2):313–330.

Roland Meyer. 2011. New wine in old wine-skins?—tagging old russian via annotation projection from modern translations. *Russian linguistics*, 35(2):267–281.

Alistair Plum, Caroline Döhmer, Emilia Milano, Anne-Marie Lutgen, and Christoph Purschke. 2024. [LuxBank: The first Universal Dependency treebank for Luxembourgish](#). In *Proceedings of the 22nd Workshop on Treebanks and Linguistic Theories (TLT 2024)*, pages 30–39, Hamburg, Germany. Association for Computational Linguistics.

Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.

Nina Yu Shokina and Susanne Mocken. 2016. A text mining system for creating electronic glossaries with application to research of church slavonic language. *Vychislitel'nyye tekhnologii*, 21(4).

Milan Straka. 2018. [UDPipe 2.0 prototype at CoNLL 2018 UD shared task](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 197–207, Brussels, Belgium. Association for Computational Linguistics.

Ann Taylor. 2007. [The York—Toronto—Helsinki parsed corpus of old english prose](#), pages 196–227. Palgrave Macmillan UK, London.

Amir Zeldes. 2017. [The GUM corpus: Creating multilayer resources in the classroom](#). *Language Resources and Evaluation*, 51(3):581–612.

Amir Zeldes and Nathan Schneider. 2023. [Are UD treebanks getting more consistent? a report card for English UD](#). In *Proceedings of the Sixth Workshop on Universal Dependencies (UDW, GURT/SyntaxFest 2023)*, pages 58–64, Washington, D.C. Association for Computational Linguistics.

A GPT initial outputs

GPT’s initial renditions of the Old English Cairo sentences are given below, in italics. Note that these include clear errors. For reference, original Old English sentences from the Cairo dataset in quotes and the final Old English translations for our UD Old English Cairo data are also supplied for easy comparison.

1. *Sēo mæden wrāt ænne ærendgecwidd tō hyre frēonde.*
“The girl wrote a letter to her friend.”
Pæt mæden wrat hyre freonde ærend-writ.
2. *Ic þence þæt hit rīne.*
“I think that it is raining.”
Ic þence þæt hit regnaþ.
3. *Hē gewende tō geswīgenne smōkunge and drincan.*
“He tried to stop smoking and drinking.”
He gesohte to geswicanne smican and drincan.
4. *Wilt þū gān?*
“Do you want to go?”
Wilt þu gan?
5. *Sam, ontyne þæt ēagþýrel!*
“Sam, open the window!”
Sam, ontyne þæt eagþyrel!
6. *Hēo dyde þæt hyre efenġemacca þæt cræt gewēasce*
“She made her husband wash the car.”
Heo dyde þæt hyre efen-gemacca þæt cræt geweasce.
7. *Pēteres nēahġebūr smealde þone stoccweall rēad.*
“Peter’s neighbor painted the fence red.”
Peteres neah-gebur hæfð þone eodor readne ongemet.
8. *Mīn fæder is cealdra þonne þīn.*
“My dad is cooler than yours.”
Mīn fæder is micelre þonne þin.
9. *Maria gewann brons, Pēter seolfor, and Ġeane gold.*
“Mary won bronze, Peter silver, and Jane gold.”
Maria gewann ær, Petrus seolfor, and Iohanna gold.
10. *Is Iguazu micel rīce oþþe lýtrel?*
“Is Iguazu a big or a small country?”
Is Iguazu micel rice oþþe lytel?
11. *Nēþer Pēter Smīþ nē Maria Brūn mihton bēon gecōrene.*
“Neither Peter Smith nor Mary Brown could be selected.”
Ne Petrus Smiþ ne Maria Brun mihton beon gecorene.
12. *Hīe nabbaþ nāne ġeþōht hwā hit wrāt.*
“They have no idea who wrote it.”
Hie nabbaþ nane ġeþoht hwa hit wrat.

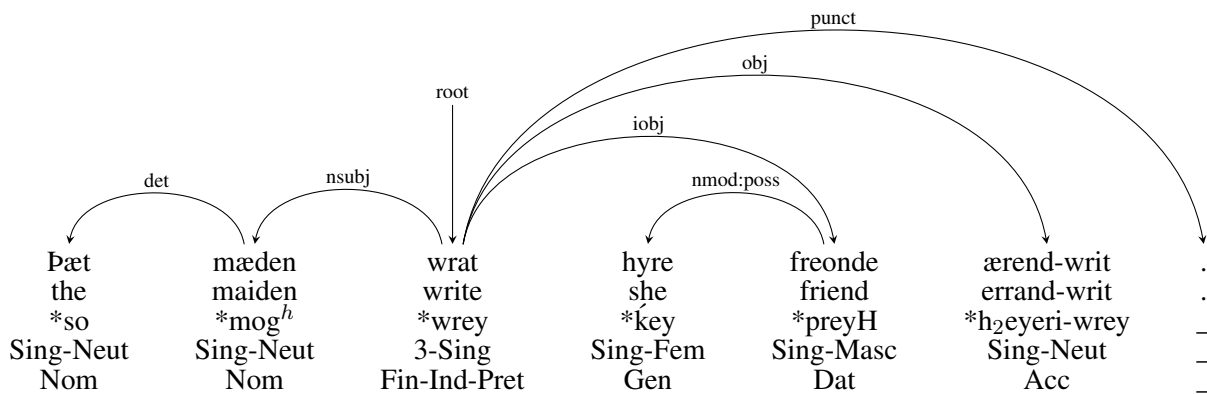


Figure 4: Dependency analysis of a ditransitive construction in OE: *The girl wrote a letter to her friend*. The glossing contains (from top to bottom) OE form, hyperlemma, root, and morphological feature annotations (person-number-gender and verbForm-mood-tense, as applicable).

13. *Hwæt locast þū on?*
 “What are you looking at?”
 Hwæt locast þu on?
14. *Hwonne þencst þū þæt þū mōst cuman?*
 “When do you think you can come?”
 Hwonne þencst þu þæt þu meaht cuman?
15. *Hē bohte cræt ac his brōþor hweōl-bearwe.*
 “He bought a car but his brother just a wheelbarrow.”
 He bohte cræt ac his broþor hweol-bearwe.
16. *Petrus and Maria beclypton heom be-tweonen and þa forlētone þa heoþe.*
 “Peter and Mary hugged each other and then left the room.”
 Petrus and Maria beclypton heom be-tweonen and þa forleton þa heoþe.
17. *Hēo scolde þæt hire hēafodgearwung dyde ac for sumre rēswan hēo nolde on þām dæge.*
 “She should have been doing her hair but for some reason she wouldn’t that day.”
 Hit gebyrede þt heo dyde hire her, ac for sume wisan nolde seo don swa.
18. *Ic ne mihte fylgan, forþām hē tō hrædlīce arn.*
 “I wasn’t able to keep up, because he ran too fast.”
 Ic ne mihte þurhwunian, forðon he to swiftlice rann.
19. *Þis gewrit is fram Pētere and hit wæs gebrōht gierstan dæge.*
 “This letter is from Peter and it was delivered yesterday.”
 Þis ærend-writ is fram Petere and hit wæs gebroht gierstan-dæg.
20. *Hēo weox on Paris, þære heofodburh Franclandes.*
 “She grew up in Paris, the capital of France.”
 Heo weox on Paris, þære heafod-burh Franc-landes.

B Sample OE sentence analysis

We provide a sample Old English sentence analysis, complete with its syntactic features, hyperlemma, Indo-European root, and morphological features in Figure 4.

C Old English specific morphological feature space

Our feature space follows that of Modern English, with the exception of case and gender, with 5 cases (nominative, accusative, genitive, dative, instrumental – the latter not attested in UD Cairo) and 3 genders (masculine, neuter, and feminine). Old English also exhibited dual number in pronouns, though this is not attested in our data. Features in tense, person, verb-form, and mood are equivalent with constraints set in Modern English UD corpora.

1. Case. Nom, Acc, Gen, Dat, Ins.
2. Gender. Masc, Fem, Neut.
3. Number. Sing, Dual, Plur.
4. Person. 1, 2, 3.
5. VerbForm. Fin, Inf, Part, Ger.
6. Mood. Ind, Imp, Sub.
7. Tense. Pres, Pret
8. Degree. Pos, Comp, Sup
9. Poss. Yes