*Proceedings of the Second Workshop on Language Technologies for Historical and Ancient Languages (LT4HALA 2022)*, pages 129–134
Language Resources and Evaluation Conference (LREC 2022), Marseille, 25 June 2022


# Linguistic Annotation of Neo-Latin Mathematical Texts: a Pilot-Study to Improve the Automatic Parsing of the *Archimedes Latinus*

**Margherita Fantoli[1], Miryam de Lhoneux[123]**
[1]KU Leuven, [2]University of Copenhagen, [3]Uppsala University
[1]Oude Markt 13, 3000 Leuven,[2] Universitetsparken 1. DK-2100 Copenhagen, [3]Thunbergsvägen 3H, 752 38 Uppsala
margherita.fantoli@kuleuven.be, ml@di.ku.dk

**Abstract**

This paper describes the process of syntactically parsing the Latin translation by Jacopo da San Cassiano of the Greek mathematical work *The Spirals* of Archimedes. The Universal Dependencies formalism is adopted. First, we introduce the historical and linguistic importance of Jacopo da San Cassiano's translation. Subsequently, we describe the deep Biaffine parser used for this pilot study. In particular, we motivate the choice of using the technique of treebank embeddings in light of the characteristics of mathematical texts. The paper then details the process of creation of training and test data, by highlighting the most compelling linguistic features of the text and the choices implemented in the current version of the treebank. Finally, the results of the parsing are discussed in comparison to a baseline and the most prominent errors are discussed. Overall, the paper shows the added value of creating specific training data, and of using targeted strategies (as treebank embeddings) to exploit existing annotated corpora while preserving the features of one specific text when performing syntactic parsing.

**Keywords:** Dependency parsing, Latin mathematical language, Universal Dependencies

## 1. Introduction

Jacopo da San Cassiano (1395-1494) translated most of the archimedean corpus from Ancient Greek to (Neo)-Latin around 1450 probably on the orders of Pope Nicholas V (d'Alessandro and Napolitani, 2012). The works of Archimedes, alongside Euclides' *Elements*, are considered a pillar of Hellenistic, and, in general, Greek mathematics (Heath, 1921). Jacopo's translation became a crucial medium for the rediscovery of Archimedes, and thus of Greek mathematics, among Humanists (Høyrup, 2019) and was used in the *editio princeps* of the Greek texts (Thomas Gechauff Venatorius, Basel 1544): the Latin translation was considered necessary to properly understand such a difficult work. Unlike modern mathematical texts, that rely heavily on symbolic notation, Ancient Greek mathematical texts are entirely written in plane natural language. Expressions that nowadays are rendered as 'AB:CD=DE:EF' were expressed as 'the line AB has to the line CD the same proportion that DE has to EF'. This creates an extremely peculiar variety of Ancient Greek (Acerbi, 2011; Acerbi, 2012; Netz, 2003), translated to Latin by Jacopo adopting the same style. The study of the linguistic features of Jacopo's translation hasn't been undertaken until now. Nonetheless, the creation of a treebank of this corpus is promising for different reasons:

- The text features a variety of Latin rarely targeted by linguistic studies and underrepresented in linguistic resources.
- The adaptation of the Latin language for translating the Greek of mathematics poses unique challenges.
- The availability of linguistically annotated Renaissance texts is still limited.

By creating the *Archimedes Latinus* treebank we aim at investigating the syntactic peculiarities of mathematical (Neo)-Latin. Given the success of the Universal Dependencies (UD) initiative (Nivre et al., 2016), and the large treebank availability (160), the Latin Archimedes treebank adopts the UD formalism. In addition, the recently created UDante treebank (Cecchini et al., 2020a) represents a milestone for Latin UD annotation. In fact, it is the first "native" UD treebank and its creation has generated the first (not-yet complete) language-specific guidelines for Latin[1]. Regularity is one of the most striking features of mathematical language, since a handful of terms and syntactic structures, indicating mathematical objects and relations, constitute the bulk of the text. Hence, we aim at verifying whether a syntactic parser, trained on a part of Jacopo's translation, can successfully parse the rest of the corpus, or at least reach results that significantly accelerate the post-correction for the treebank creation. This paper is structured as follows: section 2 describes the syntactic parser that we have finetuned for this study and introduces the concept of treebank embedding; section 3 describes the creation of training and test data; section 4 discusses the results of the parsing.

## 2. Parser

Treebank (or dataset) embeddings have been developed by (Stymne et al., 2018) on the ground of (de Lhoneux

[1]Cf. for example `https://universaldependencies.org/la/dep/obl-cmpr.html`.

et al., 2017), in order to tackle the problem of training a monolingual dependency parser using heterogenous treebanks. In fact, different UD treebanks for the same language might differ on specific aspects of the UD formalism: for instance, on the choice of the PoS for non clear-cut categories such as DET and PRON (see discussion below). Such inconsistencies might cause poor performances of parsers trained on multiple treebanks (Stymne et al., 2018). Treebank embeddings are used to prioritize, when parsing a new text, the conventions of one of the treebanks used for the training. To this goal, during training, a treebank embedding is concatenated to each word and thus one representation is learned for each of the treebank used. The representation is the same for every token of one treebank and differs from one treebank to another. When parsing a new text, a treebank identifier is given and the sentences get parsed following the 'style' of the chosen treebank. This method allows to take advantage of large training sets without overlooking treebank-specific features. The same method has been exploited to use treebanks of related languages during the training of a dependency parsing model (Smith et al., 2018). By applying this method to our case, we aim at verifying whether the parser picks up the specific, 'regular', features of the mathematical text while taking advantage of other existing treebanks. Hence, we train a multitask model to predict Parts-of-Speech (POS) and parse the text. We use the deep biaffine parser (Dozat and Manning, 2017) implementation of MaChAmp (van der Goot et al., 2021) with the use of dataset embeddings in the encoder introduced in (van der Goot and de Lhoneux, 2021). The parser uses mBERT[2] (Devlin et al., 2019) as an encoder, and a dataset embedding is concatenated to the embedding of each wordpiece before it is passed to the decoder.

# 3. Data Creation

## 3.1. Text extraction and tokenization

B. Sisana, a scholar specializing on Jacopo's work[3], has shared the critical edition of the whole corpus of Jacopo's archimedean translations, as contained in the the manuscript Nouv. Acq. Lat. 1538 (Paris, Bibliothèque nationale) identified as the autograph (d'Alessandro and Napolitani, 2012). The text is edited using MauroTeX, a specific mark-up language developed for the edition of mathematical texts[4]. Via a Python script, the text of the *Spirals* (on which this pilot study focusses) was extracted from the TeX file[5]. The tokenization and an automated PoS tagging have been performed using the Pie Latin LASLA+ model 0.0.6 (Manjavacas et al., 2019), fine-tuned on ca. 1,500,000 tokens taken from the LASLA Latin corpus[6] (Clérice, 2021).

## 3.2. Creating training data and test data

A training and test sets of sentences of *The Spirals* have been created[7]. They both qualify as Gold Standard since they have been manually annotated by a Latin philologist. The training set consists of the first 48 sentences of the book (1307 tokens), while the test set consist of 30 sentences taken from Propositions XIX and XX of the same book (913 tokens). The automatic PoS tagging of the Pie Latin LASLA+ model 0.0.6 has been corrected using Pyrrha, a language independent post correction app for PoS and lemmatization (Clérice et al., 2019). The PoS tagset used in Pyrrha has been converted to the Universal PoS tagset (Petrov et al., 2012) adopted in the UD initiative. At this point, the annotation of mathematical letters remains an open challenge. In fact, to indicate points, lines, circumferences and other mathematical objects, Ancient Greek authors use letters, in expression such as 'the line AB', often reduced to 'the AB'. In Greek manuscripts, a line is traced on the top of the letters. Jacopo follows the same convention, adding to dots around the string of letters, as visible in Figure 1. Recent studies have argued that these (groups of) letters have the purely linguistic anaphoric value of labels, since they allow to refer unambiguously to the same mathematical object across the same proof (Acerbi, 2020). This goes against the theory according to which their primary goal is to identify specific points in a diagram. It is thus not obvious whether the string of letters can be considered as a single token and what PoS should be assigned to it. Based on the graphic evidence of the manuscripts, we decided to keep the forms as a single token, instead of considering each letter as separate token corresponding to a point. The PoS has been left undetermined at this stage (X), but will be assigned either to NOUN [8] and SYM[9] following additional discussions with UD experts. In fact, UD guidelines assign the PoS SYM to

Figure 1: Snippet of the Nouv. Acq. Lat. 1538 'in tempore gk'

[2]We also experimented with latin-bert (Bamman and Burns, 2020) and used only the Latin treebanks for the training, but this led to slightly worse results so we decided to stick with mBERT and the multilingual cluster.

[3]PhD student at the Università di Roma Tre.

[4]See `https://people.dm.unipi.it/maurolic/mtex/mtexen.htm`

[5]The extracted text is available at https://github.com/mfantoli/Archimedes$_{Latinus/blob/main/texts/spirals.txt}$

[6]`http://web.philo.ulg.ac.be/lasla/`

[7]The annotated data are available here `https://github.com/mfantoli/Archimedes_Latinus`.

[8]`https://universaldependencies.org/u/pos/NOUN.html`

[9]`https://universaldependencies.org/u/pos/SYM.html`

'word-like entities that differ from ordinary words by form, function, or both' and indicate that mathematical operators are SYM. Given the specific layout of these strings in the manuscripts, the choice SYM seems appealing. On the other hand, in the UDante treebank, in expressions such as *linea recta ad A*, *A* is tagged as NOUN, and indeed single letters are considered nouns in traditional lexicographical resources. The decision does not impact the syntactic parsing, but will be defined before the publication of the final version of the data. In addition, to mirror the layout of the manuscript, the mathematical letters are kept between square brackets (*linea* [AB]). In order to create gold data for the syntactic parsing, the Biaffine parser discussed in section 2 has been trained on a cluster of ancient languages described in (Smith et al., 2018): the UD version of Latin Index Thomisticus Treebank (Passarotti, 2019), of the Perseus Ancient Greek and Latin Dependency Treebanks (Bamman and Crane, 2011), of the PROIEL Old Church Slavonic, Gothic, Latin and Ancient Greek Treebanks (Eckhoff et al., 2018), of the Late Latin Charter Treebank (Cecchini et al., 2020b; Korkiakangas, 2021), and the UDante treebank. (Cecchini et al., 2020a). The sentences of the training and test sets have been parsed using the UDante embedding as treebank 'model', and manually corrected using UD Annotatrix (Tyers et al., 2017). The annotation follows the UD-style available guidelines for Latin[10] and takes into account the choices implemented in the UDante treebank (Cecchini et al., 2020a). Nonetheless, given the still limited availability of language-specific UD guidelines for Latin, and the non-literary and non-classical linguistic features of *The Spirals*, some choices have been implemented following discussions with UD experts[11]. In the UD formalism, syntactic annotation consists in identifying typed dependency relations between the words forming the sentence. Each word of a sentence - except the root- depends on one another word (head). The relation (EDGE) between a word and its head is typed based on UD dependency relations (DEPRELs[12]). The root is the head of the sentence. The DEPREL between the term indicating the mathematical object and the label (*linea AB*) has been indicated as 'flat'[13], since it is comparable to expressions such as 'President Obama'[14]. Conventionally, we indicated as head the NOUN ('*linea*'), which generally coincides, with few exceptions, with the first word of the compound[15]. However, flat relations imply that the choice of the head is arbitrary since the two words do not hold a head-modifier relation. The length of sentence represents a second challenge: the digitized text is, in fact, the direct transcription of a Renaissance manuscript. Manuscripts tend to record only minimal punctuation (Parkes, 1992), and the transcription sticks to the original layout. This entails extremely long sentences, whose clauses are rarely separated by commas: in the test set, the median length is of 21.5 words with a maximum of 104, resulting in syntactic trees with a median depth of 5.5 layers. This will be addressed for the final version of the treebank, by adding punctuation as modern editors regularly do. The lack of punctuation entails additional difficulties in analyzing the role of Latin particles in the sentence. Words such as *enim* ('namely', 'indeed'), *autem* ('however', 'on the other hand'), *rursus* ('on the contrary'), can both linking clauses belonging to the same sentence or link one sentence with the preceding one, structuring the discourse (Kroon, 1995). In very long sentences composed of a number of juxtaposed clauses, it is challenging to establish whether the tuple (verb, particle) should receive the DEPREL 'cc'[16], in case the particle functions as coordinating conjunction, or 'discourse'[17], when functioning as discourse marker.

## 4. Training and evaluation

Once the training data and test data have been created, the Biaffine parser is again trained on the cluster of ancient languages described above, with the addition of the newly created mathematical training data. For the prediction of the PoS tags, the mathematical training data, the treebanks, and the UD versions of the LASLA corpus[18] are used as training data. The model is trained for 80 epochs[19]. The Unlabeled Attachment Score (UAS) measures the correctness of the syntactic structure (EDGES), whereas the Labeled Attachment Score (LAS) includes also the evaluation of the label attached to the dependency (DEPRELs)[20] (Buchholz and Marsi, 2006). Table 1 reports the LAS and UAS scores computed on the test data processed with this

[10] https://universaldependencies.org/guidelines.html.

[11] In particular, Flavio M. Cecchini, CIRCSE, Università Cattolica di Milano.

[12] https://universaldependencies.org/u/dep/

[13] https://universaldependencies.org/u/dep/flat.html

[14] https://universaldependencies.org/u/dep/flat.html

[15] The UD guidelines indicate that the first of the two words linked by the 'flat' DEPREL should be used as head. The few exceptions will be corrected in the final version of the Treebank.

[16] https://universaldependencies.org/docs/en/dep/cc.html.

[17] https://universaldependencies.org/docs/en/dep/discourse.html

[18] The LASLA corpus is a morphosyntactically manually annotated corpus of Latin classical texts (Denooz, 1978), see http://web.philo.ulg.ac.be/lasla/presentation-du-laboratoire/. The LASLA data were converted to UD by Flavio M. Cecchini, CIRCSE, Università Cattolica di Milano. A sample of these data has been used in the frame of Evalatin2022, https://circse.github.io/LT4HALA/2022/EvaLatin.

[19] The model will be shared soon.

[20] For this pilot study, the LAS was computed on first-level relations only, without considering any subrelation (e. g., the DEPRELS obl, obl:cmpr and obl:arg all count as obl).

parser ('Archimedes') and with UDPipe using the UD v. 2.6 for Latin (baseline)[21].

| Model | PoS | UAS | LAS |
|---|---|---|---|
| Biaffine Archimedes | 91.25 | 72.43 | 59.85 |
| IT-TB | NA | 68.60 | 55.03 |
| Perseus | NA | 68.16 | 50.44 |

Table 1: UPOS, UAS and LAS score of different parsers

The results show a significant gain with respect to the UDPipe IT-TB model, 3.83 UAS points and 4.82 LAS points. The results can still be improved significantly: at the moment such procedure can only be effectively use as a first step to accelerate the following manual annotation. Nevertheless, it seems that a multi-task learning setup is well suited to using multiple sources of data to facilitate the annotation of a new dataset. Additionally, dataset embeddings facilitate annotating new data in the style of a specific treebank.
In order to measure the impact of the addition of mathematical texts and the use of the 'mathematical embedding' on the performance of the Biaffine parser, we also evaluated the performance of the Biaffine parser with the training described in 3.2 on a very brief portion of text (ca 400 tokens, propositions VII-VIII). The results on UAS and LAS (resp.70.56 and 58.63) outperform UDPipe, but are lower than those obtained in the final stage. However, we should mention that this test-set might not be representative, since some complex sentences had to be removed due to editorial issues. The annotation of PoS scores quite high (95.6): the result, higher than with the addition of mathematical texts, can be explained by the absence, in this portion of text, of the term *spiralis* ('spiral'), which is the main source of errors for the final test-set (see below).

## 5. Error analysis

To complement the scores, we performed an analysis of the errors on the POS, dependencies and labels. The confusion matrix of the PoS is shown in Figure 2. The most frequent error is due to the mislabeling of *spiralis* ('spiral') as NOUN in the expression *linea spiralis* ('spiral line'), where it is an ADJ. The second most frequent source of errors is the confusion between DET[22] and PRON[23], which is mostly due to linguistic ambiguity, given that the same words, such as *ille* ('that','that

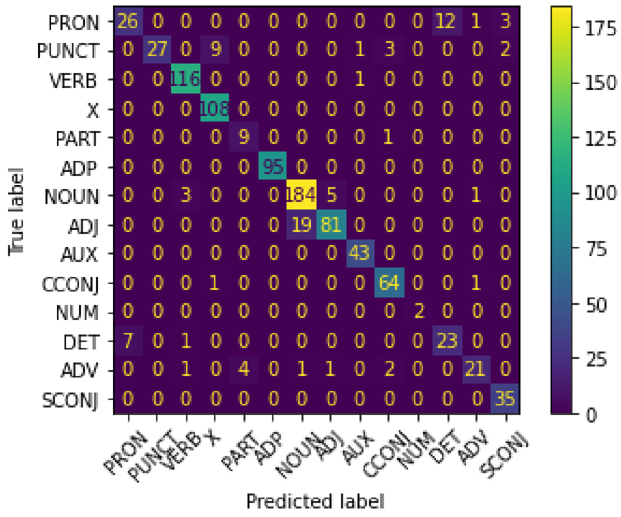


Figure 2: Confusion matrix for the PoS prediction

person') or *iste* ('this','this person'), can be used with both functions[24]. Out of the 253 cases of wrong head assignment, 57 concern mathematical labels, which is a highly specific feature of our text. In the subset of corrected predicted dependency, the most common error in DEPRELs assignment (15 times out of 127 errors) is 'nmod'[25] instead of 'flat', always between a mathematical term and its label. As it appears, most of the errors generate from the mathematical content of the text.

## 6. Conclusion

The linguistic annotation of non-classical, non-literary varieties of Latin poses major challenges, both because of the difficulty of adapting existing guidelines to these texts[26] and because of the lack of well-suited annotated data and tools to automate the process. In this pilot study we have shown the added value of creating specific training data, and of using targeted strategies (as treebank embeddings) to jointly exploit existing annotated corpora without losing the features of one specific text. Such strategy beats baseline results, and appears promising for the future. As next steps, the performance of the parser will be improved by assigning PoS to mathematical labels in the training data and by increasing the amount and variety of training data from Jacopo's translation of different work of Archimedes.
In a second stage, we will manually correct the output of the parser to provide a treebank of Jacopo's translation of at least one complete work of Archimedes. Finally, the completion of such project will result in the

[21]The PROIEL score is not recorded because the model splits long sentences at weak punctuation marks, and the PoS score is not reported for IT-TB and Perseus because of the 'X' PoS assigned in the gold data to the mathematical labels.
[22]see UD guidelines `https://universaldependencies.org/u/pos/DET.html`.
[23]see UD guidelines `https://universaldependencies.org/u/pos/PRON.html`.
[24]see UD guidelines `https://universaldependencies.org/u/pos/DET.html`.
[25]`https://universaldependencies.org/en/dep/nmod.html`
[26]see, for instance, (Korkiakangas and Passarotti, 2012) and (Grotto et al., 2021)

contribution to the guidelines for the UD-style annotation of Latin, in particular scientific Latin.

## 7. Acknowledgments

We would like to thank Flavio Cecchini for his advice on many of the annotation questions raised during this project. We would also like to thank the reviewer for their precious comments and suggestions.

## 8. Bibliographical References

Acerbi, F. (2011). The language of the "givens": its forms and its use as a deductive tool in greek mathematics. *Archive for History of Exact Sciences*, 65(2):119–153, feb.

Acerbi, F. (2012). I codici stilisticidella matematica greca:dimostrazioni, procedure, algoritmi. *Quaderni Urbinati di Cultura Classica*, (101.2):167–216.

Acerbi, F. (2020). Mathematical generality, letter-labels, and all that. *Phronesis*, 65:27–75.

Bamman, D. and Burns, P. J. (2020). Latin bert: A contextual language model for classical philology. *ArXiv*, abs/2009.10053.

Bamman, D. and Crane, G., (2011). *Language Technology for Cultural Heritage*, chapter The Ancient Greek and Latin Dependency Treebanks, pages 79–98. SpingerLink.

Buchholz, S. and Marsi, E. (2006). CoNLL-X shared task on multilingual dependency parsing. In *Proceedings of the Tenth Conference on Computational Natural Language Learning (CoNLL-X)*, pages 149–164, New York City, June. Association for Computational Linguistics.

Cecchini, F. M., Sprugnoli, R., Moretti, G., and Passarotti, M. (2020a). UDante: First Steps Towards the Universal Dependencies Treebank of Dante's Latin Works. In *Seventh Italian Conference on Computational Linguistics*, pages 1–7, Bologna. CEUR-WS.org.

Cecchini, F. M., Korkiakangas, T., and Passarotti, M. (2020b). A new Latin treebank for Universal Dependencies: Charters between Ancient Latin and Romance languages. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 933–942, Marseille, France, May. European Language Resources Association.

Clérice, T. (2021). Latin Lasla Model, Apr. DOI: `10.5281/zenodo.4661034`.

Clérice, T., Pilla, J., Jean-Baptiste-Camps, and architexte. (2019). hipster-philology/pyrrha: 2.1.0, November.

Paolo d'Alessandro et al., editors. (2012). *Archimede Latino: Iacopo da San Cassiano e il Corpus archimedeo alla metà del Quattrocento con edizione della Circuli dimensio e della Quadratura parabolae*. Les Belles Lettres, Paris.

de Lhoneux, M., Shao, Y., Basirat, A., Kiperwasser, E., Stymne, S., Goldberg, Y., and Nivre, J. (2017). From raw text to Universal Dependencies - look, no tags! In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 207–217, Vancouver, Canada, August. Association for Computational Linguistics.

Denooz, J. (1978). L'ordinateur et le latin. techniques et méthodes. *Revue de l'organisation internationale pour l'étude des langues anciennes par ordinateur*, 4:1–36.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (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, June. Association for Computational Linguistics.

Dozat, T. and Manning, C. (2017). Deep Biaffine Attention for Neural Dependency Parsing. In *Proceedings of the 5th International Conference on Learning Representations.*

Eckhoff, H. M., Bech, K., Bouma, G., Eide, K., Haug, D., Haugen, O. E., and Jøhndal, M. (2018). The PROIEL treebank family: a standard for early attestations of Indo-European languages. *Language Resources and Evaluation*, 52(1):29–65.

Grotto, F., Sprugnoli, R., Fantoli, M., Simi, M., Cecchini, F. M., and Passarotti, M. C. (2021). The annotation of liber abbaci, a domain-specific latin resource. In *Proceedings of the Eighth Italian Conference on Computational Linguistics (CLiC-it 2021). Milan, Italy, January 26-28, 2022, (MILANO – ITA, 26-28 January 2022)*, pages 1–8.

Thomas Heath, editor. (1921). *A History of Greek Mathematics*. Clarendon Press, Oxford.

Høyrup, J., (2019). *Archimedes: Reception in the Renaissance*, pages 1–7. Springer International Publishing, Cham.

Korkiakangas, T. and Passarotti, M. (2012). Challenges in annotating medieval latin charters. *Journal for Language Technology and Computational Linguistics*, 26(2):103–114, February.

Korkiakangas, T. (2021). Late latin charter treebank: contents and annotation. *Corpora*, 16(2):191–203.

Kroon, C. (1995). *Discourse particles in Latin : a study of nam, enim, autem, vero and at*. Gieben, Amsterdam.

Manjavacas, E., Kádár, Á., and Kestemont, M. (2019). Improving lemmatization of non-standard languages with joint learning. 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 1493–1503, Minneapolis, Minnesota, June. Association for Computational Linguistics.

Netz, R. (2003). *The Shaping of Deduction in Greek*

*Mathematics : A Study in Cognitive History*. Cambridge University Press.

Nivre, J., De Marneffe, M.-C., Ginter, F., Goldberg, Y., Hajic, J., Manning, C. D., McDonald, R., Petrov, S., Pyysalo, S., Silveira, N., et al. (2016). Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666.

M.B. Parkes, editor. (1992). *Pause and Effect: An Introduction to the History of Punctuation in the West*. Routledge, Berkeley-Los Angeles.

Passarotti, M. (2019). The Project of the Index Thomisticus Treebank. In Monica Berti, editor, *Digital Classical Philology. Ancient Greek and Latin in the Digital Revolution*, pages 299–319. De Gruyter, Berlin.

Petrov, S., Das, D., and McDonald, R. (2012). A universal part-of-speech tagset. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2089–2096, Istanbul, Turkey, May. European Language Resources Association (ELRA).

Smith, A., Bohnet, B., de Lhoneux, M., Nivre, J., Shao, Y., and Stymne, S. (2018). 82 treebanks, 34 models: Universal Dependency parsing with multi-treebank models. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 113–123, Brussels, Belgium, October. Association for Computational Linguistics.

Stymne, S., de Lhoneux, M., Smith, A., and Nivre, J. (2018). Parser training with heterogeneous treebanks. pages 619–625, 01.

Tyers, F. M., Sheyanova, M., and Washington, J. N. (2017). UD annotatrix: An annotation tool for Universal Dependencies. In *Proceedings of the 16th International Workshop on Treebanks and Linguistic Theories*, pages 10–17, Prague, Czech Republic.

van der Goot, R. and de Lhoneux, M. (2021). Parsing with pretrained language models, multiple datasets, and dataset embeddings. In *Proceedings of the 20th International Workshop on Treebanks and Linguistic Theories (TLT, SyntaxFest 2021)*, pages 96–104, Sofia, Bulgaria. Association for Computational Linguistics.

van der Goot, R., Üstün, A., Ramponi, A., Sharaf, I., and Plank, B. (2021). Massive choice, ample tasks (MaChAmp): A toolkit for multi-task learning in NLP. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 176–197, Online. Association for Computational Linguistics.