Symmetric Dependency Structure of Coordination: Crosslinguistic Arguments from Dependency Length Minimization

Adam Przepiórkowski
ICS Polish Academy of Sciences
and University of Warsaw

Magdalena Borysiak
University of Warsaw

Adam Okrasiński
University of Warsaw

Bartosz Pobożniak
University of Warsaw

Wojciech Stempniak
University of Warsaw

Kamil Tomaszek
University of Warsaw

Adam Głowacki
University of Warsaw

Abstract

The aim of this paper is to replicate and extend recent treebank-based considerations regarding the syntactic structure of coordination. Overall, we confirm the previous results that, given the principle of Dependency Length Minimization, corpus data suggest that the structure of coordination is symmetric. While previous work was based on 2 English datasets, we extend the investigation to 3 more English datasets, 3 Polish datasets, and UD corpora for a number of diverse languages. The results confirm the symmetric structure of coordination, but they also make it possible to question some of the previous findings regarding the exact symmetric structure of coordination.

1 Introduction

There is no agreement in theoretical linguistics about the syntactic structure of coordination. Within dependency approaches alone, 4 basic structures have been proposed with a number of variants (see Popel et al. 2013 and Przepiórkowski and Woźniak 2023; the latter henceforth abbreviated to PW23), as schematically presented in (1)–(4):

(1) **Multi-headed/London**:

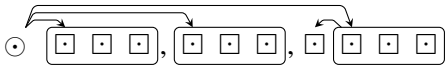

(2) **Conjunction-headed/Prague**:

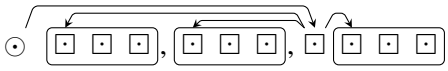

(3) **Chain/Moscow**:

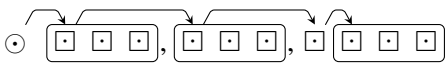

(4) **Bouquet/Stanford**:

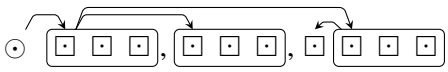

In these schemata, ⊙ marks the governor (e.g., *saw* in (5)), ⊡ marks tokens within coordination, with tokens belonging to the same conjunct grouped; the single ungrouped ⊡ is the conjunction (e.g., *and*).

(5) Maggie *saw* [[a brown dog], [a grey cat], and [a green tree]].

Moreover, these schemata follow syntactic theory in assuming that heads of conjuncts are typically near the beginning of these conjuncts in English, given that it is a head-initial language; e.g., the DP conjuncts in (5) are headed by the determiners.[1]

Prague and London approaches (1)–(2) are symmetric in the sense that all conjuncts bear the same relation to the governor of the coordinate structure: in (1) they are direct dependents of the governor, while in (2) they are all direct dependents of the conjunction. By contrast, in the asymmetric (3)–(4), only the first conjunct is a direct dependent of the governor of the coordinate structure, with the other conjuncts being direct (in (4)) or possibly indirect (in (3)) dependents of the first conjunct.

PW23 give a novel corpus-based argument for a symmetric structure of coordination. The argument assumes the principle of Dependency Length Minimization (DLM) – a robustly demonstrated tendency for natural languages to strive for maximally local dependencies.[2] As argued in Hawkins 1994 and Futrell et al. 2020, this tendency operates both at the level of use and at the level of grammar.

At the level of use, when both orders of two dependents are grammatical, the longer one of these dependents gets, the stronger the pressure for the other dependent to occur closer to the governor. For example (cf. Przepiórkowski and Woźniak 2023), consider an intransitive verb, e.g., *sing*, and its two PP dependents: a durative PP_{for}, e.g., *for two hours*, and a locative PP_{in}, e.g., the short *in that club* or

[1] This should be contrasted with Universal Dependencies (UD; https://universaldependencies.org/; Nivre et al. 2016, de Marneffe et al. 2021, Zeman et al. 2024), where the nouns – i.e., typically conjunct-final tokens – are assumed to be heads; see Osborne and Gerdes 2019 for discussion.

[2] See, e.g., Behaghel 1909, Hawkins 1994, Gibson 1998, Ferrer-i-Cancho 2004, Gildea and Temperley 2007, 2010, Liu 2008, Liu et al. 2017, Futrell and Levy 2017, Temperley and Gildea 2018, and many others.

Proceedings of the 22nd Workshop on Treebanks and Linguistic Theories (TLT 2024), pages 11 - 22

the much longer *in the most famous American jazz club*. Then, if the likelihood of the order [V PP_{for} PP_{in}] (as opposed to [V PP_{in} PP_{for}]) is p_1 for the shorter PP_{in} (i.e., for *sing [for two hours] [in that club]*), then it will be $p_2 > p_1$ for the longer PP_{in} (i.e., for *sing [for two hours] [in the most famous American jazz club]*).

Such at-use pressures may become conventionalized, i.e., they may become at-grammar tendencies. For example (cf. PW23 again), given that NPs are on average shorter than PPs (which consist at least of a preposition and an NP), DLM will be more often satisfied by [V NP PP] than by the [V PP NP] order. Hawkins (1994: 90) argues that this tendency became conventionalized in English into a general preference for the former order, active even when the lengths of the NP and the PP are equal, i.e., when there is no at-use DLM gain. For example, despite the similar lengths of the two dependents, *I sold [my mother's ring] [for five dollars]* is preferred to *I sold [for five dollars] [my mother's ring]*. On the other hand, this at-grammar pressure may be overridden by the at-use pressure when length differences are large: the [V PP NP] order becomes more natural again in *I sold [for five dollars] [my mother's silver engagement ring that she got from my father]*, despite the violation of the at-grammar preference for [V NP PP].

Now, the general idea of PW23's argument is to compare the predictions of each of the four proposed structures of coordination to what is observed in corpora. For example, consider binary coordinations in the asymmetric Stanford approach.

(6) **Bouquet/Stanford**:

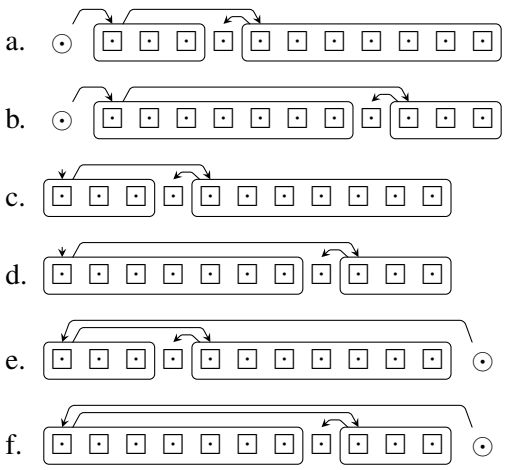

(6a–b) illustrates coordination with the governor on the left (as in (5)), (6c–d) – coordinations with no governor (e.g., coordination of sentences), (6e–f) – those with the governor on the right (e.g., *Bart and Lisa laughed*). Each pair compares two orders of conjuncts: in the first the first conjunct is shorter, in the second – the second is shorter.

If DLM operates in coordinate structures only at the level of use, then the following tendencies are predicted. First, as seen in (6a–b), when the governor is on the left, there is a pressure for the first conjunct to be shorter: the total sum of dependency lengths is smaller in (6a) than in (6b).[3] So, there is an at-use pressure for the shorter conjunct to occur as the first conjunct when the governor is on the left. Moreover, the difference between the aggregate dependency lengths in (6a–b) is equal to the difference of lengths of the two conjuncts. Hence, this pressure for the shorter conjunct to be first is greater when the conjunct length differences are greater. These considerations translate into a clear prediction: when the difference between the lengths of conjuncts is greater, the proportions of coordinations with the shorter first conjunct should be greater. Formally, let $p_L(n)$ be the proportion of those binary coordinations with a governor on the left with the absolute length difference between the two conjuncts being $n > 0$ in which the first conjunct is shorter. The prediction of the Stanford approach is that $p_L(n)$ should be a monotonically increasing function of n.

It is easy to see that exactly the same prediction is made when there is no governor (see (6c–d)) and when the governor is on the right (see (6e–f)): in all three cases, when the first conjunct is shorter, the aggregate dependency length is smaller. Moreover, in all three cases the difference between the sum of lengths is the same and equal to the length difference between the two conjuncts. That is, $p_L(n)$, $p_-(n)$ (the proportion function when there is no governor), and $p_R(n)$ (the proportion function when the governor is on the right) should all be equally monotonically increasing.

In order to verify such predictions, PW23 examined the distribution of binary coordinations in $PTB_{\&}$ (Ficler and Goldberg 2016), a version of Penn Treebank (PTB; Marcus et al. 1993) which improves on PTB by offering explicit and relatively consistent information about coordinations. Out of 21,825 binary coordinations they extracted from around 49.2K sentences in $PTB_{\&}$, 13,106 had gov-

[3] Dependencies *within* conjuncts are not shown here, as they do not depend on the order of conjuncts, i.e., they do not matter for the comparison of aggregate dependency lengths. Also, unlike some of the previous work reported below, we only consider lengths measured in words here – not syllables or characters.

ernor on the left, 4,000 had no governor, and 4,719 had governor on the right. For each subpopulation, they fitted a monofactorial logistic regression model to estimate $p_L(n)$, $p_-(n)$, and $p_R(n)$. The result was that $p_L(n)$ and $p_-(n)$ were monotonically increasing, as predicted by the Stanford approach, but $p_R(n)$ was more or less constant, with confidence bands compatible with the true $p_R(n)$ being either decreasing or increasing. However, they also performed a multifactorial binary logistic regression analysis, which showed that the slope is statistically significantly flatter when the governor is on the right than when it is on the left or missing. This is not predicted by the Stanford approach, where all three slopes should be the same.

As binary coordinations have almost exactly the same dependency relations on the two asymmetric – Stanford and Moscow – approaches, the above observations and conclusions also hold for the Moscow approach, which we will not consider further. However, the predictions of the two symmetric approaches are more interesting. The relevant schemata are presented in (7)–(8):

(7) **Conjunction-headed/Prague**:

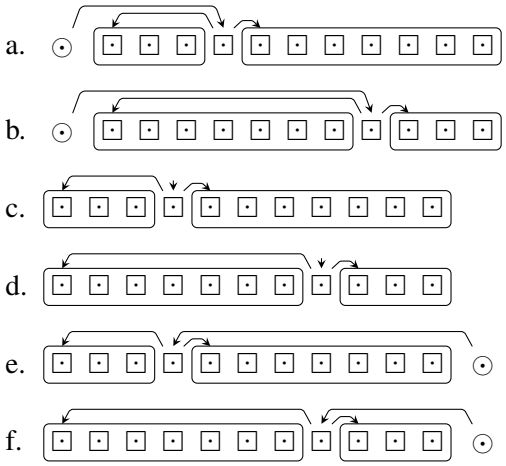

(8) **Multi-headed/London**:

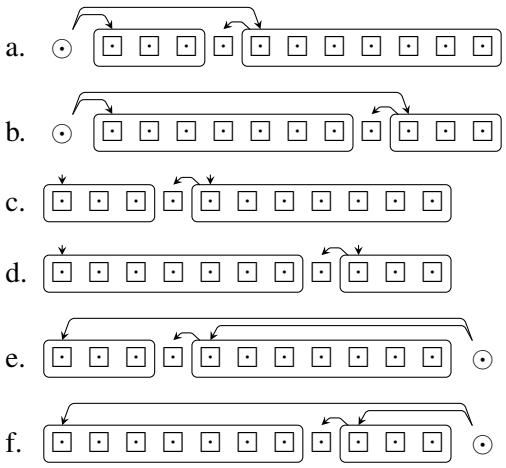

PW23 note that the Prague approach is directly compatible with their corpus analyses: shorter first conjuncts minimize aggregate dependency length when the governor is on the left (see (7a–b)) or absent (see (7c–d)), but not when it is on the right (see (7e–f)). In the latter case, the aggregate dependency lengths are the same in (7e–f). This directly corresponds to the observed monotonically increasing $p_L(n)$ and $p_-(n)$, and constant $p_R(n)$.[4]

Finally, in the case of the London approach, the prediction is that $p_L(n)$ is increasing (cf. (8a–b)), $p_-(n)$ is constant (cf. (8c–d)), and $p_R(n)$ is decreasing (cf. (8e–f)). This is not directly compatible with PW23's corpus-based models when DLM is only considered at use. However, PW23 also consider an at-grammar DLM effect, related to the well-known fact (which they confirm on the basis of PTB$_{\&}$) that most of the coordinate structures have their governor on the left in English. As – on any approach to coordination – the shorter first conjunct minimizes the aggregate dependency length in such situations, this means that in most cases it pays to have the first conjunct shorter and that this tendency could have plausibly been conventionalized to the at-grammar pressure for shorter first conjuncts in general.

The existence of such a hypothetical at-grammar tendency does not change anything in the case of asymmetric approaches (they still predict that all three $p_*(n)$ functions should be equally monotonically increasing), but it makes a difference in the case of the London approach. If such an at-grammar tendency is present, then $p_L(n)$ is still predicted to be monotonically increasing, but now also $p_-(n)$ is predicted to be monotonically increasing, by virtue of the at-grammar pressure alone. Moreover, the at-use pressure for the shorter *second* conjunct observed in (8e–f) is counterbalanced by the hypothetical at-grammar pressure for the shorter *first* conjunct, resulting in the roughly constant $p_R(n)$ observed in PTB$_{\&}$.

One of the limitations of PW23 is the relative scarcity of data: the number of coordinations with the governor on the right was not sufficient to train

[4] Note that the Prague approach predicts stronger pressure when the governor is on the left (see (7a–b); the difference between the two orders is *twice* that of the conjunct length difference) than when it is absent (see (7c–d); the difference is that of the conjunct length difference). PW23's multifactorial analysis confirms the corresponding difference of slopes between $p_L(n)$ and $p_-(n)$ when length is measured in characters or syllables, but it detects no statistically significant difference when it is measured in words.

a logistic regression model that would give a statistically significant answer concerning the monotonicity of $p_R(n)$. In an attempt to remove this limitation, Przepiórkowski et al. 2024 (henceforth, PBG24) replicate PW23's study on the basis of the Corpus of Contemporary American English (COCA; Davies 2008–2023) automatically parsed with Stanza (Qi et al. 2020) to the UD format. Unlike PW23, they considered the first and last conjunct in *all* coordinations, noting that over 86% of them were binary and that restriction to binary coordinations does not affect the results. From a subset of COCA containing almost 21.8M sentences, they extracted over 11.5M coordinations and fitted those with considerable length differences between conjuncts (at least 4 words) into a logistic regression model. As in PW23, the estimated $p_L(n)$ and $p_-(n)$ were monotonically increasing – $p_L(n)$ more so than $p_-(n)$ – but this time $p_R(n)$ was monotonically *decreasing* (statistically significantly with $p \ll 0.001$). This is clearly incompatible with asymmetric theories, on which all should be similarly increasing, not fully compatible with symmetric Prague approach, on which $p_R(n)$ should be constant if DLM only operates at use or increasing if it also operates at grammar, but fully compatible with the London approach, on the assumption that the at-grammar tendency is strong enough to make $p_-(n)$ – constant at the level of use – increasing, but not strong enough to make $p_R(n)$ – decreasing at use – constant or increasing. PBG24 conclude that their study makes it possible to sharpen the results of PW23, as it not only provides evidence for symmetric approaches to coordination in general, but for a particular such approach (London).

However, PBG24 note a major limitation of their approach that was absent in PW23, namely, the low quality of their automatically parsed data. For each governor position (left, absent, right) and each conjunct length difference (from 1 to 20 words), they sampled 15 coordinations from the 11.5M coordinations automatically extracted from COCA, resulting in 900 coordinations altogether, and checked whether they were extracted correctly, i.e., had the right information about governor position and identified the two conjuncts correctly, as only this information matters for the statistical model. They found that only slightly over 50% of coordinations were extracted correctly in this sense. While there are no reasons to think that the distribution of errors significantly influenced their results, such a problem cannot be *a priori* excluded, leading them to the conclusion that "further replication studies, also based on languages other than English, are needed to make these results even more robust" (Przepiórkowski et al. 2024: 1029).

2 New Studies

In order to validate the results of PW23 and PBG24, we performed a number of similar studies on different datasets: 2 for English, 2 for Polish (another head-initial language), and further studies based on UD corpora of a number of languages, including English and Polish. Because of relatively small sizes of the datasets for languages other than English, some of the results by themselves are not statistically significant, but taken together they largely confirm the conclusions of previous studies.[5]

2.1 English

Two English studies follow PBG24: they are based on automatically-parsed COCA, i.e., on a low quality but large resource. The difference with respect to PBG24's study is that in the current studies COCA was not parsed to the UD format.

As is well known (see, e.g., Przepiórkowski and Patejuk 2019), the representation of coordinate structures in the basic UD standard is not optimal: certain structures cannot be represented unambiguously. For example, there is just one UD representation of the sequence *lazy cats and dogs*, whether *lazy* modifies *cats* alone, or whether it modifies the whole coordinate structure (so that *dogs* are also *lazy*).[6] This is a problem, as it is not clear whether the two conjuncts in *lazy cats and dogs* are of same lengths (this is the case if *lazy* modifies the whole coordination) or whether the first conjunct is longer (if *lazy* modifies *cats* alone), and this information is crucial for the argument at hand. While PBG24 implemented various heuristics for disambiguating such representations, they are imperfect, so this ambiguity problem contributes to the low quality of input data in their study.

[5] All statistics and visualizations were performed using R (R Core Team 2024), with the stastical significance of slope differences estimated using the `emtrends` commands from the `emmeans` package (Lenth 2024).

[6] This difference is easy to represent in some other approaches, including the Prague approach and the enhanced version of UD (Schuster and Manning 2016). Unfortunately, the main dependency parsers currently in use only provide the basic UD structures.

2.1.1 COCA parsed with Stanza/SUD

In order to alleviate this problem, two different representations were used here. In the first study, the Surface-syntactic Universal Dependencies (SUD; Gerdes et al. 2018, 2021a) format was used, which makes it possible to represent information about shared dependents explicitly.[7] We trained Stanza on a treebank consisting of SUD versions of three English UD corpora: EWT (Silveira et al. 2014), GUM (Zeldes 2017), and (the English part of) ParTUT (Sanguinetti and Bosco 2014), all downloaded from https://surfacesyntacticud.github.io/data/.

In order to assess the quality of coordinations extracted using this SUD-trained parser, we also trained Stanza on the original UD versions of the same corpora, and compared the 1526 coordinations extracted from the testing parts of these corpora by the two trained parsers. The two parsers agreed on 1075 coordinations. In the case of the remaining 451, SUD-based procedure correctly identified 260 (57.6%) coordinations, and UD-based procedure – 252 (55.9%) coordinations. That is, the coordination extraction process based on the SUD-trained parser turned out to be only slightly better than that based on the UD-trained parser. (This difference was not statistically significant, according to McNemar's test.) Hence, while we did not evaluate the quality of extracted coordinations using the same procedure as PBG24, we do not expect coordinations based on SUD-trained Stanza to be of significantly better quality than those based on UD-pre-trained Stanza in PBG24.

Despite this only marginal improvement, the SUD-trained Stanza was used to replicate PBG24's study. The whole COCA was parsed and, as a result, 14,341,063 coordinations were identified, including 12,476,392 binary coordinations. Three logistic regression models were trained, as before, with results presented in the left column of Figure 1. The relations between the three slopes are as expected: the slope is most positive in the case of $p_L(n)$ (top graph) and least positive in the case of $p_R(n)$ (bottom graph). The fact that all slope differences are highly statistically significant ($p \ll 0.001$) is consistent with symmetric approaches, but not with asymmetric approaches.

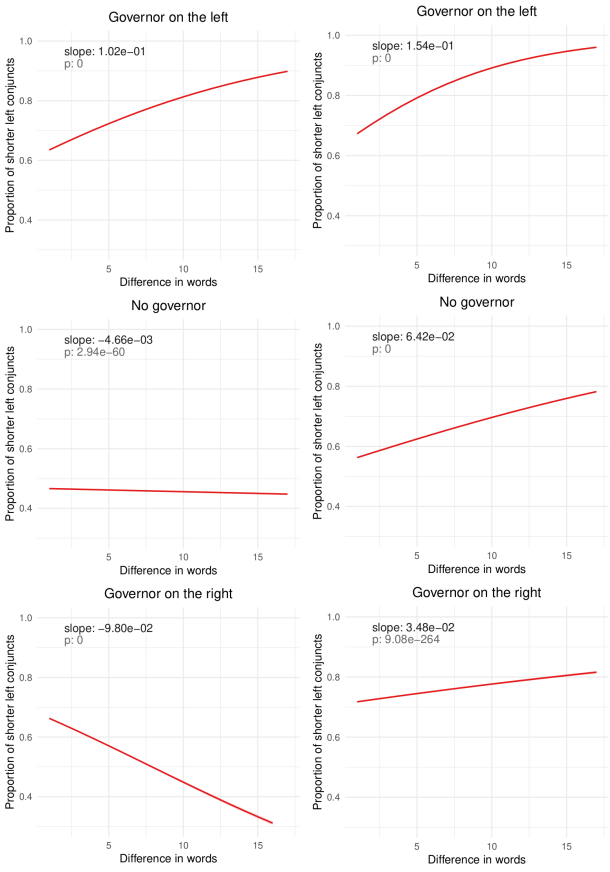

Figure 1: Logistic regression models of COCA coordinations extracted with Stanza trained on SUD (left column) and UD (right column) corpora

Moreover, the significantly negative slope of $p_R(n)$ is only consistent with the symmetric London approach. However, what is unexpected and not witnessed before is that also the slope of $p_-(n)$ was significantly negative. This is incompatible not only with asymmetric approaches and the Prague approach, which all predict that it should be positive, but also with the London approach, on which it should be constant (if there is no at-grammar pressure) or positive (if there is additional at-grammar pressure). This effect is specific to SUD-trained Stanza.[8]

Interestingly, when the same full COCA was parsed with Stanza trained by us on the UD versions of the same training corpora, the slopes of all logistic regression models were significantly pos-

[7] Also other aspects of SUD representations, especially, the fact that constructions are headed by function rather than content words (e.g., PPs are headed by prepositions rather than nouns), make the resulting structures less ambiguous and easier to work with.

[8] Moreover, it seems that, at least to some extent, this effect was caused by the inclusion of spoken parts of COCA in the current study, unlike in PBG24, where only written parts of COCA were processed. After removing two conversational genres – spoken and TV/movies – the slope of $p_-(n)$, while still significantly negative, was much flatter (-0.00104 vs. -0.00466 in Figure 1), with p not reaching the < 0.001 significance level.

itive (see the right column of Figure 1), unlike in the SUD-based study (see the left column again), but also unlike in PBG24, where a large subset of COCA was parsed with Stanza pre-trained on UD and where $p_R(n)$ was monotonically decreasing. However, while this difference awaits explanation, the positive slope of $p_R(n)$ is compatible with both symmetric approaches to coordination: on the assumption of *any* at-grammar pressure, the slope of $p_R(n)$ is expected to be positive on the Prague approach, and if this at-grammar pressure is sufficiently strong, then the positive slope of $p_R(n)$ is also expected on the London approach. Moreover, it is important to note that the relations between slopes of these UD-based models are again as expected by symmetric theories of coordination: the slope is most positive in the case of $p_L(n)$ (top graph) and least positive in the case of $p_R(n)$ (bottom graph), with all relevant differences statistically significant ($p \ll 0.001$).

2.1.2 COCA Parsed with BNP

Another way to avoid the problems of UD representation of coordination was to use a constituency parser. To this end, we utilized the Berkeley Neural Parser (BNP; Kitaev and Klein 2018, Kitaev et al. 2019) with the `benepar_en3` model. All of COCA apart from the spoken genre was parsed – around 59.5M sentences. Only simple binary coordinations were extracted – constituents consisting of three children, where the middle child is a conjunction (e.g., *Lisa and Bart*) and constituents consisting of four children, where the first and third child constitute a conjunction (e.g., *either Marge or Homer*). This way the problem of the exact extents of conjuncts was avoided. However, unlike dependency representations, the PTB format produced by BNP does not contain a clear information about governors, so heuristics similar to those used in PW23 were employed. In the process, information about 13,543,340 coordinations was extracted.

The quality of the resulting data was evaluated using exactly the same procedure as in PBG24: for each governor position (left, absent, right) and each conjunct length difference (from 1 to 20 words), 15 coordinations were sampled and checked for correctness (understood as in PBG24: the right conjuncts and the right position of the governor). The data quality was much higher than in the case of Stanza-parsed dataset used in PBG24: 78.11% of coordinations were judged as correct here, as opposed to 50.1% in PBG24.

The results are analogous to those based on UD-trained Stanza reported in the previous section: 1) all three slopes were significantly positive (with $p \ll 0.001$), 2) that of $p_L(n)$ was most positive (0.112), followed by $p_-(n)$ (0.085), and by $p_R(n)$ (0.029), with all differences highly statistically significant ($p \ll 0.001$). Again, this is compatible with both symmetric approaches to coordination (assuming different strengths of at-grammar pressure), but not with asymmetric approaches.

2.2 Polish

Two Polish studies follow PW23: they are based on manually-annotated treebanks, i.e., on small but relatively high-quality resources.

2.2.1 Składnica Constituency Parsebank

The first is based on Składnica, a manually-disambiguated constituency parsebank of Polish (Woliński et al. 2011, 2018) containing 14K sentences. As this is a much smaller corpus than $\text{PTB}_{\&}$ (49.2K sentences), the first and last conjuncts of all coordinations, not just binary, were taken into account, resulting in 5395 extracted coordinations (including 4800 binary; vs. 21,825 in $\text{PTB}_{\&}$).

The results are similar to those obtained by PW23. First of all, both $p_L(n)$ and $p_-(n)$ are monotonically increasing with statistically highly significant ($p \ll 0.001$) positive slopes; this is compatible with all approaches, although in the case of the London approach only with the assumption of an at-grammar tendency for shorter first conjuncts. Second, the slope of $p_L(n)$ is statistically significantly ($p < 0.05$) greater than that of $p_-(n)$ (0.18 vs. 0.09); this is only explained by the symmetric approaches. Third, while the slope of $p_R(n)$ is also positive (0.025), this value is not significantly different than 0 ($p > 0.05$); this is again more in line with symmetric approaches. Finally, while the difference of slopes of $p_-(n)$ and $p_R(n)$ is not statistically significant, the difference of slopes of $p_L(n)$ and $p_R(n)$ is ($p < 0.05$), which is not compatible with asymmetric approaches (on which all slopes should be the same), but immediately explained by both symmetric approaches.

2.2.2 Polish Dependency Bank

The second study is based on Polish Dependency Bank (PDB; Wróblewska 2014), a pre-UD dependency treebank in which coordinations are annotated according to the Prague approach (so they were free from the ambiguity problem mentioned

above). The version of PDB used in this study contains over 22K sentences. Again, all coordinations were taken into account: 13,247 were extracted, including 11,635 binary coordinations.

The results of this study are similar to those of the previous one. First, all three slopes are monotonically increasing, but this time all positive slopes are statistically significant ($p \ll 0.001$ for $p_L(n)$ and $p_-(n)$, $p < 0.05$ for $p_R(n)$). Second, the relation between the three slopes is as expected by symmetric theories of coordination: greatest for $p_L(n)$ (0.093), smaller for $p_-(n)$ (0.073), and smallest for $p_R(n)$ (0.055); however, this time the differences between these slopes did not turn out to be statistically significant.[9]

In summary, the results based on PDB alone are not sufficient to distinguish between symmetric and asymmetric approaches to coordination: the relevant differences, while in line with symmetric approaches, are not statistically significant. However, these results are compatible with those based on Składnica, where most of the crucial differences are statistically significant and, hence, provide an argument from Polish for the symmetric structure of coordination.

2.3 Partial Summary and Discussion

The results of previous work and our own studies are presented in Table 1.[10] In the $L/-$ column, '+' means that the slope of $p_L(n)$ is greater (more positive) than that of $p_-(n)$, and '−' that it is smaller (more negative), and analogously in the next two columns. Recall that the prediction of both symmetric approaches is that the slope of $p_L(n)$ is greatest and that of $p_R(n)$ smallest, so an ideal confirmation of such approaches would have a sequence of statistically significant +'s in these three columns. On the other hand, according to asymmetric approaches there should be no slope differences, so a sequence of statistically insignificant differences is expected. The final R column presents the sign of the slope of $p_R(n)$; the negative sign, '−', is compatible with the London approach, but not with the Prague approach. The number of asterisks reflects levels of statistical significance: *** for $p < 0.001$, ** for $p < 0.01$, * for $p < 0.05$; additionally, when $p > 0.1$, + or − is in grey.[11]

Table 1: Summary of studies described above: number of sentences, number of extracted coordinations, comparisons of slopes (see explanation in text)

	sents	coords	$L/-$	$-/R$	L/R	R
PW23	49.2K	21.8K	−	$+^{***}$	$+^{**}$	−
PBG24	21.8M	11.5M	$+^{***}$	$+^{***}$	$+^{***}$	$-^{***}$
St./SUD	69.2M	14.3M	$+^{***}$	$+^{***}$	$+^{***}$	$-^{***}$
St./UD	69.2M	10.8M	$+^{***}$	$+^{***}$	$+^{***}$	$+^{***}$
BNP	59.5M	13.5M	$+^{***}$	$+^{***}$	$+^{***}$	$+^{***}$
Składnica	14.0K	5.4K	$+^{*}$	+	$+^{*}$	+
PDB	22.2K	13.2K	+	+	+	$+^{*}$

What all English models have in common is that the differences between the slopes are as predicted by the symmetric approaches to coordination: when the difference is significant, it is always +. However, the studies reported here also show that the effect of parser is clearly visible, with the slope of $p_R(n)$ – crucial for a potential argument for the superiority of the London symmetric approach over the Prague approach – sometimes significantly negative (PBG24, Stanza/SUD), and sometimes significantly positive (Stanza/UD, BNP). Given that – as shown by the evaluation of extracted coordinations – the quality of input to these models was highest in the case of COCA parsed with BNP, these results seem to be most reliable. Hence, the conclusion of PBG24 that – given the negative slope of $p_R(n)$ – the London approach is the only one compatible with corpus data might have been premature. That is, the current conclusion must be that asymmetric approaches are clearly incompatible with corpus data, but – contrary to the conjecture of PBG24 – the resulting models are not sufficiently reliable to distinguish between the two symmetric approaches. Note that this conclusion is compatible with PW23's results, which were inconclusive about the slope of $p_R(n)$, as well as with the results of our Polish studies, according to which the slope of $p_R(n)$ is positive (significantly so, according to the PDB-based study).

2.4 Other Languages (UD Corpora)

We also performed similar studies on the basis of UD corpora of 10 languages (version 2.14; Zeman et al. 2024). We only considered clearly head-ini-

[9] Recall that we assume that lengths are measured in words. When they are measured in syllables, the difference between $p_L(n)$ and $p_R(n)$ turns out to be statistically significant ($p < 0.05$), while the character metric renders the difference between $p_L(n)$ and $p_-(n)$ statistically significant ($p < 0.05$).

[10] The small ratios of coordinations to sentences in St./(S)UD and BNP rows is probably caused by the inclusion of conversational genres (spoken in all three, TV/movies also in St./(S)UD), characterized by a very large number of very short – coordination-free – sentences.

[11] PBG24 do not report the levels of statistical significance for slope differences; we estimated these levels on the basis of their raw data, made available to us.

tial languages with at least 700K tokens in UD corpora, i.e., 5 Romance languages (Italian, Latin, Portuguese, Romanian, Spanish), 2 Germanic (English, Icelandic), and 3 Slavic (Czech, Polish – exceptionally, even though it had less than 700K tokens, Russian).[12] See Table 2.

Table 2: Sizes of – and results based on – UD datasets: number of tokens, number of extracted coordinations, comparisons of slopes (see explanation in text)

	tokens	coords	$L/-$	$-/R$	L/R	R
it	864K	25,426	$+^{***}$	$+$	$+^{*}$	$+^{**}$
la	983K	39,510	$+^{*}$	$+^{***}$	$+^{***}$	$-$
pt	1,361K	29,255	$+^{***}$	$+$	$+^{**}$	$+^{**}$
ro	938K	37,247	$+$	$+^{***}$	$+^{***}$	$+$
es	1,002K	28,666	$+^{***}$	$+$	$+^{*}$	$+^{**}$
en	718K	21,013	$-$	$+^{**}$	$+^{**}$	$-$
is	1,183K	43,852	$+^{***}$	$-$	$+^{*}$	$+^{*}$
cs	2,249K	90,566	$-^{***}$	$+^{***}$	$+^{***}$	$-^{*}$
pl	497K	16,684	$-$	$+^{*}$	$+$	$+$
ru	1,896K	61,004	$+$	$+^{***}$	$+^{***}$	$-$

While many differences are statistically insignificant, a fact that may be explained by the relatively small sizes of corpora used, it is clear that the results of this study are overall only compatible with the symmetric approaches.

This is most clear in the case of the Romance languages, where all differences are in the positive direction expected by symmetric approaches and the L/R difference is always statistically significant. While only in the case of Latin are all three differences statistically significant, for all Romance languages at least one of the differences $L/-$ and $-/R$ is highly statistically significant ($p < 0.001$), contra asymmetric approaches.

All statistically significant differences are in the 'right' positive direction also in the case of the two Germanic languages, English and Icelandic, and similarly for two of the Slavic languages, Russian and Polish, even if only one difference reaches the level of statistical significance in the case of Polish (probably because of the very small dataset).

[12]We used Typometrics (https://typometrics.elizia.net/; Gerdes et al. 2021b) to estimate headedness: we considered a language head-initial if it scored over 50% on two measures: the percentage of adpositional constructions with the adposition preceding its proper noun object (ADP-comp:obj-PROPN) and the percentage of verbal phrases with the verb preceding its proper noun object (VERB-comp:obj-PROPN). While all selected languages scored close to 100% on the adpositional measure, they differed widely on the verbal measure: from 51% for Latin, 72% for Czech, 81% for Russian, and 85% for Polish, to over 99% for English, Portuguese, and Italian. By contrast, two prototypically head-final languages – Korean and Turkish – scored 0% on the adpositional measure and, respectively, 0% and 3% on the verbal measure. German is not included, as it scored below 50% on the verbal measure.

Finally, Czech is an outlier in this study, in that the $L/-$ difference is highly significantly *negative*. However, the difference of slopes of $p_L(n)$ and $p_-(n)$ is relatively small: 0.0404 vs. 0.0553. Moreover, while both these slopes are positive, the slope of $p_R(n)$ is statistically significantly ($p < 0.05$) negative, -0.0167, which speaks not only against asymmetric approaches, on which it should be positive, but also – a little ironically – against the Prague approach, on which the slope of $p_R(n)$ should be 0 or positive. We leave the investigation of this outlier for future work.

3 Conclusion

At the most general level, the main contribution of this paper is a demonstration that extensive replication is crucial not only in psychology, medicine, and social sciences, but also in formal and computational linguistics. PBG24's replication of PW23's argument for the symmetry of coordination seemed to narrow down potentially valid representations of coordination from the two symmetric approaches to just the London approach, but the current more extensive replication invalidates this conjecture. While two of our studies (Stanza/SUD and UD/cs) result in models with negative slopes of $p_R(n)$, compatible only with the London approach, 7 other studies (Stanza/UD, BNP, PDB, UD/it, UD/pt, UD/es, UD/is) – including two based on English data, just as Stanza/SUD and PBG24 – result in models with significantly positive slopes of $p_R(n)$. Clearly, the choice of parser and dataset is important for the argument, and future research should determine how exactly it influences the results.

Nevertheless, the current studies add strong cross-linguistic arguments for the main claim of PW23 and PBG24, namely, that corpora provide quantitative evidence for the symmetry of coordination. Apart from UD/cs, where the unexpected statistically significant $L/-$ difference was observed, and PDB, where relevant differences were not statistically significant, in all other 13 models statistically significant slope differences were found that are only compatible with symmetric approaches.

An important limitation of this paper is that it only considers head-initial languages, as the above reasoning assumes that heads of conjuncts are conjunct-initial on average. An investigation of the structure of coordination in two head-final languages, Korean or Turkish, may be found in Stempniak 2024a.

Acknowledgements

The corpus-based studies discussed in this paper were performed within the Bachelor's Group Research Project "Corpus-based investigation of the symmetry of coordination" carried out in 2022–2024 at the Cognitive Science programme at the University of Warsaw. The project was supervised by the first author, all authors were participants. Particular studies are described in more detail in:

- Borysiak 2024 – COCA parsed with Stanza/SUD (§2.1.1),[13,14]
- Pobożniak 2024 – COCA Parsed with BNP (§2.1.2),[15]
- Okrasiński 2023 – Składnica Constituency Parsebank (§2.2.1),[16]
- Tomaszek 2023 – Polish Dependency Bank (§2.2.2),[17]
- Stempniak 2024b – UD Corpora (§2.4).[18]

We are grateful to Bartosz Maćkiewicz for his statistical assistance in some of these studies, as well as to TLT 2024 reviewers for their comments.

References

Otto Behaghel. 1909. Beziehungen zwischen Umfang und Reihenfolge von Satzgliedern. *Indogermanische Forschungen*, 25:110–142.

Magdalena Borysiak. 2024. Dependency structure of English coordination: A surface-syntactic approach. Bachelor's thesis, University of Warsaw.

Mark Davies. 2008–2023. The Corpus of Contemporary American English (COCA). Available online at https://www.english-corpora.org/coca/.

Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. Universal Dependencies. *Computational Linguistics*, 47(2):255–308.

Ramon Ferrer-i-Cancho. 2004. Euclidean distance between syntactically linked words. *Physical Review E*, 70:056135.

Jessica Ficler and Yoav Goldberg. 2016. Coordination annotation extension in the Penn Tree Bank. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 834–842, Berlin, Germany.

[13] https://github.com/bmagdab/sud-coords

[14] https://github.com/glowak/dlm

[15] https://github.com/KattGaii/coca-thesis

[16] https://github.com/Adokr/korpus

[17] https://github.com/kvmilos/PracaLicencjacka

[18] https://github.com/wjstempniak/Dependency-Structure-of-Coordination

Richard Futrell and Roger Levy. 2017. Noisy-context surprisal as a human sentence processing cost model. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2009): Volume 1, Long Papers*, pages 688–698, Valencia, Spain.

Richard Futrell, Roger P. Levy, and Edward Gibson. 2020. Dependency locality as an explanatory principle for word order. *Language*, 96(2):371–412.

Kim Gerdes, Bruno Guillaume, Sylvain Kahane, and Guy Perrier. 2018. SUD or Surface-Syntactic Universal Dependencies: An annotation scheme near-isomorphic to UD. In *Proceedings of the Second Workshop on Universal Dependencies (UDW 2018)*, pages 66–74. Association for Computational Linguistics.

Kim Gerdes, Bruno Guillaume, Sylvain Kahane, and Guy Perrier. 2021a. Starting a new treebank? Go SUD! In *Proceedings of the Sixth International Conference on Dependency Linguistics (DepLing, Syntax Fest 2021)*, pages 35–46, Sofia, Bulgaria.

Kim Gerdes, Bruno Guillaume, Sylvain Kahane, and Guy Perrier. 2021b. Typometrics: From implicational to quantitative universals in word order typology. *Glossa: A Journal of General Linguistics*, 6(1):1–31.

Edward Gibson. 1998. Linguistic complexity: Locality of syntactic dependencies. *Cognition*, 68(1):1–76.

Daniel Gildea and David Temperley. 2007. Optimizing grammars for minimum dependency length. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics*, pages 184–191, Prague.

Daniel Gildea and David Temperley. 2010. Do grammars minimize dependency length? *Cogntiive Science*, 34(2):286–310.

John A. Hawkins. 1994. *A Performance Theory of Order and Constituency*. Cambridge University Press, Cambridge.

Nikita Kitaev, Steven Cao, and Dan Klein. 2019. Multilingual constituency parsing with self-attention and pre-training. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3499–3505, Florence, Italy. Association for Computational Linguistics.

Nikita Kitaev and Dan Klein. 2018. Constituency parsing with a self-attentive encoder. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686, Melbourne, Australia. Association for Computational Linguistics.

Russell V. Lenth. 2024. *emmeans: Estimated Marginal Means, aka Least-Squares Means*. R package version 1.10.0.

Haitao Liu. 2008. Dependency distance as a metric of language. *Journal of Cognitive Science*, 9(2):159–191.

Haitao Liu, Chunshan Xu, and Junying Liang. 2017. Dependency distance: A new perspective on syntactic patterns in natural languages. *Physics of Life Reviews*, 21:171–193.

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

Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. Universal Dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation, LREC 2016*, pages 1659–1666, Portorož, Slovenia. European Language Resources Association (ELRA).

Adam Okrasiński. 2023. Minimalizacja długości zależności w strukturach współrzędnie złożonych: badanie korpusowe na podstawie banku drzew Składnica. Bachelor's thesis, University of Warsaw.

Timothy Osborne and Kim Gerdes. 2019. The status of function words in dependency grammar: A critique of Universal Dependencies (UD). *Glossa: A Journal of General Linguistics*, 4(17).

Bartosz Pobożniak. 2024. Analysing dependency structure of coordination using a constituency parser and Dependency Length Minimization. Bachelor's thesis, University of Warsaw.

Martin Popel, David Mareček, Jan Štěpánek, Daniel Zeman, and Zděněk Žabokrtský. 2013. Coordination structures in dependency treebanks. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 517–527, Sofia, Bulgaria.

Adam Przepiórkowski, Magdalena Borysiak, and Adam Głowacki. 2024. An argument for symmetric coordination from Dependency Length Minimization: A replication study. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1021–1033, Torino, Italy. ELRA and ICCL.

Adam Przepiórkowski and Agnieszka Patejuk. 2019. Nested coordination in Universal Dependencies. In *Proceedings of the Third Workshop on Universal Dependencies (UDW, SyntaxFest 2019)*, pages 58–69. Association for Computational Linguistics.

Adam Przepiórkowski and Michał Woźniak. 2023. Conjunct lengths in English, Dependency Length Minimization, and dependency structure of coordination. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15494–15512, Toronto, Canada. 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.

R Core Team. 2024. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna.

Manuela Sanguinetti and Cristina Bosco. 2014. PartTUT: The Turin University Parallel Treebank. In Roberto Basili, Cristina Bosco, Rodolfo Delmonte, Alessandro Moschitti, and Maria Simi, editors, *Harmonization and Development of Resources and Tools for Italian Natural Language Processing within the PARLI Project*, pages 51–69. Springer-Verlag.

Sebastian Schuster and Christopher D. Manning. 2016. Enhanced English Universal Dependencies: An improved representation for natural language understanding tasks. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation, LREC 2016*, pages 2371–2378, Portorož, Slovenia. European Language Resources Association (ELRA).

Natalia Silveira, Timothy Dozat, Marie-Catherine de Marneffe, Samuel Bowman, Miriam Connor, John Bauer, and Christopher D. Manning. 2014. A gold standard dependency corpus for English. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014*, pages 2897–2904, Reykjavík, Iceland. European Language Resources Association (ELRA).

Wojciech Stempniak. 2024a. Dependency structure of coordination in head-final languages: A Dependency-Length-Minimization-based study. In *Proceedings of the 22nd Workshop on Treebanks and Linguistic Theories (TLT 2024)*, Hamburg, Germany.

Wojciech Stempniak. 2024b. Struktura zależnościowa koordynacji – analiza korpusów Universal Dependencies. Bachelor's thesis, University of Warsaw.

David Temperley and Daniel Gildea. 2018. Minimizing syntactic dependency lengths: Typological/cognitive universal? *Annual Review of Linguistics*, 4:67–80.

Kamil Tomaszek. 2023. Minimalizacja długości zależności w strukturach współrzędnie złożonych: badanie korpusowe na podstawie Polish Dependency Bank. Bachelor's thesis, University of Warsaw.

Marcin Woliński, Katarzyna Głowińska, and Marek Świdziński. 2011. A preliminary version of Składnica—a treebank of Polish. In *Proceedings of the 5th Language & Technology Conference: Human Language Technologies as a Challenge for Computer Science and Linguistics*, pages 299–303, Poznań, Poland.

Marcin Woliński, Elżbieta Hajnicz, and Tomasz Bartosiak. 2018. A new version of the Składnica treebank of Polish harmonised with the Walenty valency dictionary. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, pages 1839–1844, Paris, France. European Language Resources Association (ELRA).

Alina Wróblewska. 2014. *Polish Dependency Parser Trained on an Automatically Induced Dependency Bank*. Ph.D. dissertation, Institute of Computer Science, Polish Academy of Sciences, Warsaw.

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

Daniel Zeman, Joakim Nivre, Mitchell Abrams, Elia Ackermann, Noëmi Aepli, Hamid Aghaei, Željko Agić, Amir Ahmadi, Lars Ahrenberg, Chika Kennedy Ajede, Gabrielė Aleksandravičiūtė, Ika Alfina, Avner Algom, Erik Andersen, Lene Antonsen, Katya Aplonova, Angelina Aquino, Carolina Aragon, Glyd Aranes, Maria Jesus Aranzabe, Bilge Nas Arıcan, Þórunn Arnardóttir, Gashaw Arutie, Jessica Naraiswari Arwidarasti, Masayuki Asahara, Deniz Baran Aslan, Cengiz Asmazoğlu, Luma Ateyah, Furkan Atmaca, Mohammed Attia, Aitziber Atutxa, Liesbeth Augustinus, Elena Badmaeva, Keerthana Balasubramani, Miguel Ballesteros, Esha Banerjee, Sebastian Bank, Verginica Barbu Mititelu, Starkaður Barkarson, Rodolfo Basile, Victoria Basmov, Colin Batchelor, John Bauer, Seyyit Talha Bedir, Kepa Bengoetxea, Yifat Ben Moshe, Gözde Berk, Yevgeni Berzak, Irshad Ahmad Bhat, Riyaz Ahmad Bhat, Erica Biagetti, Eckhard Bick, Agnė Bielinskienė, Kristín Bjarnadóttir, Rogier Blokland, Victoria Bobicev, Loïc Boizou, Emanuel Borges Völker, Carl Börstell, Cristina Bosco, Gosse Bouma, Sam Bowman, Adriane Boyd, Anouck Braggaar, Kristina Brokaitė, Aljoscha Burchardt, Marie Candito, Bernard Caron, Gauthier Caron, Lauren Cassidy, Tatiana Cavalcanti, Gülşen Cebiroğlu Eryiğit, Flavio Massimiliano Cecchini, Giuseppe G. A. Celano, Slavomír Čéplö, Neslihan Cesur, Savas Cetin, Özlem Çetinoğlu, Fabricio Chalub, Shweta Chauhan, Ethan Chi, Taishi Chika, Yongseok Cho, Jinho Choi, Jayeol Chun, Juyeon Chung, Alessandra T. Cignarella, Silvie Cinková, Aurélie Collomb, Çağrı Çöltekin, Miriam Connor, Daniela Corbetta, Marine Courtin, Mihaela Cristescu, Philemon Daniel, Elizabeth Davidson, Mathieu Dehouck, Martina de Laurentiis, Marie-Catherine de Marneffe, Valeria de Paiva, Mehmet Oguz Derin, Elvis de Souza, Arantza Diaz de Ilarraza, Carly Dickerson, Arawinda Dinakaramani, Elisa Di Nuovo, Bamba Dione, Peter Dirix, Kaja Dobrovoljc, Timothy Dozat, Kira Droganova, Puneet Dwivedi, Hanne Eckhoff, Sandra Eiche, Marhaba Eli, Ali Elkahky, Binyam Ephrem, Olga Erina, Tomaž Erjavec, Aline Etienne, Wograine Evelyn, Sidney Facundes, Richárd Farkas, Federica Favero, Jannatul Ferdaousi, Marília Fernanda, Hector Fernandez Alcalde, Jennifer Foster, Cláudia Freitas, Kazunori Fujita, Katarína Gajdošová, Daniel Galbraith, Federica Gamba, Marcos Garcia, Moa Gärdenfors, Sebastian Garza, Fabrício Ferraz Gerardi, Kim Gerdes, Filip Ginter, Gustavo Godoy, Iakes Goenaga, Koldo Gojenola, Memduh Gökırmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta González Saavedra, Bernadeta Griciūtė, Matias Grioni, Loïc Grobol, Normunds Grūzītis, Bruno Guillaume, Céline Guillot-Barbance, Tunga Güngör, Nizar Habash, Hinrik Hafsteinsson, Jan Hajič, Jan Hajič jr., Mika Hämäläinen, Linh Hà Mỹ, Na-Rae Han, Muhammad Yudistira Hanifmuti, Takahiro Harada, Sam Hardwick, Kim Harris, Dag Haug, Johannes Heinecke, Oliver Hellwig, Felix Hennig, Barbora Hladká, Jaroslava Hlaváčová, Florinel Hociung, Petter Hohle, Jena Hwang, Takumi Ikeda, Anton Karl Ingason, Radu Ion, Elena Irimia, Ọlájídé Ishola, Kaoru Ito, Siratun Jannat, Tomáš Jelínek, Apoorva Jha, Anders Johannsen, Hildur Jónsdóttir, Fredrik Jørgensen, Markus Juutinen, Sarveswaran K, Hüner Kaşıkara, Andre Kaasen, Nadezhda Kabaeva, Sylvain Kahane, Hiroshi Kanayama, Jenna Kanerva, Neslihan Kara, Ritván Karahóǧa, Boris Katz, Tolga Kayadelen, Jessica Kenney, Václava Kettnerová, Jesse Kirchner, Elena Klementieva, Elena Klyachko, Arne Köhn, Abdullatif Köksal, Kamil Kopacewicz, Timo Korkiakangas, Mehmet Köse, Natalia Kotsyba, Jolanta Kovalevskaitė, Simon Krek, Parameswari Krishnamurthy, Sandra Kübler, Oğuzhan Kuyrukçu, Aslı Kuzgun, Sookyoung Kwak, Veronika Laippala, Lucia Lam, Lorenzo Lambertino, Tatiana Lando, Septina Dian Larasati, Alexei Lavrentiev, John Lee, Phuong Lê Hồng, Alessandro Lenci, Saran Lertpradit, Herman Leung, Maria Levina, Cheuk Ying Li, Josie Li, Keying Li, Yuan Li, KyungTae Lim, Bruna Lima Padovani, Krister Lindén, Nikola Ljubešić, Olga Loginova, Stefano Lusito, Andry Luthfi, Mikko Luukko, Olga Lyashevskaya, Teresa Lynn, Vivien Macketanz, Menel Mahamdi, Jean Maillard, Aibek Makazhanov, Michael Mandl, Christopher Manning, Ruli Manurung, Büşra Marşan, Cătălina Mărănduc, David Mareček, Katrin Marheinecke, Stella Markantonatou, Héctor Martínez Alonso, Lorena Martín Rodríguez, André Martins, Jan Mašek, Hiroshi Matsuda, Yuji Matsumoto, Alessandro Mazzei, Ryan McDonald, Sarah McGuinness, Gustavo Mendonça, Tatiana Merzhevich, Niko Miekka, Karina Mischenkova, Margarita Misirpashayeva, Anna Missilä, Cătălin Mititelu, Maria Mitrofan, Yusuke Miyao, AmirHossein Mojiri Foroushani, Judit Molnár, Amirsaeid Moloodi, Simonetta Montemagni, Amir More, Laura Moreno Romero, Giovanni Moretti, Keiko Sophie Mori, Shinsuke Mori, Tomohiko Morioka, Shigeki Moro, Bjartur Mortensen, Bohdan Moskalevskyi, Kadri Muischnek, Robert Munro, Yugo Murawaki, Kaili Müürisep, Pinkey Nainwani, Mariam Nakhlé, Juan Ignacio Navarro Horñiacek, Anna Nedoluzhko, Gunta Nešpore-Bērzkalne, Manuela Nevaci, Luong Nguyễn Thị, Huyền Nguyễn Thị Minh, Yoshihiro Nikaido, Vitaly Nikolaev, Rattima Nitisaroj, Alireza Nourian, Hanna Nurmi, Stina Ojala, Atul Kr. Ojha, Adédayọ̀ Olúòkun, Mai Omura, Emeka Onwuegbuzia, Noam Ordan, Petya Osenova, Robert Östling, Lilja Øvrelid, Şaziye Betül Özateş, Merve Özçelik,

Arzucan Özgür, Balkız Öztürk Başaran, Teresa Paccosi, Alessio Palmero Aprosio, Hyunji Hayley Park, Niko Partanen, Elena Pascual, Marco Passarotti, Agnieszka Patejuk, Guilherme Paulino-Passos, Giulia Pedonese, Angelika Peljak-Łapińska, Siyao Peng, Cenel-Augusto Perez, Natalia Perkova, Guy Perrier, Slav Petrov, Daria Petrova, Andrea Peverelli, Jason Phelan, Jussi Piitulainen, Tommi A Pirinen, Emily Pitler, Barbara Plank, Thierry Poibeau, Larisa Ponomareva, Martin Popel, Lauma Pretkalniņa, Sophie Prévost, Prokopis Prokopidis, Adam Przepiórkowski, Tiina Puolakainen, Sampo Pyysalo, Peng Qi, Andriela Rääbis, Alexandre Rademaker, Mizanur Rahoman, Taraka Rama, Loganathan Ramasamy, Carlos Ramisch, Fam Rashel, Mohammad Sadegh Rasooli, Vinit Ravishankar, Livy Real, Petru Rebeja, Siva Reddy, Mathilde Regnault, Georg Rehm, Ivan Riabov, Michael Rießler, Erika Rimkutė, Larissa Rinaldi, Laura Rituma, Putri Rizqiyah, Luisa Rocha, Eiríkur Rögnvaldsson, Mykhailo Romanenko, Rudolf Rosa, Valentin Roșca, Davide Rovati, Ben Rozonoyer, Olga Rudina, Jack Rueter, Kristján Rúnarsson, Shoval Sadde, Pegah Safari, Benoît Sagot, Aleksi Sahala, Shadi Saleh, Alessio Salomoni, Tanja Samardžić, Stephanie Samson, Manuela Sanguinetti, Ezgi Sanıyar, Dage Särg, Baiba Saulīte, Yanin Sawanakunanon, Shefali Saxena, Kevin Scannell, Salvatore Scarlata, Nathan Schneider, Sebastian Schuster, Lane Schwartz, Djamé Seddah, Wolfgang Seeker, Mojgan Seraji, Syeda Shahzadi, Mo Shen, Atsuko Shimada, Hiroyuki Shirasu, Yana Shishkina, Muh Shohibussirri, Dmitry Sichinava, Janine Siewert, Einar Freyr Sigurðsson, Aline Silveira, Natalia Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Kiril Simov, Maria Skachedubova, Aaron Smith, Isabela Soares-Bastos, Shafi Sourov, Carolyn Spadine, Rachele Sprugnoli, Vivian Stamou, Steinþór Steingrímsson, Antonio Stella, Milan Straka, Emmett Strickland, Jana Strnadová, Alane Suhr, Yogi Lesmana Sulestio, Umut Sulubacak, Shingo Suzuki, Daniel Swanson, Zsolt Szántó, Chihiro Taguchi, Dima Taji, Yuta Takahashi, Fabio Tamburini, Mary Ann C. Tan, Takaaki Tanaka, Dipta Tanaya, Mirko Tavoni, Samson Tella, Isabelle Tellier, Marinella Testori, Guillaume Thomas, Sara Tonelli, Liisi Torga, Marsida Toska, Trond Trosterud, Anna Trukhina, Reut Tsarfaty, Utku Türk, Francis Tyers, Sumire Uematsu, Roman Untilov, Zdeňka Urešová, Larraitz Uria, Hans Uszkoreit, Andrius Utka, Elena Vagnoni, Sowmya Vajjala, Rob van der Goot, Martine Vanhove, Daniel van Niekerk, Gertjan van Noord, Viktor Varga, Uliana Vedenina, Eric Villemonte de la Clergerie, Veronika Vincze, Natalia Vlasova, Aya Wakasa, Joel C. Wallenberg, Lars Wallin, Abigail Walsh, Jing Xian Wang, Jonathan North Washington, Maximilan Wendt, Paul Widmer, Shira Wigderson, Sri Hartati Wijono, Seyi Williams, Mats Wirén, Christian Wittern, Tsegay Woldemariam, Tak-sum Wong, Alina Wróblewska, Mary Yako, Kayo Yamashita, Naoki Yamazaki, Chunxiao Yan, Koichi Yasuoka, Marat M. Yavrumyan, Arife Betül Yenice, Olcay Taner Yıldız, Zhuoran Yu, Arlisa Yuliawati, Zdeněk Žabokrtský, Shorouq Zahra, Amir Zeldes, He Zhou, Hanzhi Zhu, Anna Zhuravleva, and Rayan Ziane. 2024. Universal Dependencies 2.14. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.