

Quality and Efficiency of Manual Annotation: Pre-annotation Bias

Marie Mikulová, Milan Straka, Jan Štěpánek, Barbora Štěpánková, Jan Hajič

Institute of Formal and Applied Linguistics

Computer Science School, Faculty of Mathematics and Physics, Charles University, Prague

{mikulova,straka,stepanek,stepankova,hajic}@ufal.mff.cuni.cz

Abstract

This paper presents an analysis of annotation using an automatic pre-annotation for a mid-level annotation complexity task - dependency syntax annotation. It compares the annotation efforts made by annotators using a pre-annotated version (with a high-accuracy parser) and those made by fully manual annotation. The aim of the experiment is to judge the final annotation quality when pre-annotation is used. In addition, it evaluates the effect of automatic linguistically-based (rule-formulated) checks and another annotation on the same data available to the annotators, and their influence on annotation quality and efficiency. The experiment confirmed that the pre-annotation is an efficient tool for faster manual syntactic annotation which increases the consistency of the resulting annotation without reducing its quality.

Keywords: annotation, syntax, pre-annotation bias, annotation efficiency, annotation quality, inter-annotator agreement

1. Introduction

Developing high-quality annotated corpora representing natural language phenomena that can be used by machine learning tools or explored by linguists solving various research tasks is time-consuming and expensive. This paper presents results of a carefully designed experiment that shed some light on the issue of efficiency vs. quality of manual syntactic annotation. For this purpose, we used manually annotated parts of the Prague Dependency Treebank - Consolidated, PDT-C in sequel (Hajič et al., 2020)¹ under various experimental set-ups: annotation with no support (“from-scratch mode”), annotation supported by the use of automatic pre-annotation (using a high-accuracy parser), availability of other annotation on the same data, and “on-line” checking rules (implemented in such a way that the annotator can check the manual annotation or manual post-editing immediately). The experiments focus on dependency syntax annotation: each token of a sentence is assigned its head and specific type of the dependency relation.

With the experiment, we aim to answer three main research questions, the first of which is crucial:

- **Does the quality of manual annotation remain acceptably high if pre-annotation is used?** How much the manual correction of automatic pre-annotation increases quality? The concern is that annotators might tend to simply go along with the pre-annotation, which would lead to a lower quality than they could produce by a fully manual annotation (cf. similar question in Rehbein et al. (2009)).
- What effect do other available tools (checking rules, other annotation on the same data) have on

the quality and efficiency of the annotation and to what extent?

- Which set-up is most useful for manual dependency syntax annotation? We want to find out how to best prepare the annotation environment and workflow for the syntactic annotation of 2,000,000 tokens in the PDT-C project (see Sect. 3) and possibly generalize the findings for other similar projects.

To evaluate the results, we use the usual methods for measuring accuracy and inter-annotator agreement in the dependency syntax tasks (UAS, LAS, Cohen’s kappa).

The experiment confirmed the usefulness of automatic pre-annotation for efficiency of annotation and quality of the result. The influence of the other support tools on efficiency and quality of annotation was not as positive as expected. However, when used together they significantly increase the quality of the annotation.

In the following Sect. 2, we summarize previous related work on similar tasks. In Sect. 3, the annotation task is briefly described. Annotation support tools used within the experiment (parser, automatic checks, availability of other annotation) are presented in Sect. 4. The core of the article is Sect. 5, in which the experiment is described and the results are presented and evaluated. Conclusions and future work are summarized in Sect. 6.

2. Related Work

Various time-saving methods while maintaining data quality have already been explored. An overview of studies on this topic is given by Grouin et al. (2014). A number of previous works have demonstrated the usefulness of automatic pre-annotation by verifying it with various annotation experiments.

¹<https://hdl.handle.net/11234/1-3185>

The experiments were performed on different annotation tasks (POS annotation (Fort and Sagot, 2010), semantic role labeling (Chou et al., 2006; Rehbein et al., 2009), name entity annotation (Grouin et al., 2014; Rosset et al., 2013), treebank construction (Marcus et al., 1993; Chiou et al., 2001; Gupta et al., 2010; Voutilainen and Purtonen, 2011)) for different purposes (comparison of annotation quality by experienced and inexperienced annotators (Gupta et al., 2010), the effect of different types of pre-annotation, the difference between using a high-quality and lower-quality parser (Rehbein et al., 2009; Chiou et al., 2001)), which corresponds to the different set-ups of the experiments. Experimental design was chosen not only in line with research question that the authors wanted to address, but also in line with the number of annotators, the amount of data and time available.

The previous works also confirmed that a clear and controlled, well performed experimental design is necessary for minimizing side effects on the results of the experiment (annotator influence, learning effect, inappropriate or/and different data; cf. the design of the experiment in Rehbein et al. (2009), where an ongoing learning effect (repeated annotation of the same data by the same annotators) seems to have led to a distortion of the results).

The approaches that are closest to ours are that of treebank construction. One of the earliest studies that investigated the usefulness of automatic pre-annotation for treebank construction was carried out by Marcus et al. (1993). In the context of the development of the Penn Treebank, they compared the semi-automatic approach to a fully manual annotation and found that the semi-automatic method resulted both in a significant reduction of annotation time, and in increased inter-annotator agreement and accuracy. Chiou et al. (2001) investigated the effect of pre-annotation in the context of the development of the Chinese Penn Treebank. They experimented with two different parsers (using a high-accuracy and lower-accuracy parser) and found that both significantly shortened the total annotation time while maintaining accuracy. Gupta et al. (2010) described a semi-automatic approach to expedite the process of manual annotation of a Hindi Dependency Treebank. In addition to the influence of automatic pre-annotation, they also experimented with two groups of annotators: experienced and inexperienced. From their observations, it seems that a semi-automatic process is an effective way of doing dependency annotation when the annotators are trained and experienced. The experiment also confirmed the fact that treebank annotation is not a trivial task and supervision is required for carrying out the task in an efficient manner. Voutilainen and Purtonen (2011) documented a double-blind experiment on syntactic annotation of Finnish texts. The results suggest that as a result of reviews and negotiations, inter-annotator agreement can be much higher than the corresponding labeled attachment scores re-

ported for state-of-the-art dependency parsers. Metrics for measuring annotation quality have been discussed in Artstein and Poesio (2008). In a more recent work, Skjærholt (2014) pointed out the inadequacy of existing metrics for syntactic annotation and proposes his own metric adapted to the structure of syntactic annotation. However, the UAS and LAS (unlabeled and labeled attachment scores, respectively) for measuring accuracy (Yang and Li, 2018; Braggaar and van der Goot, 2021) and Cohen’s kappa (Cohen, 1960) for measuring inter-annotator agreement (Yang and Li, 2018; Nguyen, 2018) are still widely used metrics for the annotation quality evaluation in the dependency syntax area.

3. Annotation Task

The experiment focuses on annotation at the dependency syntax layer within the multi-layered treebank for Czech – PDT-C. Its version 1.0 features a fully manual morphological and syntactic-semantic annotation (Hajič et al., 2020). For version 2.0, the goal is to extend the fully manual mid-level dependency syntax annotation to all parts of PDT-C (this assumes an annotation of 2,000,000 tokens). For our experiments described in this paper, we have used a sample from the PCEDT part of PDT-C (Czech translation of the Wall Street Journal portion of Penn Treebank (Marcus et al., 1999)).

A dependency syntax annotation usually consists of determining the head for each node and assigning a syntactic function (attribute *afun*) describing the relation between the dependent node and its head. In PDT-style dependency syntax annotation, every word (token) of a sentence (including punctuation marks) is represented by a node of the tree and at the same time, no added nodes are allowed. The original style of this layer annotation is best described in Hajič et al. (1999)² based on principles first formulated in Hajič (1998).

The rules follow (wherever possible) traditional principles of the Czech grammar (especially as described in Šmilauer (1947)). In a prototypically structured sentence, a predicate (*Pred*) is the head of the sentence and it depends on the technical root. *Afuns* (25 types are distinguished) primarily describe dependency syntax relations such as subject (*Sb*), object (*Obj*), adverbial (*Adv*), attribute (*Attr*). Also all other nodes are assigned a head and a dependency relation (sometimes only in a purely technical way), e.g. auxiliary verbs (*AuxV*), prepositions (*AuxP*), subordinate conjunctions (*AuxC*), coordination (*Coord*) and apposition (*Apos*) conjunctions, particles (*AuxZ*) and punctuation (*AuxG*, *AuxX*, *AuxK*).

Additional syntactic features such as ellipsis (*_E*) and parenthesis (*_P*) and members of coordination (*_Co*) or apposition (*_Ap*) structure are denoted by a specific affixes attached to appropriate *afun*. These three types

²<https://ufal.mff.cuni.cz/pdt-c/publications/PDT20-a-man-en.pdf>

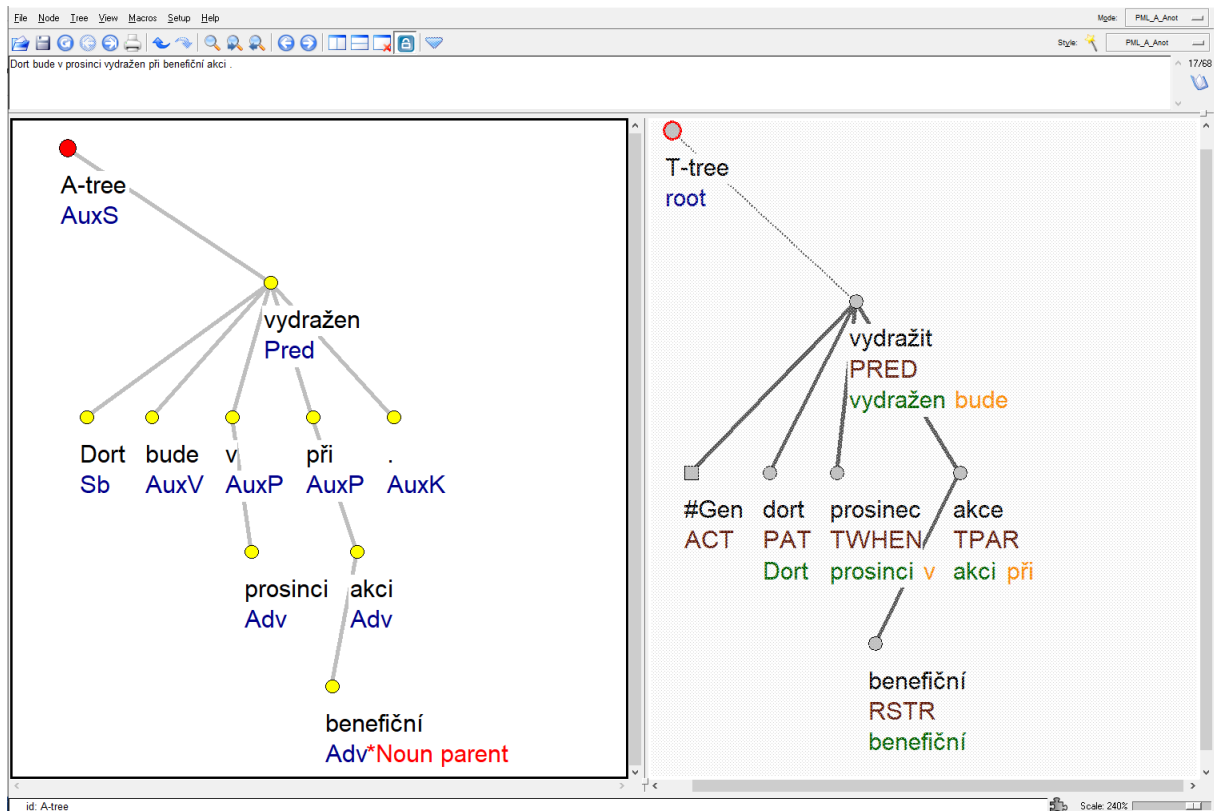


Figure 1: Annotation tool TrEd: Dependency syntax annotation of the sentence *Dort bude v prosinci vydražen při benefiční akci* ‘The cake will be auctioned at a charity event in December.’ (on the left) with syntactic-semantic annotation of the same sentence available in a separate window (on the right), displayed in the stylesheet available to annotators during the annotation. In the dependency syntax annotation, the output of the automatic “online” checking procedure is also shown: there is an error in the incorrectly assigned *afun* *Adv* to a node that depends on the noun. According to the annotation guidelines, any node that depends on a noun gets the *afun* *Attr*.

of affixes can combine with each other, increasing the total number of possible *afuns* to 200.³

An example of dependency syntax annotation displayed in the annotation tool (Sect. 4) is shown in Fig. 1 (compared to syntactic-semantic annotation of the same sentence; read more in Sect. 4.3).

4. Annotation Tools Used

For manual annotation, we use the TrEd annotation tool (Pajas and Štěpánek, 2008). The basic editing capabilities of TrEd allow the user to easily modify the tree structure with drag-and-drop operations. The annotation process is greatly accelerated by a set of functions

³Some combinations of *afun* and affix are in principle unacceptable, but the annotation tool (Sect. 4) allows to create any of them, so when measuring the results of our experiment (Sect. 5.2), we take into account both the number of basic *afuns* (25) and the number of all possible *afuns* ($25 * 2^3 = 200$). Unacceptable combinations of *afun* and affix are then subject to automatic checks (Sect. 4.2). Also, it should be noted that the rules for some *afuns*, ellipsis and affixes differ slightly from the original guidelines as referenced above.

called macros (activated by keyboard shortcuts), customized to simplify the most common tasks done by the annotators (e.g., in our annotation task (Sect. 3), macros are available for assigning any of the *afun* values). TrEd provides many tools to customize the editor for various tasks. When annotating dependency syntax, we also utilize the possibility to write macros that incorporate consistency tests into the annotation process (Sect. 4.2). TrEd can also display more than one annotation (more than one tree) in one editing window. An annotator can thus check his/her annotation against another (Sect. 4.3). The annotation tool TrEd as we use it for the annotation task described here is shown in Fig. 1.

4.1. Pre-annotation

To perform automatic pre-annotation, we selected a dependency parser achieving the best published performance on the Prague Dependency Treebank (PDT). To our best knowledge, that is the *UDPipe 2* dependency parser combined with RobeCzech (Straka et al., 2021), where RobeCzech is a Czech variant of RoBERTa.

We trained the parser on the PDT-style syntax including the *afun* affixes, using the train portion of PDT-C 1.0

(Hajič et al., 2020), reaching UAS 94.18% and LAS 91.37% on the PDT evaluation test set (the *etest* section). To ascertain the effect of out-of-domain data, we also evaluated the parser performance on 2,000 manually annotated sentences from PCEDT part of PDT-C, obtaining UAS 93.25% and LAS⁴ 90.93%, close to the PDT result despite the change of domain.

4.2. Annotation Checking Rules

To get a higher quality and a greater consistency of annotated data, a set of automatic “online” checking procedures has been proposed and created in accordance with the annotation guidelines and incorporated into the annotation process to prevent the annotator from making accidental mistakes, e.g. because of poor attention. The “online” checks are based on the fact that many annotation rules imply that particular phenomena cannot (or have to) occur in the annotation output. They mainly work with the combination of attribute values and the structure of a tree. In our experiments, 9 sets of automatic checking rules were available.

Some of the checking rules utilize the multi-layered structure of the PDT-C and use information from the lower morphological layer (e.g. morphological categories, lemmas). For example a rule states that an attribute (afun *Attr*) never depends on a verb and a node that depends on the noun is an attribute (afun *Attr*), or a rule states that a nominal part of a predicate (afun *Pnom*) always depends on a verb *být* ‘to be’ and its variants.

Many checks combine several rules, for example a check examines whether the afun *Pred* for predicate is a head of the tree, except cases when it is denoted as a parenthesis, or a member of coordination or apposition by affixes *_P*, *_Co*, or *_Ap*. Further checking rules are defined by enumeration, for example a check states that the head of a tree has only a limited set of possible afuns: *Coord* for coordination head, *Apos* for apposition head, *Pred* for predicate, *Denom* for noun phrase, *Part1* for interjection. Also some combinations of afun and affix are excluded, for example the punctuation marks *AuxX* and *AuxG* are never complemented with any affixes.

It is recommended to the annotators to run the checks after annotation of every single tree and consequently check and fix possible errors. In Fig. 1, we can see an example of output of the automatic checks. The automatic “online” checking procedure revealed an error in the incorrectly assigned afun *Adv* to a node that depends on the noun. As mentioned above, according to the annotation guidelines, any node that depends on a noun gets the afun *Attr*.

⁴Please note that the LAS score referred to here is using a full set of the original afuns, including the original affixes, such as *_Co*; later in this paper, we define LAS differently and add one more metric (Sect. 5.2.1).

4.3. Availability of Other Annotation

The dependency syntax annotation can take advantage of the availability of other manual annotation on the same data. Especially the higher, syntactic-semantic layer manual annotation with its tree structure can serve as a guidance because both layers are derived from the same syntactic theory developed in the framework of Functional Generative Description (Sgall et al., 1986). The main difference between the two syntactic layers lies in the fact that at the dependency syntax layer, every word of a sentence is represented by a node of its own and at the same time, no added nodes are allowed (cf. Sect. 3), while at the syntactic-semantic layer, only the content words are represented by a separate node and the function words such as prepositions, conjunctions, auxiliary or modal verbs do not have their own node: their meaning is captured by attributes of the nodes for the respective content words. Moreover, the cases of surface deletion are resolved by inserting new nodes into the tree structure.⁵

As an illustration, note the differences between the two layers in Fig. 1 depicting annotation of the sentence *Dort bude v prosinci vydražen při benefiční akci*. ‘The cake will be auctioned at a charity event in December’.⁶ Fig. 1 shows that at the dependency syntax layer, the auxiliary verb *bude* ‘will’ has its own node with label *AuxV*, while at the syntactic-semantic layer, its function is captured within the node for the predicate. It is similar for prepositions (*v* ‘in’, *při* ‘at’) that have their nodes (with label *AuxP*) at the dependency syntax layer, but not at the syntactic-semantic one. Otherwise, there is an added node for the so called general actor (with lemma *#Gen*) at the syntactic-semantic layer, which is not expressed in the surface shape of the sentence. The semantic labels such as *ACT* (Actor), *PAT* (Patient), *TWHEN* (When) and *RSTR* (Restriction) usually correspond to afuns *Sb*, *Obj*, *Adv* and *Attr* respectively at the dependency syntax layer. However, there are several exceptions such as in case of passive constructions, etc.

Annotators see both trees side by side during the annotation, so they can check their annotation against the other one.⁷ Although the principles of the annotation on both layers are different and the resulting tree structure may also be very dissimilar, we assume that the availability of (largely) semantic annotation during

⁵Detailed annotation rules can be found in the annotation guidelines (Mikulová et al., 2006; Mikulová, 2014).

⁶The TrEd concept of stylesheets allows to visually differentiate nodes, edges by color, shape, etc. and also visualize cross-layer relations. The trees in Fig. 1 are shown in the stylesheet available to annotators during the annotation. Under a lemma in the syntactic-semantic tree (on the right side), the function words belonging to the given node are displayed in orange.

⁷The annotators do not have to have a full knowledge of the other annotation structure and guidelines; it is sufficient they know the principles in order to be able to “read” the annotation as shown in TrEd.

the dependency syntax annotation speeds up the understanding of a sentence structure (especially in the case of long sentences), helps in ambiguous and complicated cases, both in solving the syntactic structure and in determining afuns; and finally increases the consistency of the annotated data.

5. Experiment

In this section, we describe the experiment, its design (Sect. 5 1), results (Sect. 5 2) and evaluation (Sect. 5 3). Input data, individual experimental annotations, and a complete and detailed overview of the measured results are freely available (as a supplemental material to this paper) via the LINDAT/CLARIAH-CZ repository (Mikulová et al., 2022).⁸

5.1. Design

In line with the research questions that we want to address and the annotators that we have available, we choose the following experiment design.

In the experiment, we examine four annotation tasks: annotation with no support (abbreviated *no_supp*), annotation with checking rules available (abbreviated *rules*; cf. Sect. 4.2), annotation with syntactic-semantic annotation on the same data available (abbreviated *annot*; cf. Sect. 4.3), and annotation with checking rules and syntactic-semantic annotation on the same data available (abbreviated *rul_annot*). For each task, the data was prepared in two modes: annotation on pre-parsed data (abbreviated *pre-parsed*; cf. Sect. 4.1) and annotation “from scratch” (abbreviated *from-scratch*). In the experiment, we thus examined eight different annotation set-ups (four tasks in two modes).

Task	no_supp		rules		annot		rul_annot	
Data	D1	D2	D3	D4	D5	D6	D7	D8
Pre-parsed	a1	a3	a1	a3	a1	a3	a1	a3
	a2	a4	a2	a4	a2	a4	a2	a4
From-scratch	a3	a1	a3	a1	a3	a1	a3	a1
	a4	a2	a4	a2	a4	a2	a4	a2

Table 1: Data distribution

We prepared eight datasets (*D1 - D8*) of 1,250 words, two sets for each task, i.e., a total of 2,500 words per task. We prepared each dataset in both modes (for pre-parsed and from-scratch annotation). The datasets were selected to be as similar as possible in terms of the number of sentences (each dataset contains approximately 60 sentences), number of nodes per sentence, depth of tree, etc. Each dataset was a coherent, continuous text. With this selection, we eliminated the data effect on the resulting values.

Each task was annotated by four annotators (abbreviated *a1*, *a2*, *a3*, *a4*) in stable (unchanged) pairs. The annotation pairs were composed so that they were experientially balanced. Each pair thus experiences all eight annotation set-ups and no annotator annotated any data twice (cf. Tab. 1 of data distribution). With

⁸<http://hdl.handle.net/11234/1-4647>

Task/Mode	Pre-parsed	From-scratch
UAS		
no_supp	96.5 ±0.46	96.5 ±0.48
rules	96.8 ±0.42 p=30.9%	95.9 ±0.50 p=79.2%
annot	97.5 ±0.48 p= 7.1%	97.0 ±0.48 p=24.5%
rul_annot	97.6 ±0.37 p= 3.5%	97.8 ±0.32 p= 1.5%
LAS		
no_supp	95.0 ±0.55	94.5 ±0.55
rules	95.3 ±0.49 p=30.6%	93.9 ±0.59 p=74.9%
annot	96.5 ±0.58 p= 3.3%	95.9 ±0.55 p= 4.7%
rul_annot	96.6 ±0.45 p= 1.3%	96.8 ±0.38 p< 0.1%
FULL		
no_supp	94.5 ±0.56	94.3 ±0.56
rules	95.1 ±0.50 p=23.6%	93.5 ±0.60 p=80.1%
annot	96.1 ±0.60 p= 3.5%	95.2 ±0.63 p=14.4%
rul_annot	96.4 ±0.47 p= 0.8%	96.6 ±0.40 p< 0.1%

Table 2: **Accuracy** of annotation in various set-ups; the standard deviations are estimated using bootstrap re-sampling with 1M samples, and *p* stands for a p-value that the given configuration is statistically significantly better than the corresponding *no_supp* configuration, estimated by a Monte Carlo permutation test with 1M samples (Fay and Follmann, 2002).

this setting, we eliminated the learning effect and reduced the impact of the experience and “style” of the annotators on the resulting (averaged) values.

For each dataset, a gold standard data annotation was also created. As follows from the data distribution (cf. Tab. 1), each dataset was annotated four times (twice in from-scratch mode and twice in pre-parsed mode). The annotations were then diff-ed and disagreements were resolved by the guideline setters.

5.2. Results

5.2.1. Measuring accuracy

We measured the accuracy of the annotation against the gold data. We measured the unlabeled and labeled attachment score for each dataset in each task for each mode, separately for each annotator. For each annotation set-up (for each task in the given mode) we obtained four values (cf. a1 to a4 in each cell of Tab. 1), from which we then calculated the average value.

We distinguish two labeled attachment scores. The first score, abbreviated LAS, measures accuracy on 25 basic labels (afuns), ignoring the three types of affixes. Afuns including affixes (increasing the number of labels to 200, cf. Sect. 3) are taken into account in the LAS score variant abbreviated as FULL.

All three scores are calculated in the same way: the number of agreements divided by the total number of edges. The agreement means the same head for a node for UAS, the same head and the same afun for LAS, and the same head and the same afun including affixes for FULL. The results are presented in Tab. 2.

5.2.2. Inter-annotator agreement

Inter-annotator agreement among annotators is a standard way to determine consistency in the process of annotation. We use a widely used measure, Cohen’s kappa, κ (Cohen, 1960) for measuring agreement between annotators. Although there has been criticism of the inadequacy of this method for structure annotation (Skjærholt, 2014), it provides a standardized way of estimating the agreement rate between annotators. Similarly to Sect. 5.2.1, we measured three different kappas: unlabeled, labeled, and full. But the details were quite different: for the traditional formula

$$\kappa = 1 - \frac{(1 - p_0)}{(1 - p_e)}$$

the values of p_0 (actual agreement) and p_e (probability of a random agreement) were calculated as follows:

- For unlabeled kappa, $p_0 = a_P/n$ where a_P is the number of nodes with the same head in both annotations, while n is the total number of nodes (as nodes are not added or removed, the number is always the same for both the annotators); $p_e = \bar{s}^{-1}$ where \bar{s} is the average sentence size.
- For labeled kappa, $p_0 = a_L/a_P$, i.e. we only consider the edges that are common to both the annotations and we check their labels (a_L is the number of agreements on afuns); $p_e = 1/25$ where 25 is the number of all possible labels (afuns).
- The full labeled kappa is similar: $p_0 = a_F/a_P$ and $p_e = 1/200$ (where a_F is the number of agreements on labels including affixes). Clearly, the probability of a random agreement is very low for such a high number of labels.

We measured the inter-annotator agreement for each pair of annotators. For each task in the given mode, we obtained two values (cf. pairs a1-a2 and a3-a4 in each cell of Tab. 1), from which we then calculated the average values. The results are presented in Tab. 3.

5.2.3. Annotation time

In Tab. 4, we report the total time taken in annotation for each annotation set-up of each annotator.

5.3. Evaluation

In this section, we judge the results of the experiment (presented above in Sect. 5.2) with respect to the research questions stated in the Sect. 1.

5.3.1. Does annotation quality remain acceptably high if pre-annotation is used?

To answer the main research question, we compare the values obtained for the pre-parsed and for from-scratch data for accuracy (cf. Tab. 2), inter-annotator agreement (cf. Tab. 3), and annotation time (cf. Tab. 4).

From Tab. 2, it is clear that the annotation accuracy remained, more or less, the same for both modes of annotation, with a fluctuation (as measured by the standard

Task/Mode	Pre-parsed	From-scratch
Unlabeled κ		
no_supp	0.96	0.95
rules	0.97	0.94
annot	0.97	0.96
rul_annot	0.97	0.97
Labeled κ		
no_supp	0.99	0.96
rules	0.98	0.97
annot	0.99	0.98
rul_annot	0.99	0.99
Full-labeled κ		
no_supp	0.98	0.96
rules	0.98	0.97
annot	0.98	0.97
rul_annot	0.98	0.99

Table 3: **Agreement rate** in various set-ups

Task/Mode		Pre-parsed	From-scratch
no_supp	a1	66	150
	a2	125	231
	a3	80	140
	a4	200	280
rules	a1	90	156
	a2	216	252
	a3	60	150
	a4	150	180
annot	a1	111	130
	a2	212	205
	a3	70	180
	a4	210	230
rul_annot	a1	80	150
	a2	200	245
	a3	50	160
	a4	180	155

Table 4: **Time** in various set-ups (in mins)

deviation) of around 0.5 percentage points. Cf. also Fig. 2 which depicts overall accuracy only on the annotation of no-support task.

To find out how much manual correction of automatic pre-annotation increases the quality, we also measured the accuracy of the parser used (i.e., the accuracy of the input data for tasks in ”pre-parsed” mode) against the gold data. An overview of the obtained values for parser, annotation on pre-parsed data (no-support task) and annotation on from-scratch data (no-support task) is shown in Fig. 2.

Lower parsing accuracy (compared to the values stated in Sect. 4.1) are due to the fact that we have changed the domain and also some annotation rules (cancellation of some afuns, addition of new ones, etc., even if only a small minority, occurrence-wise) compared to the previous dataset on which the parser has been trained. The graph shows that manual annotation (whether on pre-parsed data or not) produces higher quality data than the output from the parser.

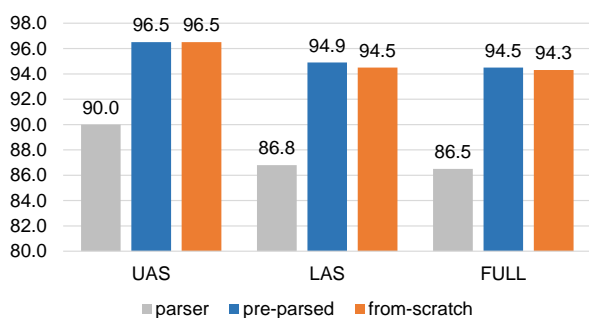


Figure 2: **Accuracy** of the annotation of no-support task in comparison with accuracy of the parser

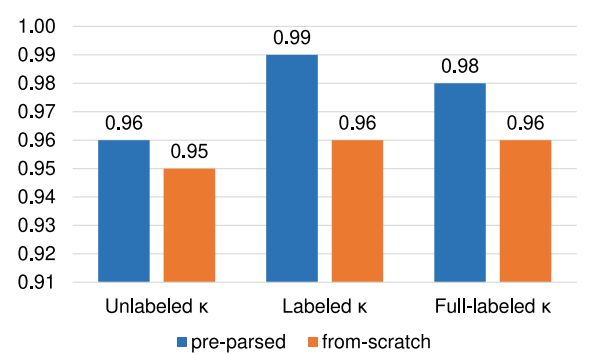


Figure 3: **Agreement rate** of no-support annotation

Fig. 3 depicts total agreement rate for annotation on the no-support task. From the Fig. 3, it is obvious that the pre-annotation contributes to the consistency of the annotation. This can be attributed to the fact that judgments of annotators are influenced by the output of automatic pre-annotation. However, as stated above, pre-annotation does not affect the quality of the result.

We can also observe that the variance in the annotation speed between individual annotators is large, from 50 to 200 minutes per task of pre-parsed data and from 130 to 280 minutes per task of from-scratch data (cf. Tab. 4). On average, the annotation on from-scratch data takes almost 1.7 times longer than annotation on pre-parsed data.

It follows from these results that the answer to the first research question is in the affirmative: the quality of the annotation is not hurt by the pre-parsing mode while the speed is up and inter-annotator agreement increases.

5.3.2. What effect do the used tools have on annotation quality and efficiency?

To answer the question how the quality and speed of annotation changes when using other support tools, we compare in detail the figures obtained for each task in Tables 2, 3, and 4.

As mentioned in the previous section, the annotation accuracy remained, more or less, the same for both modes. As can be seen from the figures in Tab. 2, we

can observe how much the addition of a support tool (*rules*, *annot*, *rul_annot*) has increased the quality of annotation. While the differences seem small, some of them are significant when compared to the *no_supp* task (see the *p*-values in Tab. 2). Detailed insight through the significance tests shows that especially the combination of the “online” checking rules and the possibility to easily visually check the other type of annotation (the syntactic-semantic one) and make changes under its guidance (line *rul_annot* in Tab. 2) did improve the quality of the annotation significantly for all of UAS, LAS, and FULL metrics, in some cases very convincingly (as demonstrated by a very small *p*-value). Even the use of the other annotated data alone (line *annot* in Tab. 2) improved some metrics in at least one mode significantly, even if less convincingly: LAS for both the pre-parsed and from-scratch mode, and FULL for the pre-parsed mode; UAS for the pre-parsed mode lies on the boundary of acceptable significance.

The apparently small influence when using the checking rules as the only support (line *rules* in Tab. 2) can be attributed to the fact that there were relatively only a few sets of checks available at the time of the experiment. The more interesting it is then to observe that the combination of the “online” checking rules and the availability of the other annotation (*rul_annot*) has increased the quality so convincingly.

Regarding the consistency of annotation, the figures in Tab. 3 show the total agreement rate for each task on pre-parsed and from scratch data. We can observe that the use of annotation at the syntactic-semantic layer as a support lead to higher consistency of the annotation, especially on from-scratch input data. If pre-annotation is used, the effect of other tools (including other annotation support) on the consistency of annotation is unclear.




Figure 4: **Time** of annotation

It is hard to draw any conclusions regarding the correlation of speed of the annotation and the use of the support tools, due to the high variation of the speed of annotation across the annotators. But perhaps this is in fact good news: the annotation time does not increase (too much) due to the use of checking rules, neither when annotators have the other annotation available to check against (cf. Tab. 4 and detailed graph in Fig. 4).

5.3.3. Which set-up is most useful for manual dependency syntax annotation?

The experiment confirms the usefulness of automatic pre-annotation for efficiency of annotation and quality of the result. It significantly increases the speed of annotation and does not reduce the quality of data.

The influence of the support tools is not as positive as expected, especially when using the checking rules alone. This contradicts our previous experience (e.g. Mikulová and Štěpánek (2010)). Perhaps it will help to develop more checking rules throughout the annotation project, once it starts in full. Regarding the use of annotation at the syntactic-semantic layer, the agreement rate figures measured on the from-scratch data show that the other annotation support has the potential to increase the consistency of the annotation, as well as the quality if used in conjunction with the checking rules. Therefore, and because these tools do not increase annotation time too much (cf. Fig. 4), we will use them in the annotation project proper (together with pre-parsing).




Figure 5: Extrapolated hours of annotation needed

A graph in Fig. 5 depicts the extrapolated time for annotation using just manual work (from scratch with no annotation support) vs. the time needed when using pre-parsed data (with and without the additional annotation support) as presented in this paper. We leave the calculation of financial savings to the readers, based on their usual hourly rates. Please note that any hourly savings of the primary annotation induce additional savings of technical support, supervision etc., which is however harder to quantify.

In our case, for the PDT-C corpora annotation with a total of about 2 million tokens, the primary saving will be approx. 2,000 hours (5,400–3,400) per single annotation pass (~1.3 person-years, or at least about \$21,000) if we compare the annotation with no support and the annotation with all available support; with the usual double annotation, it is then also double the savings in both time and cost.

6. Conclusions and Future Work

In the paper, we presented pre-annotation bias experiment for a mid-level annotation complexity task—dependency syntax annotation within the Prague Dependency Treebank - Consolidated 2.0 project which is now ongoing. We judged efficiency and quality of annotation performed under various experimental set-ups: annotation with no support (“from-scratch”) and annotation supported by the use of automatic pre-annotation (using a high-accuracy parser), availability of other annotation on the same data, and “online” checking rules (implemented in such a way that the annotator can check the manual annotation immediately).

The experiment confirmed that pre-annotation using UDPipe 2 is an efficient tool for manual syntactic annotation, which increases the speed and consistency of the resulting annotation without reducing its quality. Thus we can conclude it is highly recommended to use automatic pre-annotation in syntactic annotation projects, even if the other tools are not available.

The influence of the other support tools (online annotation checking rules and the availability of a related, even if largely not directly comparable syntactic-semantic annotation on the same data) on efficiency and quality of annotation was not as positive as expected. However, when used together they significantly increase the quality of the annotation, in both the pre-parsed and from-scratch modes. As these tools do not increase annotation time (much), we will use them in the upcoming large-scale PDT-C dependency syntax annotation process.

We plan to repeat the experiment at half-time of the project to re-evaluate the influence of the chosen annotation set-up on the quality and efficiency of the annotation.

7. Acknowledgements

We thank the annotators for their commitment in the implementation of the experiment.

The research and language resource work reported in the paper has been supported by the LINDAT/CLARIAH-CZ project funded by Ministry of Education, Youth and Sports of the Czech Republic (project LM2018101) and by the EXPRO project LUSyD, funded by the Grant Agency of the Czech Republic as project No. GX20-16819X.

8. Bibliographical References

- Artstein, R. and Poesio, M. (2008). Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.
- Braggaar, A. and van der Goot, R. (2021). Creating a Universal Dependencies Treebank of Spoken Frisian-Dutch Code-switched Data.
- Chiou, F.-D., Chiang, D., and Palmer, M. (2001). Facilitating treebank annotation using a statistical parser. In *Proceedings of the First International*

- Conference on Human Language Technology Research*.
- Chou, W.-C., Tsai, R. T.-H., Su, Y.-S., Ku, W., Sung, T.-Y., and Hsu, W.-L. (2006). A semi-automatic method for annotating a biomedical Proposition Bank. In *Proceedings of the Workshop on Frontiers in Linguistically Annotated Corpora 2006*, pages 5–12, Sydney, Australia, July. Association for Computational Linguistics.
- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Fay, M. P. and Follmann, D. A. (2002). Designing Monte Carlo Implementations of Permutation or Bootstrap Hypothesis Tests. *The American Statistician*, 56(1):63–70.
- Fort, K. and Sagot, B. (2010). Influence of pre-annotation on POS-tagged corpus development. In *Proceedings of the Fourth Linguistic Annotation Workshop*, pages 56–63, Uppsala, Sweden, July. Association for Computational Linguistics.
- Grouin, C., Lavergne, T., and Névéol, A. (2014). Optimizing annotation efforts to build reliable annotated corpora for training statistical models. In *Proceedings of LAW VIII-The 8th Linguistic Annotation Workshop*, pages 54–58.
- Gupta, M., Yadav, V., Husain, S., and Sharma, D. M. (2010). Partial parsing as a method to expedite dependency annotation of a Hindi treebank. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta, May. European Language Resources Association (ELRA).
- Hajič, J., Panevová, J., Buráňová, E., Urešová, Z., and Bémová, A. (1999). Annotation at analytical level. Technical report, Institute of Formal and Applied Linguistics, Charles University, Prague, Czech Republic.
- Hajič, J., Bejček, E., Hlaváčová, J., Mikulová, M., Straka, M., Štěpánek, J., and Štěpánková, B. (2020). Prague Dependency Treebank - Consolidated 1.0. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5208–5218, Marseille, France, May. European Language Resources Association.
- Hajič, J. (1998). Building a syntactically annotated corpus: The Prague Dependency Treebank. In *Issues of Valency and Meaning. Studies in Honour of Jarmila Panevová (ed. Eva Hajičová)*. Karolinum, Charles University Press, Prague, ISBN 80-7184-601-5.
- Marcus, M., Santorini, B., and Marcinkiewicz, M. A. (1993). Building a large annotated corpus of English: The Penn Treebank.
- Mikulová, M. and Štěpánek, J. (2010). Ways of evaluation of the annotators in building the Prague Czech-English Dependency Treebank. In *Proceedings of the 7th International Conference on Language Resources and Evaluation (LREC 2010)*, pages 1836–1839, Valletta, Malta. European Language Resources Association.
- Mikulová, M., Bémová, A., Hajič, J., Hajičová, E., Havelka, J., Kolářová, V., Kučová, L., Lopatková, M., Pajas, P., Panevová, J., Razímová, M., Sgall, P., Štěpánek, J., Urešová, Z., Veselá, K., and Žabokrtský, Z. (2006). Annotation on the tectogrammatical level in the Prague Dependency Treebank. Annotation manual. Technical Report 30, Institute of Formal and Applied Linguistics, Charles University, Prague, Czech Republic.
- Mikulová, M. (2014). Annotation on the tectogrammatical level. Additions to annotation manual (with respect to PDTSC and PCEDT). Technical Report TR-2013-52, Institute of Formal and Applied Linguistics, Charles University, Prague, Czech Republic.
- Nguyen, K.-H. (2018). BKTtreebank: Building a Vietnamese Dependency Treebank. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May. European Language Resources Association (ELRA).
- Pajas, P. and Štěpánek, J. (2008). Recent advances in a feature-rich framework for treebank annotation. In Donia Scott et al., editors, *The 22nd International Conference on Computational Linguistics - Proceedings of the Conference*, volume 2, pages 673–680, Manchester, UK. The Coling 2008 Organizing Committee.
- Rehbein, I., Ruppenhofer, J., and Sporleder, C. (2009). Assessing the benefits of partial automatic pre-labeling for frame-semantic annotation. In *Proceedings of the Third Linguistic Annotation Workshop (LAW III)*, pages 19–26, Suntec, Singapore, August. Association for Computational Linguistics.
- Rosset, S., Grouin, C., Lavergne, T., Ben Jannet, M., Leixa, J., Galibert, O., and Zweigenbaum, P. (2013). Automatic named entity pre-annotation for out-of-domain human annotation. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 168–177, Sofia, Bulgaria, August. Association for Computational Linguistics.
- Sgall, P., Hajičová, E., and Panevová, J. (1986). *The Meaning of the Sentence and Its Semantic and Pragmatic Aspects*. Academia/Reidel Publishing Company, Prague/Dordrecht.
- Skjærholt, A. (2014). A chance-corrected measure of inter-annotator agreement for syntax. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 934–944, Baltimore, Maryland. Association for Computational Linguistics.
- Šmilauer, V. (1947). Novočeská skladba (Syntax of Modern Czech). *Prague: Academia*.
- Straka, M., Náplava, J., Straková, J., and Samuel, D. (2021). Robeczech: Czech RoBERTa, a monolin-

gual contextualized language representation model. In Kamil Ekštejn, et al., editors, *Text, Speech, and Dialogue*, pages 197–209, Cham. Springer International Publishing.

- Voutilainen, A. and Purtonen, T. (2011). A double-blind experiment on interannotator agreement: The case of dependency syntax and Finnish. In *Proceedings of the 18th Nordic Conference of Computational Linguistics (NODALIDA 2011)*, pages 319–322.
- Yang, A. and Li, S. (2018). SciDTB: Discourse dependency TreeBank for scientific abstracts. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 444–449, Melbourne, Australia, July. Association for Computational Linguistics.

9. Language Resource References

- Hajič, Jan and Bejček, Eduard and Bémová, Alevtina and Buráňová, Eva and Fučíková, Eva and Hajičová, Eva and Havelka, Jiří and Hlaváčová, Jaroslava and Homola, Petr and Ircing, Pavel and Kárník, Jiří and Kettnerová, Václava and Klyueva, Natalia and Kolářová, Veronika and Kučová, Lucie and Lopatková, Markéta and Mareček, David and Mikulová, Marie and Mírovský, Jiří and Nedoluzhko, Anna and Novák, Michal and Pajas, Petr and Panevová, Jarmila and Peterek, Nino and Poláková, Lucie and Popel, Martin and Popelka, Jan and Romportl, Jan and Rysová, Magdaléna and Semecký, Jiří and Sgall, Petr and Spoustová, Johanka and Straka, Milan and Straňák, Pavel and Synková, Pavlína and Ševčíková, Magda and Šindlerová, Jana and Štěpánek, Jan and Štěpánková, Barbora and Toman, Josef and Urešová, Zdeňka and Vidová Hladká, Barbora and Zeman, Daniel and Zikánová, Šárka and Žabokrtský, Zdeněk. (2020). *Prague Dependency Treebank - Consolidated 1.0 (PDT-C 1.0)*. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University; <http://hdl.handle.net/11234/1-3185>.
- Mitchell P. Marcus and Beatrice Santorini and Mary Ann Marcinkiewicz and Ann Taylor. (1999). *Penn Treebank-3 (LDC99T42)*. Linguistic Data Consortium, Philadelphia, PA, USA, ISLRN 141-282-691-413-2.
- Mikulová, Marie and Straka, Milan and Štěpánek, Jan and Štěpánková, Barbora and Hajič, Jan. (2022). *Quality and Efficiency of Manual Annotation: Data from the Pre-annotation Bias Experiment (part of the PDT-C 2.0 project)*. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University; <http://hdl.handle.net/11234/1-4647>.