LETTER
Correcting Syntactic Annotation Errors Using a Synchronous Tree Substitution Grammar

Yoshihide KATO[†a] *and* **Shigeki MATSUBARA**[†], *Members*

SUMMARY   This paper proposes a method of correcting annotation errors in a treebank. By using a synchronous grammar, the method transforms parse trees containing annotation errors into the ones whose errors are corrected. The synchronous grammar is automatically induced from the treebank. We report an experimental result of applying our method to the Penn Treebank. The result demonstrates that our method corrects syntactic annotation errors with high precision.
key words: error correction, synchronous tree substitution grammar, syntactic annotation, treebank, parallel corpus

1. Introduction

Annotated corpora play an important role in the fields such as theoretical linguistic researches or the development of NLP systems. However, they often contain annotation errors which are caused by a manual or semi-manual mark-up process. These errors are problematic for corpus-based researches.

To solve this problem, several error detection and correction methods have been proposed so far [1]–[8]. These methods detect corpus positions which are marked up incorrectly, and find the correct labels (e.g. pos-tags) for those positions. However, the methods cannot correct errors in structural annotation. This means that they are insufficient to correct annotation errors in a treebank.

This paper proposes a method of correcting errors in structural annotation. Our method is based on a synchronous grammar formalism, called *synchronous tree substitution grammar* (STSG) [9], which defines a tree-to-tree transformation. By using an STSG, our method transforms parse trees containing errors into the ones whose errors are corrected. The grammar is automatically induced from the treebank. To select STSG rules which are useful for error correction, we define a score function based on the occurrence frequencies of the rules. An experimental result shows that the selected rules archive high precision.

This paper is organized as follows: Section 2 gives an overview of previous work. Section 3 explains our method of correcting errors in a treebank. Section 4 reports an experimental result using the Penn Treebank.

2. Previous Work

This section summarizes previous methods for correcting er-

Fig. 1   An example of an annotation error in a treebank.

rors in corpus annotation and discusses their problem.

Some research addresses the detection of errors in pos-annotation [2], [3], syntactic annotation [4], [5], [7], and dependency annotation [8]. These methods only detect corpus positions where errors occur. It is unclear how we can correct the errors.

Several methods can correct annotation errors [1], [6]. These methods are to correct tag-annotation errors, that is, they simply suggest a candidate tag for each position where an error is detected. The methods cannot correct syntactic annotation errors, because syntactic annotation is structural. There is no approach to correct structural annotation errors.

To clarify the problem, let us consider an example. Figure 1 depicts two parse trees annotated according to the Penn Treebank annotation [*]. The parse tree (a) contains errors and the parse tree (b) is the corrected version. In the parse tree (a), the positions of the two subtrees (, ,) are erroneous. To correct the errors, we need to move the subtrees to the positions which are directly dominated by the node PRN. This example demonstrates that we need a framework of transforming tree structures to correct structural annotation er-

*[*]0 and *T* are null elements.*

rors.

3. Correcting Errors by Using Synchronous Grammar

To solve the problem described in Sect. 2, this section proposes a method of correcting structural annotation errors by using a synchronous tree substitution grammar (STSG) [9]. An STSG defines a tree-to-tree transformation. Our method induces an STSG which transforms parse trees containing errors into the ones whose errors are corrected.

3.1 Synchronous Tree Substitution Grammar

First of all, we describe the STSG formalism. An STSG defines a set of tree pairs. An STSG can be treated as a tree transducer which takes a tree as input and produces a tree as output. Each grammar rule consists of the following elements:

- a pair of trees called *elementary trees*
- a one-to-one alignment between nodes in the elementary trees

For a tree pair $\langle t, t' \rangle$, the tree t and t' are called *source* and *target*, respectively. The non-terminal leaves of elementary trees are called *frontier nodes*. There exists a one-to-one alignment between the frontier nodes in t and t'. The rule means that the structure which matches the source elementary tree is transformed into the structure which is represented by the target elementary tree. Figure 2 shows an example of an STSG rule. The subscripts indicate the alignment. This rule can correct the errors in the parse tree (a) depicted in Fig. 1.

An STSG derives tree pairs. Any derivation process starts with the pair of nodes labeled with special symbols called *start symbols*. A derivation proceeds in the following steps:

1. Choose a pair of frontier nodes $\langle \eta, \eta' \rangle$ for which there exists an alignment.
2. Choose a rule $\langle t, t' \rangle$ s.t. $label(\eta) = root(t)$ and $label(\eta') = root(t')$ where $label(\eta)$ is the label of η and $root(t)$ is the root label of t.
3. Substitute t and t' into η and η', respectively.

Figure 3 shows a derivation process in an STSG.

In the rest of the paper, we focus on the rules in which the source elementary tree is not identical to its target, since such identical rules cannot contribute to error correction.

3.2 Inducing an STSG for Error Correction

This section describes a method of inducing an STSG for error correction. The basic idea of our method is similar to the method presented by Dickinson and Meurers [4]. Their method detects errors by seeking word sequences satisfying the following conditions:

- The word sequence occurs more than once in the corpus.
- Different syntactic labels are assigned to the occurrences of the word sequence.

Unlike their method, our method seeks word sequences whose occurrences have different partial parse trees. We call a collection of these word sequences with partial parse trees *pseudo parallel corpus*. Moreover, our method extracts STSG rules which transform the one partial tree into the other.

3.2.1 Constructing a Pseudo Parallel Corpus

Our method firstly constructs a pseudo parallel corpus which represents a correspondence between parse trees containing errors and the ones whose errors are corrected. The procedure is as follows: Let T be the set of the parse trees occurring in the corpus. We write $Sub(\sigma)$ for the set which consists of the partial parse trees included in the parse tree σ. A pseudo parallel corpus $Para(T)$ is constructed as follows:
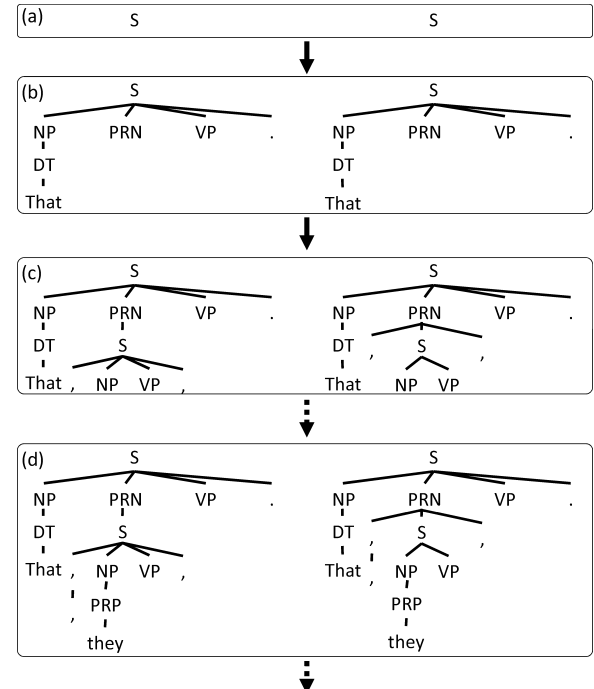

Fig. 3 A derivation process of tree pairs in an STSG.


Fig. 2 An example of an STSG rule.

Fig. 4 An example of a partial parse tree pair in a pseudo parallel corpus.

**Fig. 5** Another example of a parse tree containing a word sequence ", they say ,".

$$Para(T) = \{\langle \tau, \tau' \rangle \mid \tau, \tau' \in \bigcup_{\sigma \in T} Sub(\sigma)$$
$$\wedge \tau \neq \tau'$$
$$\wedge yield(\tau) = yield(\tau')$$
$$\wedge root(\tau) = root(\tau')\}$$

where $yield(\tau)$ is the word sequence dominated by τ.

Let us consider an example. If the parse trees depicted in Fig. 1 exist in the treebank T, the pair of partial parse trees depicted in Fig. 4 is an element of $Para(T)$. We also obtain this pair in the case where there exists not the parse tree (b) depicted in Fig. 1 but the parse tree depicted in Fig. 5, which contains the word sequence ", they say ,".

3.2.2 Inducing a Grammar from a Pseudo Parallel Corpus

Our method induces an STSG from the pseudo parallel corpus according to the method proposed by Cohn and Lapata [10]. Cohn and Lapata's method can induce an STSG which represents a correspondence in a parallel corpus. Their method firstly determine an alignment of nodes between pairs of trees in the parallel corpus and extracts STSG rules according to the alignments.

For partial parse trees τ and τ', we define a node alignment $C(\tau, \tau')$ as follows:

$$C(\tau, \tau') = \{\langle \eta, \eta' \rangle \mid \eta \in Node(\tau)$$
$$\wedge \eta' \in Node(\tau')$$
$$\wedge \eta \text{ is not the root of } \tau$$
$$\wedge \eta' \text{ is not the root of } \tau'$$
$$\wedge label(\eta) = label(\eta')$$
$$\wedge yield(\eta) = yield(\eta')\}$$

where $Node(\tau)$ is the set of the nodes in τ, and $yield(\eta)$ is the word sequence dominated by η. Figure 4 shows an example of a node alignment. The subscripts indicate the alignment.

An STSG rule is extracted by deleting nodes in a partial parse tree pair $\langle \tau, \tau' \rangle \in Para(T)$. The procedure is as follows:

- For each $\langle \eta, \eta' \rangle \in C(\tau, \tau')$, delete the descendants of η and η'.

For example, the rule shown in Fig. 2 is extracted from the pair shown in Fig. 4.

3.3 Rule Selection

Some rules extracted by the procedure in Sect. 3.2 are not useful for error correction, since the pseudo parallel corpus contains tree pairs whose source tree is correct or whose target tree is incorrect. The rules which are extracted from such pairs can be harmful. To select rules which are useful for error correction, we define a score function which is based on the occurrence frequencies of elementary trees in the treebank. The score function is defined as follows:

$$Score(\langle t, t' \rangle) = \frac{f(t')}{f(t) + f(t')}$$

where $f(\cdot)$ is the occurrence frequency in the treebank. The score function ranges from 0 to 1. We assume that the occurrence frequency of an elementary tree matching incorrect parse trees is very low. According to this assumption, the score function $Score(\langle t, t' \rangle)$ is high when the source elementary tree t matches incorrect parse trees and the target elementary tree t' matches correct parse trees. Therefore, STSG rules with high scores are regarded to be useful for error correction.

4. An Experiment

To evaluate the effectiveness of our method, we conducted an experiment using the Penn Treebank [11].

We used 49208 sentences in Wall Street Journal sections. We induced STSG rules by applying our method to the corpus. We obtained 8776 rules. We measured the precision of the rules. The precision is defined as follows:

$$precision =$$
$$\frac{\text{\# of the positions where an error is corrected}}{\text{\# of the positions to which some rule is applied}}$$

Since it is time-consuming and expensive to evaluate all of the rules, we only evaluated the first 100 rules which are ordered by the score function described in Sect. 3.3. These rules were applied at 331 positions. The precision of the rules is 71.9%. For each rule, we measured the precision of it. 70 rules achieved 100 % precision. These results demonstrate that our method can correct syntactic annotation errors with high precision. Moreover, 30 rules of the 70 rules transformed bracketed structures. This fact shows that the treebank contains structural errors which cannot be dealt

Fig. 6 Examples of error correction rules induced from the Penn Treebank.

with by the previous methods.

Figure 6 depicts examples of error correction rules induced in this experiment. Rule (1) and (2) are rules which transform bracketed structures. Rule (3) simply replaces a node label. Rule (1) corrects an incorrect position of a prepositional phrase (PP) attachment. Rule (2) deletes a useless node NP in a subject position. Rule (3) replaces a node label NP with the correct label PP.

5. Conclusion

This paper proposes a method of correcting errors in a treebank by using a synchronous tree substitution grammar. Our method constructs a pseudo parallel corpus from the treebank and extracts STSG rules from the parallel corpus. The experimental result demonstrates that we can obtain error correction rules with high precision.

In future work, we will explore a method of increasing the recall of error correction by constructing a wide-coverage STSG.

Acknowledgments

References

[1] E. Eskin, "Detecting errors within a corpus using anomaly detection," Proc. 1st North American Chapter of the Association for Computational Linguistics Conference, pp.148–153, 2000.

[2] T. Nakagawa and Y. Matsumoto, "Detecting errors in corpora using support vector machines," Proc. 19th Internatinal Conference on Computatinal Linguistics, pp.709–715, Taipei, Taiwan, Aug. 2002.

[3] M. Dickinson and D. Meurers, "Detecting errors in part-of-speech annotation," Proc. 10th Conference of the European Chapter of the Association for Computational Linguistics, pp.107–114, 2003.

[4] M. Dickinson and D. Meurers, "Detecting inconsistencies in treebanks," Proc. 2nd Workshop on Treebanks and Linguistic Theories, 2003.

[5] T. Ule and K. Simov, "Unexpected productions may well be errors," Proc. 4th International Conference on Language Resources and Evaluation, pp.1795–1798, 2004.

[6] M. Murata, M. Utiyama, K. Uchimoto, H. Isahara, and Q. Ma, "Correction of errors in a verb modality corpus for machine translation with a machine-learning method," ACM Trans. Asian Language Information Processing (TALIP), vol.4, no.1, pp.18–37, 2005.

[7] M. Dickinson and W.D. Meurers, "Prune diseased branches to get healthy trees! how to find erroneous local trees in a treebank and why it matters," Proc. 4th Workshop on Treebanks and Linguistic Theories, Barcelona, Spain, 2005.

[8] A. Boyd, M. Dickinson, and D. Meurers, "On detecting errors in dependency treebanks," Research on Language and Computation, vol.6, no.2, pp.113–137, 2008.

[9] J. Eisner, "Learning non-isomorphic tree mappings for machine translation," Proc. 41st Annual Meeting of the Association for Computational Linguistics (ACL), Companion Volume, pp.205–208, Sapporo, July 2003.

[10] T. Cohn and M. Lapata, "Sentence compression as tree transduction," J. Artificial Intelligence Research, vol.34, no.1, pp.637–674, 2009.

[11] M.P. Marcus, B. Santorini, and M.A. Marcinkiewicz, "Building a large annotated corpus of English: the Penn Treebank," Computational Linguistics, vol.19, no.2, pp.310–330, 1993.