Automatic Detection and Correction of Syntax Errors in Tutoring Systems for Language Instruction *

Henk Harkema
harkema@ucla.edu

Abstract

This paper presents a method for the detection and correction of syntactic errors of competence, as produced by language learners using a tutoring system for second language instruction. The method is inspired by the realization that the correction of a syntactic error essentially is the translation of an incorrect sentence into a correct sentence. The method is based on explicit grammars for correct as well as incorrect sentences. These grammars are tuned to one another, so as to allow for a systematic mapping of structures describing incorrect sentences onto structures describing correct sentences.

1 Introduction

The acquisition of a second language is a complex process which involves learning about the various aspects connected to the successful use of the language, such as the lexicon, the syntax and the communicative aspects of language. With regard to the syntax of the language, it is assumed that the learner entertains a set of hypotheses about the structure of the language. Some of the hypotheses may be incorrect, causing the learner to produce erroneous sentences. For the learner to achieve syntactic accuracy in the second language, these wrong hypotheses need to be amended. One way of doing this, is to offer the learner the correct variants of his erroneous sentences. When language learning takes place in the context of a computer-based tutoring system for second language instruction, the system should include a procedure for providing these corrections.

In this paper, a particular method for the detection and correction of syntactic errors is presented. Besides this introduction, this paper has three parts. The following section is devoted to a brief discussion of aspects of

*The work reported on in this paper was done while the author was attached to the Institute for Perception Research, Eindhoven, the Netherlands.

tutoring systems and second language learning which are relevant to error correction. The method itself is then presented in the third section. The final section contains some concluding remarks.

2 Second Language Learning

This section offers a brief overview of the general architecture of tutoring systems for second language instruction, a concise discussion of the role of error correction in second language learning and a characterization of the kinds of errors that the method is able to deal with.

2.1 Tutoring Systems

A tutoring system for second language instruction is a computer system designed for teaching a foreign language to its users. A tutoring system consists of several components, embodying the various forms of knowledge essential to teaching a language: tutoring strategies, learning goals and course material. These components interact with one another to deliver instruction which best fits the needs of the student using the system (Hoe et al., 1995; Swarz and Yazdani, 1992; Wenger, 1987). The tutoring system must be able to respond adequately to all input from the user, either correct or incorrect. To this goal, the system also includes a parser, for analyzing sentences produced by the user, and a generator, for generating sentences, e.g. corrections of erroneous sentences.

Generally, instruction in tutoring systems takes the form of a sequence of lessons and exercises. In the lessons, the course material is presented to the student. The exercises provide the student with an opportunity to put his newly acquired knowledge of language to use. Also, the exercises enable the system to monitor the progress of the student. A well-designed tutoring system is adaptive: it takes into account the knowledge level of the student using the system to generate exercises which are specifically tailored to the needs of the student. To make an exercise efficient and attractive for the student, it should have a rather loose format, so as not to limit his linguistic creativity. This quality of exercises, together with adaptivity, result in exercises for which it is not possible to know beforehand what the correct answers will be. Therefore it is necessary that the correctness of an answer is checked afterwards by performing a grammatical analysis of the answer.

2.2 Role of Error Correction

According to current cognitive approaches to second language learning, the learner is involved in a process of actively constructing hypotheses about the structure of the language, by generalizing over the linguistic input received.

New input might confirm the hypotheses held by the learner, or it might urge the learner to adapt some hypothesis or dismiss it altogether. In this view, errors are inevitable during learning and their occurrence is a positive sign that the learner is busy testing hypotheses. Moreover, errors point to gaps in the linguistic knowledge of the learner, to which further instruction should be directed. The complex of hypotheses entertained by the learner is a language in its own right, called interlanguage (Selinker, 1983).

For first language acquisition it is generally believed that error correction is of minor importance; a first language is learned from positive evidence only. In the acquisition of a second language however, negative evidence seems to play a significant role (Carroll and Swain, 1993; Dekeyser, 1993; White, 1991). Some argue that second language acquisition differs from first language acquisition in that the innate linguistic knowledge is no longer available to the learner (Clahsen and Muysken, 1986). Others hold the view that second language learners still have access to innate linguistic knowledge, but that certain learning principles which are used for exploiting this knowledge, do not operate as effectively as in first language acquisition. As a result of this, the second language learner might adopt hypotheses about the structure of the language which in fact are consistent with the input and yet are incorrect (White, 1989). One way of inciting the learner to reject such hypotheses is explicit correction of his errors.

It is important to realize that correction takes three steps. First of all, the error must be detected. Detection of an error is done by comparing a sentence with general norms for grammaticality. If the sentence deviates from these norms, it contains an error. The next step is diagnosing the error, i.e. determining the particular way in which the sentence is deviant. For this, one needs to be aware of what the student intended to express with his incorrect sentence and what hypotheses he holds about the meaning of the words and the structure of the sentences in the language. With this knowledge, the incorrect sentence can then be mapped onto another sentence which in a syntactically correct way expresses what the student intended to express with his incorrect sentence—the former sentence is a correction of the latter sentence. The student's intentions thus ultimately determine the correction for an incorrect sentence.

2.3 Syntactic Errors of Competence

The sentences produced by a language learner reflect his linguistic competence as well as his performance. Competence errors signify a lack of linguistic knowledge or a linguistic misconception. The instruction system should respond to these errors, so as to allow the user to adapt his linguistic competence in the direction of that of a native speaker. Errors of performance are unrelated to linguistic knowledge. Therefore, correction of these errors is not crucial to language learning.

The method focuses on the detection and correction of syntactic errors of competence within sentences. Syntactic errors are related to the presence, the order and the relationships between the constituents in a sentence. Syntactic errors are subdivided into structural errors and morpho-syntactic errors. A sentence whose constituent structure is not in accordance with the rules of the language, suffers from a structural error. Structural errors include missing constituents, illicitly added constituents and incorrectly ordered constituents[1]:

(1) a. *Was a riot last night.
(There was a riot last night.)

b. *He worked until fell over.
(He worked until he fell over.)

c. *Escaped the professor from prison.
(The professor escaped from prison.)

d. *He says that he no money has.
(He says that he has no money.)

e. *Do I must take my entrance exam?
(Must I take my entrance exam?)

Morpho-syntactic errors make up the other subclass of syntactic errors. These errors are the result of misapplications of syntactic rules which affect the morphology of a word, such as the rule for agreement in number and person between a finite verb and its subject. In the case of a morpho-syntactic error, the constituent structure of the sentence is all right, but the relations between the constituents, as expressed in their form, are not as required by the grammar. Most errors concerning agreement, tense or case requirements within a sentence are instances of morpho-syntactic errors. Some examples are:

(2) a. *He walk in the park.
(He walks in the park.)

b. *Why did he went?
(Why did he go?)

c. *They should left.
(They should leave.)

The method for detection and correction presented in this paper employs a parser, drawing on a lexicon and elementary grammars, which represent knowledge about respectively the words of the language and the structure

[1] The examples in this section are taken from Burt and Kiparsky (1972).

of the sentences in the language. Any error whose detection and correction requires knowledge or machinery beyond a lexicon, a grammar and a parser, unfortunately must go untreated. Among the errors not covered by the method are spelling errors, semantic errors, pragmatic errors and discourse errors (see (Harkema, forthcoming) for details).

3 Detection and Correction

This section is devoted to the description of a method for the detection and correction of syntactic errors of competence. The method is based on explicit grammars for correct as well as incorrect sentences. These grammars are tuned to one another, so as to allow for a systematic mapping of structures describing incorrect sentences onto structures describing correct sentences.

Most of the existing methods for dealing with incorrect input in tutoring systems for second language instruction focus on diagnosing errors and providing the student with an explanation of what is wrong with his sentence (e.g. (Catt and Hirst, 1990; Garman et al., 1993)). Finding a correction for an incorrect sentence does not only require diagnosing the error, but also establishing a relationship between an incorrect sentence and a correct sentence. Other methods which do provide corrections (e.g. (Labrie and Singh, 1991)), are often used in very restricted exercises or only treat morpho-syntactic errors, which simplifies the task of finding a correction (e.g. (Labrie and Singh, 1991)). The method proposed in this paper can correct any sentence as long as it is equipped with the appropriate grammars.

3.1 Strict and Tolerant Grammars

A student using an tutoring system for second language instruction formulates his sentences in response to the exercises presented to him by the system. Usually, these exercises deal with a limited set of syntactic constructions of the language to be learned, and hence with only a fragment of the entire second language. For each language fragment covered in an exercise, two types of grammar are written: a strict grammar and a tolerant grammar.

A strict grammar describes the correct sentences of the language fragment in question. All sentences falling outside the strict grammar will be considered by the system to be incorrect. It is therefore essential that the language fragment is such that its correct sentences can be described in a grammar in an exhaustive way. A tolerant grammar describes the possible errors that a language learner might make with respect to the language fragment. A tolerant grammar thus contains rules characterizing incorrect syntactic constructions. An error which is covered by a tolerant grammar

permits correction. Anticipating errors by explicitly coding them in a tolerant grammar makes sense, because syntactic errors of competence are to a large extent systematic and predictable (Catt and Hirst, 1990; Littlewood, 1984). While learning a language, the student employs learning strategies which give rise to specific learner errors. Two such learning strategies are transfer and generalization.

The first language of a second language learner is a rich source of hypotheses about the second language. Transfer occurs when a student takes a rule from the first language and applies it to the second language. In the event that this rule does not apply to the second language, because the second language differs from the first language with respect to this rule, the result will be a transfer error. For example, in Dutch subordinate clauses the verb turns up at the end of the clause, as in (3):

(3) Ik denk dat de man de jongen zag.

A Dutch learner of English who assumes that this word order is also valid for English, is prone to produce the following erroneous sentence:

(4) *I think that the man the boy saw.

Another learning strategy is generalization. This strategy leads to errors of overgeneralization: the student takes a rule from the second language and applies it in a situation in which the rule is not applicable, either because there is another rule for this particular situation, or the situation is an exception to the rule. The Dutch language, for example, does not have an equivalent of the English progressive tense. Consequently, in teaching English to Dutch students, much attention is paid to the rule for forming progressives. As a result, these students tend to use this rule inappropriately, producing sentences like:

(5) *We are not knowing the rules.

3.2 Outline of the Method

Detection and correction of syntactic errors by means of a strict and a tolerant grammar proceeds along the following lines. First, the sentence produced by the student is parsed according to the strict grammar. If this is successful, the sentence is correct and no correction has to be generated (see Figure 1).

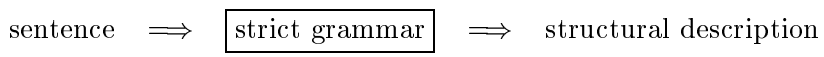

Figure 1: Dealing with a correct sentence.

If no parse can be found, an error is detected. In this case, the sentence is parsed with the tolerant grammar. If successful, this yields a structural

description of the sentence. Since the sentence is erroneous, the description will contain references to rules from the tolerant grammar describing incorrect syntactic constructions. Next, this structure is mapped onto a structure containing only correct rules from the strict grammar. This structure is then used to generate a sentence which is the correction of the erroneous sentence by the student (see Figure 2).

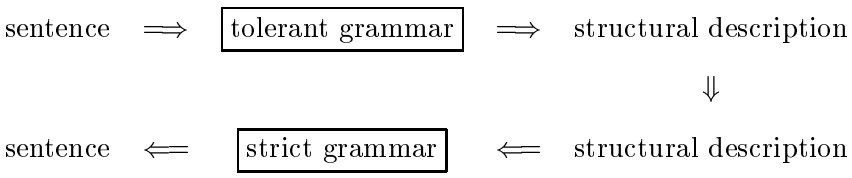

Figure 2: Correcting an incorrect sentence.

If, however, the erroneous sentence cannot be parsed with the tolerant grammar either, the detected error cannot be corrected (see Figure 3).

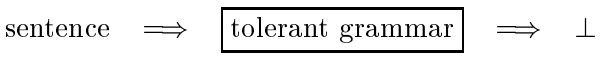

Figure 3: Detecting an error which cannot be corrected.

All sentences which are not covered by the strict grammar are detectably incorrect; of these incorrect sentences, only those which are described by the tolerant grammar are correctable.

3.3 Compositionality of Correction

The heart of the matter is the mapping of a sentence described by the tolerant grammar onto a sentence described by the strict grammar. This mapping is done according to a principle borrowed from the machine translation system Rosetta (Rosetta, 1994). In this system, a sentence is connected to its translation by the principle of compositionality of translation. This principle states that two expressions are translations of one another if they are built up from parts which are translations of one another, by means of rules which are translations of one another.

Correction can be thought of as uni-directional translation. Applied to correction, the notion of compositionality yields the principle of compositionality of syntactic correction: a sentence is a correction of some syntactically incorrect sentence, if the correct sentence is constructed from parts which are corrections of the parts of the incorrect sentence, using rules which are corrections of the rules used for the construction of the incorrect sentence. (To be exact, the correction of a correct part or rule, is understood to be the part or rule itself.) This principle summarizes the approach to correction taken in this paper. The advantage of the compositional view on correction is that it allows one to look beyond the surface form of a sentence. An

incorrect sentence is not just a correct sentence with a wrong surface form; it is assumed that an incorrect sentence is the result of the misapplication of rules or the application of wrong rules. Therefore, it seems right to give a characterization of correction on the level of syntactic rules.

The principle of compositionality of syntactic correction refers to parts and rules of sentences. What parts and rules are involved in the construction of a particular sentence, is determined by the grammar of the language.

3.4 String Grammars

The strict and tolerant grammars are written in the string grammar format of Rosetta (1994); the actual Rosetta system however uses a more powerful grammar formalism. A string grammar is specified by a set of basic expressions, a set of syntactic rules and a distinguished syntactic category. The basic expressions are the smallest expressions of the language. A basic expression is of the form $C(\alpha)$, where α is a string and C its category. The set of basic expressions makes up the lexicon of the language, although a sentence might also contain strings which are not basic expressions, as will become clear shortly. A syntactic rule describes how a larger expression is constructed from a number of smaller expressions, by specifying the categories of the argument strings, the category of the resulting string and the way in which the argument strings combine to form the resulting string. The basic operation performed by a syntactic rule is concatenation of its argument strings. Furthermore, a rule can add a string which does not appear as an argument to the rule. Such a string is said to be introduced syncategorematically, since the string lacks a syntactic category. The language generated by a string grammar $G = (V, R, S)$ is the set of expressions of distinguished category S which can be built from the basic expressions in V by applying rules from R. The power of string grammars as defined here is equivalent to that of context-free grammars. For an example of a string grammar, see Section 3.6 of this paper.

A string grammar defines a set of sentences together with their derivational structure. The derivational structure of a sentence is conveniently represented in a syntactic derivation tree (Montague, 1973). This tree shows the basic expressions and the syntactic rules involved in the construction of the sentence. The leaves of a syntactic derivation tree are labeled by basic expressions and the inner nodes of the tree are labeled by rules. Strings that are introduced syncategorematically never appear in syntactic derivation trees. Applying the rules in the tree recursively, with the basic expressions at the leaves as the ultimate arguments, yields the sentence represented by the tree. The structural descriptions of a sentence referred to in the outline of the method (see Figures 1 and 2) are syntactic derivation trees.

3.5 Homomorphy

As explained earlier, correction is construed as a mapping from a sentence described by a tolerant grammar onto a sentence described by a strict grammar. To be able to accomplish this mapping in accordance with the principle of compositionality of syntactic correction, the pair of strict and tolerant grammars has to satisfy two requirements: (i) for each basic expression in the tolerant grammar, there must be a corresponding basic expression in the strict grammar, and (ii) for each syntactic rule in the tolerant grammar, there must be a corresponding syntactic rule in the strict grammar.

For correction, the syntactic derivation tree of an incorrect sentence is mapped onto a syntactic derivation tree of a correct sentence by replacing each rule in the former tree by its corresponding rule from the strict grammar and by replacing each basic expression in this tree by its corresponding basic expression from the strict grammar. A syntactic rule in the tolerant grammar might be incorrect i.e. describing an incorrect syntactic construction; its corresponding syntactic rule in the strict grammar is always correct. Of course, a basic expression or syntactic rule from the tolerant grammar might be related to more than one basic expression or syntactic rule from the strict grammar. In this manner an (unambiguous) incorrect sentence can have more than one correction.

For a pair of strict and tolerant grammars satisfying both requirements above, the former grammar is structurally similar to the latter; the strict grammar is said to be homomorphic to the tolerant grammar. A pair of grammars is homomorphic if there exists a homomorphism from one grammar to the other. A homomorphism is a structure-preserving map between algebraic structures. When both grammars are specified as algebras, the requirements above are necessary conditions for the existence of a homomorphism from the tolerant grammar to the strict grammar (see (Janssen, 1983)).

One of the requirements on syntactic correction is that it should not affect meaning: the meaning of a corrected sentence should be identical to the meaning of the sentence it is the correction of. Correction essentially is changing the form of a sentence while preserving its meaning. Therefore, the correspondence between the basic expressions and the syntactic rules in the grammars should be based on meaning. Note, however, that the meaning of a syntactically incorrect sentence is not straightforwardly derivable from the proper meanings of the basic expressions and the syntactic rules involved in the construction of the sentence; the meaning of an erroneous sentence should be determined relative to what the student thinks these basic expressions and rules mean. Technically, in order to preserve the meaning of a sentence during correction, a basic expression c_t from a tolerant grammar is to correspond to a basic expression c_s from the strict grammar, such that the meaning of c_s is the meaning that the student thinks

c_t has, and *ditto* for the correspondence between syntactic rules. With these connections between the tolerant and the strict grammar, correction of an incorrect sentence yields a sentence which, in a syntactically correct way, expresses the meaning the student intended his incorrect sentence to have.

3.6 Example of a Correction

Suppose the tutoring system includes an exercise in which the student has to complete the sentence *I think that ....* The student is also given a set of words to choose from, e.g. *the*, *boy*, *man*, *saw*. Each word can be used more than once in a sentence. The actual set of words offered can of course be adapted according to the vocabulary already acquired by the student. The correct response to this exercise is a subordinate clause, so the strict grammar for this exercise is a grammar describing the subordinate clauses that can be constructed from the words given. The strict grammar G_S is specified by the set of basic expressions V_S, the set of syntactic rules R_S and the distinguished category S:

$$\begin{aligned} V_S &= N(man), N(boy), V(saw) \\ R_S &= \left\{ \begin{array}{ll} S_S: & NP(\alpha) + VP(\beta) \Rightarrow S(\alpha\ \beta) \\ S_{NP}: & N(\alpha) \Rightarrow NP(the\alpha) \\ S_{VP}: & NP(\alpha) + V(\beta) \Rightarrow VP(\beta\ \alpha) \end{array} \right\} \end{aligned}$$

A conceivable incorrect response to this exercise is a sentence in which the verb is at the end of the verb phrase. This is a plausible transfer error, likely to be made by Dutch learners of English. The tolerant grammar G_T thus has one additional rule, describing this erroneous construction.

$$\begin{aligned} G_T &= \{V_T, R_T, S\} \text{ where} \\ &\quad V_T = V_S \\ &\quad R_T = R_S \cup \{T_{VP}: NP(\alpha) + V(\beta) \Rightarrow VP(\alpha\ \beta)\} \end{aligned}$$

The correspondences between the basic expressions and the rules of the two grammars are as expected:

$$\begin{aligned} &N(man) \rightsquigarrow N(man) \\ &N(boy) \rightsquigarrow N(boy) \\ &V(saw) \rightsquigarrow V(man) \end{aligned}$$

$$\begin{aligned} &S_S \rightsquigarrow S_S \\ &S_{NP} \rightsquigarrow S_{NP} \\ &S_{VP} \rightsquigarrow S_{VP} \\ &T_{VP} \rightsquigarrow S_{VP} \end{aligned}$$

Now assume that the student actually makes the anticipated error, e.g. the completed sentence is:

(6) I think that the man the boy saw.

Then parsing the subordinate clause with the strict grammar fails; the error is detected. Parsing the clause with respect to the tolerant grammar yields a syntactic derivation tree containing the rule T_{VP} for the incorrect verb phrase (see Figure 4). Subsequently, this syntactic derivation tree is mapped onto a new structure by replacing the basic expressions from the tolerant grammar by their corresponding basic expressions from the strict grammar and replacing the syntactic rules from the tolerant grammar by their corresponding rules from the strict grammar. In this particular case, the mapping has only one palpable effect: rule T_{VP}, describing an incorrect verb phrase, is replaced by rule S_{VP}, describing a correct verb phrase. The structure obtained after the replacement of basic expressions and syntactic rules, is then used for generating the corrected sentence. In this case the correct subordinate clause is *the man saw the boy* (see Figure 4).

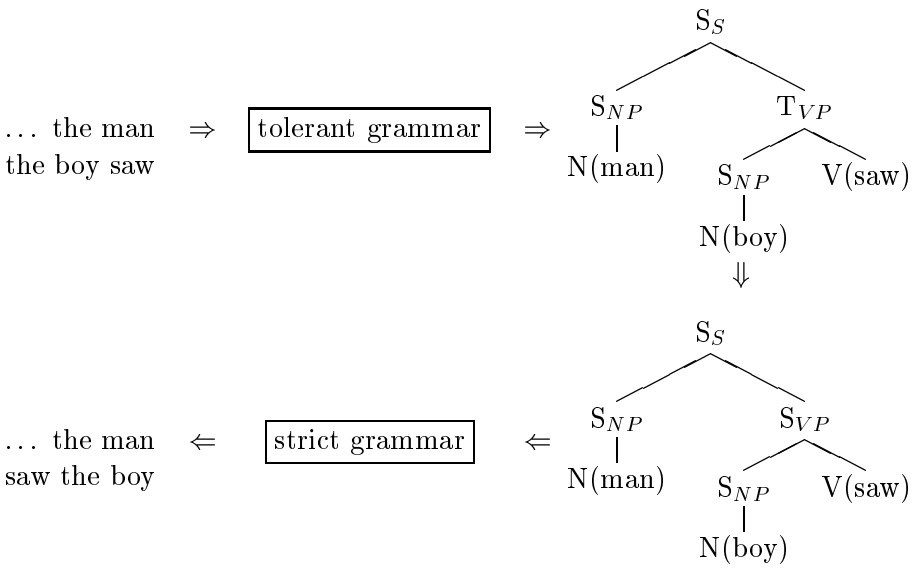

Figure 4: Correcting the subordinate clause *... the man the boy saw.*

3.7 Possible Corrections

Using homomorphic string grammars, three kinds of structural syntactic errors can be corrected: missing function words, spurious function words and incorrect order of constituents. The method also corrects morpho-syntactic errors.

3.7.1 Structural Errors

Correction of errors concerning the order of constituents in a sentence is effected by differing orders of concatenation of the arguments in the rule from the tolerant grammar and its corresponding rule from the strict grammar. An example of a correction of this kind is found in the previous section. There is a limitation as to the change in word order that can be introduced in this way: a constituent which is an argument to some rule in the tolerant grammar, must also be an argument to the corresponding rule in the strict grammar, whence a constituent can only change places with sister constituents. Should the occasion arise, this limitation can usually be overcome by rewriting the grammars. A disadvantage of rewriting however is that it is likely to compromise the linguistic generalizations according to which the grammars are arranged. As an abstract example, consider the string grammar containing the rules

$$A_1(\alpha) + A_2(\beta) \Rightarrow A(\alpha\ \beta)$$
$$B_1(\alpha) + B_2(\beta) \Rightarrow B(\alpha\ \beta)$$
$$A(\alpha) + B(\beta) \Rightarrow C(\alpha\ \beta)$$

and the basic expressions $A_1(a_1)$, $A_2(a_2)$, $B_1(b_1)$ and $B_2(b_2)$, generating the string $a_1\ a_2\ b_1\ b_2$. Assuming straightforward correspondences between the basic expressions, no grammar generating a permutation of the string above in which either the two a's or or the two b's are not adjacent (e.g. $a_1\ b_1\ b_2\ a_2$), is homomorphic to the given grammar, unless the rules are rewritten into one rule of the form:

$$A_1(\alpha) + A_2(\beta) + B_1(\gamma) + B_2(\delta) \Rightarrow C(\alpha\,\beta\,\gamma\,\delta)$$

thus loosing categories A_1, A_2, B_1 and B_2.

Missing function words are added by syncategorematic introduction of the missing word in the strict grammar. A simple example of a missing function word in Dutch is a singular noun phrase without a determiner. Deletion of illicit function words is achieved using the same mechanism, viz. through syncategorematic introduction of the pertinent word in the tolerant grammar. In this manner, an erroneous sentence like

(7) *Who do you think that will eat the apple?

can be corrected into

(8) Who do you think will eat the apple?

3.7.2 Morpho-syntactic Errors

The homomorphic grammar approach to error correction outlined above can be extended to also include the treatment of morpho-syntactic errors. To

this goal, the basic expressions of a string grammar are enriched with attributes. Agreement between constituents, or more generally, the presence of a particular value for some attribute, is enforced by the syntactic rules of the grammar. A syntactic rule selects its arguments not only by specifying their category, as before, but also by prescribing specific values for certain attributes of the arguments. Morpho-syntactic errors are described in the tolerant grammar by rules which select argument expressions with inappropriate attribute values. These tolerant rules are easily derived from the correct rules in the strict grammar by tampering with the values of the attributes occurring in the rules as desired.

In general, a sentence with an agreement error can be corrected in more than one way, because the value of any of the occurrences of the attribute involved can be changed in order to restore agreement. For example, the sentence $^{*}I$ *walks* can be corrected by either changing the person of the pronominal subject—changing I into *he*—or by changing the person of the verb—changing *walks* into *walk*. The corrections thus found are respectively *he walks* and *I walk*. For an agreement error which has more than one possible correction, often one of the corrections is superior to the others, because that correction is more likely to correspond to the intended meaning of the incorrect sentence than the other corrections. In the example above, the student most probably made an error with regard to the inflection of the verb, rather than picking the wrong pronoun. Therefore the correction *I walk* is to be preferred over the correction *he walks*.

To incorporate preferences for particular corrections in the method, the occurrences of attributes in the syntactic rules are marked either weak or strong and so are the attributes associated with the basic expressions. This distinction is used to restrict correspondences between the basic expressions and between the syntactic rules in the tolerant grammar and the strict grammar. Given a correct syntactic rule requiring agreement for a particular attribute occurring in the argument expressions of the rule, one of these occurrences is marked strong and the other occurrences are marked weak. A tolerant rule and its corresponding strict rules must show identical values for each of the strong occurrences of the attributes associated with the expressions appearing in the rule. The values of the weak occurrences may differ. Once an attribute of a complex expression is declared weak or strong in a rule, this property percolates down through the rules to the basic expressions. If in a typical rule for agreement $R_C : A(a : i) + B(a : i) \Rightarrow C$, attribute a is declared weak for A, then in a rule $R_A : D(a : i) + E(a : i) \Rightarrow A(a : i)$, attribute a is also weak for D and E. Alternatively, if in rule R_C attribute a is strong for A, then in rule R_A one of the occurrences of a must be strong as well; the other is weak. Analogously to corresponding syntactic rules, a basic expression in the tolerant grammar and its corresponding basic expression in the strict grammar are allowed to differ only with respect to the values of their weak attributes. The values of the strong attributes are required to

be the same.

As a consequence of these restrictions on correspondences between basic expressions and between syntactic rules, the mapping of derivation trees will not affect the value of any of the strong attributes appearing in the basic expressions of the sentence to be corrected; only weak attributes possibly change their values. Declaring an attribute strong or weak thus determines the way in which an agreement error will be repaired. Given an agreement error with regard to some attribute which involves two basic expressions, correction will imply a change in the basic expression for which the occurrence attribute is declared weak, such that the value match the value of the occurrence which is declared strong. If the attribute concerned is marked strong for both basic expressions, nothing will change and no correction will be found. When both attributes are weak, all possible corrections will be generated. In the example from the beginning of this section, it is assumed that agreement errors for pronominal subjects and verbs derive from incorrect verb inflections. Therefore it makes sense to mark the agreement attributes of the verb, i.c. *person* and *number*, as weak, and those of the subject as strong. Accordingly, the incorrect sentence **I walks* will be corrected into *I walk*. For morpho-syntactic errors other than agreement errors, e.g. errors with regard to case, the relevant attribute must be weak, otherwise its value cannot be changed. (For more details, see (Harkema, forthcoming).)

4 Concluding Remarks

In this paper, a method for the detection and correction of grammatical errors in natural language sentences has been presented. The method is designed for use in tutoring systems for second language instruction. This application has three properties which render the general problem of correction tractable. Firstly, the errors to be corrected by the method are to a very large extent systematic and predictable. Secondly, the grammars have to describe fragments of languages rather than entire languages. Finally, each sentence is produced by the student in the context of an exercise, which makes it possible to trace the intended meaning of the sentence. These three characteristics of the problem enable an anticipatory approach to correction in which there is an explicit grammar for incorrect sentences which is systematically related to a grammar describing correct sentences. The proposed method has several features:

- The method performs a full grammatical analysis of the input sentence. As a result, the method, when equipped with the proper grammars, is able to deal with fairly unrestricted input from the student. This makes the method fit for use in a wide range of exercise formats.

- There is no focus on specific kinds of syntactic errors. The method works for any pair of homomorphic grammars.
- By taking into account strategies for second language learning when writing tolerant grammars, these grammars will represent a credible model of the learner's linguistic competence. The tolerant grammars reflect the interlanguage of the learner.
- Having explicit strict grammars for generating corrections ensures that the generated sentences are correct indeed.
- There is a mathematically rigorous relationship between a tolerant grammar and a strict grammar and consequently between an erroneous sentence and its corrections. This facilitates implementation of the method.

The method is to be used in exercises whose linguistic complexity justifies the full grammatical analysis of the input sentence that is performed by the method. On the other hand, successful application of the method also comes with a restriction on the language fragment, which seems to ask for simple exercises: the correct sentences of the language fragment must be described exhaustively in a strict grammar. Actually writing some comprehensive grammars will show whether there are language fragments which satisfy this restriction and yet are non-trivial. Currently, the method is implemented in a modest exercise dealing with Dutch subordinate clause (see Section 3.6). After some extensions, this implementation can be the vehicle for an empirical evaluation of the method, by letting students actually use the implemented exercises. Such an evaluation should provide the answers to two important questions. The first question is about the coverage of the tolerant and strict grammars: how hard is it to come up with tolerant grammars which describe a considerable number of the syntactic errors the students actually make and how many incorrect sentences are judged to be correct due to some spurious analysis with the strict grammar? A related question is to what extent the format of an exercise might be helpful in keeping the sentences of the student within the bounds of the tolerant and strict grammars. The second important question regards the adequacy of tolerant grammars: does the tolerant grammar assign correct analyses to incorrect sentences, or, in other words, do the generated corrections match the students' intended utterances?

Another line of future investigation concerns the power of homomorphic string grammars. From a quick scan of the list of erroneous sentences in (Burt and Kiparsky, 1972), it appears that most of the syntactic errors described therein can be treated by a pair of homomorphic grammars, but each of these sentences contains a single error. In reality incorrect sentences are often fraught with more than just one error. Accommodating multiple

errors in a tolerant grammar is possible, but often requires rewriting of the strict grammar in order to maintain homomorphy. Rewriting a grammar to maintain homomorphy, however, nearly always entails a loss of structure of the grammar. It might turn out to be the case that the simultaneous correction of particular errors can only be achieved through a structureless strict grammar, i.e. a grammar which is simply an enumeration of possible sentences. This, of course, would be an undesirable situation. The problem could be solved by defining a more intricate correspondence between the tolerant and the strict grammars or by using a grammar formalism whose rules have operations more powerful than simple concatenation of arguments.

References

Burt, M.K. and C. Kiparsky. 1972. *The Gooficon, a repair manual for English.* Newbury house publishers.

Carroll, S. and M. Swain. 1993. Explicit and implicit negative feedback, an empirical study of the learning of linguistic generalizations. *Studies in second-language acquisition*, 15: 357–386.

Catt, M. and G. Hirst. 1990. An intelligent CALI system for grammatical error diagnosis. *CALL*, 3: 3–26.

Clahsen, H. and P. Muysken. 1986. The availability of Universal Grammar to adult and child learners: a study of the acquisition of German word order. *Second language research*, (2):93–119.

Dekeyser, R.M. 1993. The effect of error correction on L2 grammar knowledge and oral proficiency. *The modern language journal*, 77: 501–514.

Garman, J., J. Martin, P. Merlo, and A. Weinberg. 1993. A principle-based parser for foreign language training in German and Arabic. In *Proceedings of the third international workshop on parsing technologies.*

Harkema, H. forthcoming. A method for correcting syntax errors in tutoring systems for second language instruction. Technical report, Institute for Perception Research.

Hoe, R. Van, E. Appelo, H. Harkema, J. Masthoff, J.R. de Pijper, and J. Rous. 1995. Appeal: a multimedia environment for adaptive learning. In *Proceedings of the World Conference on Educational Multimedia and Hypermedia.*

Janssen, T.M.V. 1983. *Foundations and applications of Montague grammar.* Ph.D. thesis, University of Amsterdam.

Labrie, G. and L.P.S. Singh. 1991. Parsing, error diagnostics and instruction in a French tutor. *CALICO journal*, pages 9–25.

Littlewood, W. 1984. *Foreign and second language learning. Language acquisition research and its implications for the classroom.* Cambridge language teaching library.

Montague, R. 1973. The proper treatment of quantification in ordinary English. In J. Hintikka, J. Moravcsik, and P. Suppes, editors, *Approaches to natural language*. Reidel.

Rosetta, M.T. 1994. *Compositional translation.* Kluwer academic publishers.

Selinker, L. 1983. Interlanguage. In B. Robinett and J. Schachter, editors, *Second language learning: contrastive analysis, error analysis and related aspects.* University of Michigan Press.

Swarz, M.L. and M. Yazdani, editors. 1992. *Intelligent tutoring systems for foreign language learning.* Springer Verlag.

Wenger, E. 1987. *Artificial intelligence and tutoring systems.* Morgan Kaufman publishers.

White, L. 1989. *Universal grammar and second language acquisition.* John Benjamins, Amsterdam.

White, L. 1991. Adverb placement in second language acquisition: some effects of positive and negative evidence in the classroom. *Second language research*, 7: 133–161.