Model–based Diagnosis Under Structural Uncertainty

Wolfgang Menzel & Ingo Schröder[1]

Abstract. Model–based diagnosis techniques are built on the common assumption that the structure of the system under observation is known in advance. Unfortunately, this assumption does not hold for the diagnosis of grammatical faults in natural language utterances. In such an application the structural identification of components and of their mutual relationship becomes an integral part of the diagnostic task itself. An approach to the constraint–based diagnosis of faulty natural language utterances is presented, and its application in tutoring systems for foreign language learning is discussed.

1 INTRODUCTION

Typical approaches to model–based diagnosis share the common assumption that any malfunction of a system is caused by faulty behavior of some of its components. If we assume, furthermore, the structure of the object under consideration to be (almost) fixed and explicitly known to an external observer, a diagnosis becomes possible by assigning different modes of behavior to the individual components and checking the consistency of these assignments with the system description and the observations made so far [1, 9]. While these assumptions seem plausible for most of the technical devices and situations to which a model–based diagnosis is applied, they are no longer justified if natural language utterances are going to be subject of diagnostic efforts. This happens e. g. in advanced language tutoring systems which are expected to comment on the quality of a student's contribution without constraining the flexibility of language use too much.

Usually, there is a natural uncertainty about the structure of the diagnostic object (i. e. the incoming utterance), both in terms of structural decomposition and type identity of its components.

Pragmatically sensible exercises, as they are emphasized in a communication–oriented approach to language learning, open a considerable space of possible responses from which a student may choose her answer. Ideally, this space is restricted only by the rules of grammar and some general maxims for pragmatically adequate behavior in that particular setting instead of imposing rather limited and rigid schemata, which are frequently used to implement typical slot–filling exercises in today's language learning systems.

Since the internal structure of an utterance is not directly visible to an observer, a considerable degree of ambiguity is introduced, in particular

- *lexical ambiguity*, i.e. the identity of system components depend on the context in which they appear and
- *structural ambiguity* with respect to the syntactic and semantic relations established among system components.

[1] Fachbereich Informatik, Universität Hamburg, Vogt–Kölln–Straße 30, 22527 Hamburg, Germany, E–Mail {menzel|ingo.schroeder}@informatik.uni-hamburg.de

ECAI 98. *13th European Conference on Artificial Intelligence*
Edited by Henri Prade
Published in 1998 by John Wiley & Sons, Ltd.

Under such conditions, the no–function–in–structure principle of model–based diagnosis can not be enforced any longer. In addition to these elements of uncertainty, cases of inadmissible and inappropriate structure can be observed quite frequently in student utterances. Even if the space of acceptable structural configurations can be described by licensing constraints (grammar), the problem of structural identification (parsing) remains to be solved. It becomes especially difficult in the presence of faulty utterances which force the parser to consider two incompatible alternatives at the same time:

- relaxing certain well–formedness conditions to be able to establish a structural interpretation in spite of ungrammatical constructions
- insisting on the integrity of virtually the same conditions to efficiently guide the parsing procedure

This results in a rather difficult relationship between parsing and diagnosis which creates a kind of vicious circle: On the one hand, diagnosis requires a structural hypothesis to find out which well–formedness constraints are violated (i. e. parsing has to precede diagnosis), whereas at the same time parsing itself needs strong predictions about plausible constraint violations to avoid misleading decisions (i. e. diagnosis has to precede parsing).

All these peculiarities lead to a situation which, if translated into an electrical analog, would yield a circuit to be diagnosed which not only may contain components with faulty behavior, but whose detailed internal structure is unknown, i. e. the wiring and type identity of system components is only partially observable, and which may implement a faulty design.

Obviously, this is not a typical situation for technical diagnosis tasks. However, it can be imagined that similar situations might well occur during certain phases of experimental design processes where both the design and the behavior of system components are subject to diagnosis [9].

In the following, we describe an approach for model–based diagnosis of natural language utterances in the presence of ambiguity and faulty structures. We start with the simple case of a fixed and well–formed structure in Section 2. Using the common approach of reasoning about structure and behavior, Section 3 presents a procedure for structural identification, which is based on a parsing scheme using constraint satisfaction techniques. It allows to find a plausible interpretation of the incoming utterance even in the presence of considerable constraint violations, but in some cases yields only approximate diagnostic results (cf. Section 4). In Section 5, finally, the two techniques are integrated into a two stage diagnosis, which combines their respective advantages: structural flexibility and high diagnostic quality.

2 STRUCTURE–BASED DIAGNOSIS

Model–based diagnosis of grammatical faults becomes possible if subdomains of grammar can be identified

1. for which an exhaustive description of well–formedness is feasible and
2. which allow the design of didactically relevant exercises with little structural variation.

Surprisingly enough, even in the highly complex domain of natural language small subsets of grammatical regularities do exist for which both conditions hold. Among them are agreement and government phenomena, ordering regularities and conditions on the existence or absence of specific surface forms [4, 5]. These well–formedness conditions as well as the lexical information provided by the word forms in the utterance are mapped onto the components of the diagnostic problem.

Applications first of all can be found in exercises for the composition of complex phrases (e. g. noun or prepositional phrases), including inflectional regularities and morpho–syntactic aspects of pronominal reference.

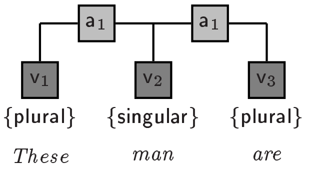

Figure 1. Component structure of a simple agreement diagnosis problem.

*(1) * These man are tall.*

The well–formedness conditions for a simple sentence like example 1 can be described by two components type–agree(a_i), $i \in \{1, 2\}$ for number agreement between determiner and noun as well as between noun and verb, respectively, and three components type–value(v_i), $i \in \{1, 2, 3\}$ for lexical value assignment for the three forms *these*, *man*, and *are*. By separately modeling agreement conditions and lexical value assignments a diagnosis of two different error types becomes possible.

1. The student violated or ignored a certain grammatical regularity (e. g. an agreement constraint) during sentence production.
2. The student used an erroneous lexical representation (e. g. for number features) when checking for consistency.

Within the knowledge about well–formedness the agreement conditions correspond to *rules* of grammar, whereas lexical assignments belong to the class of *factual* information. Note that from a diagnostic point of view only the lexical assignment components are directly *visible*, since they correspond to explicitly given surface word forms in the utterance. The agreement components, however, are established indirectly by the underlying and therefore *invisible* structural relationships. This difference is negligible as long as the (fixed) structure of the sentence is given. It becomes a crucial, though, if the structure is unknown to the diagnostic system which has to infer even the existence of a component by means of techniques for structural identification. In our example the structural relationships are given by the assumptions that *these* is the determiner of the noun *man*, and *man* in turn serves as the subject of the finite verb *are*. This structural backbone can be described by *dependency relations* which uniquely subordinate each word form to another one or to an additional root node (cf. Figure 2a).

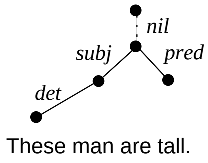

(a)

$$\begin{aligned} v_1 &= \langle det, 2\rangle \\ v_2 &= \langle subj, 3\rangle \\ v_3 &= \langle nil, 0\rangle \\ v_4 &= \langle pred, 3\rangle \end{aligned}$$

(b)

Figure 2. (a) Syntactic dependency tree for an example utterance: For each word form a unique subordination and a label, which characterizes the kind of subordination, are to be found. (b) Labellings for a set of constraint variables: Each variable corresponds to a word form and takes a pairing consisting of a label and an index (corresponding to the superordinated word form) as a value.

Components come along with ports which have to be connected properly yielding the structure of the diagnostic problem in Figure 1:[2]

$\mathrm{out}(v_1)=\mathrm{in}(a_1) \wedge$
$\mathrm{out}(v_2)=\mathrm{out}(a_1)=\mathrm{in}(a_2) \wedge$
$\mathrm{out}(v_3)=\mathrm{out}(a_2)$

Finally, the component behavior has to be specified at least for the mode of well–formedness:

$\text{type–agree}(X) \wedge \neg \mathrm{ab}(X) \rightarrow \mathrm{in}(X)=\mathrm{out}(X)$
$\text{type–value}(X) \wedge \neg \mathrm{ab}(X) \rightarrow \mathrm{out}(X) \in \mathrm{in}(X)$

If desired, this description can be completed with a specification of the abnormal mode of behavior.

The observation of the system under diagnosis (i. e. the utterance) consists of sets of possible feature values, which the dictionary assigns to the lexical items introduced by the student.

$\mathrm{in}(v_1)=\{\mathrm{plural}\} \wedge \mathrm{in}(v_2)=\{\mathrm{singular}\} \wedge \mathrm{in}(v_3)=\{\mathrm{plural}\}$[3]

From this system of axioms five minimal diagnoses can be derived, among them

$D_1=\{a_1,a_2\}$: The student ignored both number agreement conditions.

$D_2=\{v_2\}$: The student did not know the correct number feature of *man*.

Afterwards, an appropriate one (probably D_2) can be selected for explanatory purposes. More ambitious exercises can be devised for languages with richer an inflectional system e. g. German or the Slavic languages [5]. Furthermore, the basic idea easily carries over to other types of well–formedness conditions, namely linear precedence constraints and simple types of co–occurrence requirements, thus accommodating even a few though rather simple cases of structural variation.

The approach yields an extremely high diagnostic precision, which even in complex and therefore almost unexplainable error configurations allows to generate highly reliable correction proposals. For this purpose, the desired feature values as determined by the diagnosis procedure can be used immediately to retrieve word form substitutes from the dictionary.

Moreover, the somewhat redundant modeling of agreement and value assignments as both being defeasible components gives rise to *multiple error descriptions* among which a suitable one can be selected according to the desired perspective:

[2] A analogous approach has been proposed for error diagnosis in algebraic exercises [8].

[3] Note that the typical case in inflectionally rich languages is an ambiguous assignment, as for e. g. *fish* in(number–value)={singular, plural} holds, thus yielding indefinite theories as the regular case.

- rule violations (well–formedness constraints)
- erroneous factual information (value assignments)
- minimum effort repair proposals

Additional interpretations are available due to the ability of model–based diagnosis to derive alternative solutions which might lead to competing candidates for error explanation. This allows to adopt a hierarchical error explanation scheme which offers increasingly detailed error descriptions to the student in order to guide her towards an acceptable solution:

- a general well–formedness judgment
- an explanation in terms of well–formedness constraints
- an explanation in terms of lexical features
- alternative views on the error
- the precise location of a minimum effort repair proposal
- the repair proposal itself

The most important drawback of the model–based approach derives indeed from the need to presume an almost fixed structure which restricts its applicability to a few, carefully designed types of language training exercises. Nevertheless, it still offers the student a considerable freedom of lexical choice.

3 ROBUST PARSING

Well–formedness constraints, which are necessary for the purpose of model–based diagnosis, are of a secondary nature since even their existence depends on the structural properties of the utterance under diagnosis. Therefore, a component of structural identification is required. It is expected to find the most plausible structural hypothesis despite the presence of ungrammaticality in the utterance and, therefore, must be ready to treat partial constraint violations as the normal rather than the exceptional case.

For this reason, a parsing scheme based on constraint satisfaction techniques has been adopted [3, 2] where constraints are used to license dependency relations between word forms of the utterance. To facilitate the parsing of ill–formed input, where classical constraint satisfaction problems (CSP) become over–constrained, the approach has been extended to support the use of graded constraints thus being able to distinguish between more or less serious constraint violations.

Each word form w_i of the utterance $\langle w_1, \ldots, w_n\rangle$ has to be subordinated to another word form or an artificial node w_0. In addition, each subordination relation should be characterized by a label $l \in L$. Therefore, n constraint variables are established, for each of which an unambiguous value taken from the domain $L \times \{0, 1, \ldots, n\}$[4] has to be found (cf. Figure 2). By the application of constraints to local dependency structures the search space is narrowed down until eventually a unique solution can be extracted which contains the least serious constraint violations.

For reasons of efficiency, constraint propagation is restricted to at most binary constraints.

In order to guide the candidate elimination procedure every constraint receives a weight from the interval $[0, 1]$ in which zero denotes a hard constraint in the classical sense which must hold under any circumstance. The higher the weight, the more acceptable a constraint violation is considered. Constraints with a weight close to zero are used to model usual well–formedness conditions, which can be violated only if no other interpretation of the utterance is available:[5]

{X} : SubjNumber : Subj : 0.1 :
    X.label=subj → X↓num=X↑num
'The subject agrees with the verb with respect to number.'

Larger weights are attached to constraints which model preferences for certain structural configurations:

{X} : SubjOrder : Subj : 0.9 :
    X.label=subj → X↓pos<X↑pos
'The subject is placed in front of the verb.'

Due to the existence of preferences the CSP becomes over–constrained even for well–formed utterances. In an attempt to find the optimum value assignment a unique solution to the CSP can be determined which might serve as a structural description of the student's utterance during further processing.

Note that the condition part of these constraints explicitly exhibits the primary nature of structural relations with different patterns of behavior depending on it.

Hard constraints usually serve to eliminate linguistically meaningless structures from the candidate space:

{X} : SubjInit : Subj : 0.0 :
    X.label=subj → X↓cat=NOUN ∧ X↑cat=FINVERB
'A subject is a noun, and it modifies a finite verb.'

{X, Y} : SubjUnique : Subj : 0.0 :
    X.label=subj ∧ X↑id=Y↑id → Y.label≠subj
'The subject is unique for a given verb.'

Hard constraints on structural configurations lead in some cases of ungrammaticality to a situation where the constraint satisfaction procedure fails to determine a complete structural description for the utterance. However, a number of partial analyses can be obtained in such cases.

(Partial) Solutions of the CSP are compared by means of a score which is computed by accumulating the weights of all constraints violating a particular lexically instantiated dependency structure.

Due to the use of graded constraints a certain degree of robustness is introduced into the parsing process which allows to determine the intended structure even for erroneous utterances. An experiment with the systematically distorted sentence

(2) Der Mann besichtigt den Marktplatz.
(The man visits the marketplace.)

and a grammar of modest size (consisting of approximately 160 constraints) revealed that deviations with error degree[6] up to two were tolerated almost completely [6].

The level of robustness can be raised even higher by introducing a strong support from semantic (e. g. sortal restrictions) and background knowledge which is derived from the exercise environment. This was achieved by establishing a multi–level architecture which integrates several representational layers (e. g. for syntax and semantics) into a single CSP. The different layers are symmetrically coupled by means of mapping constraints, which facilitate a bidirectional information flow between parallel structural representations. Thus, semantic support helps to overcome possible syntactic difficulties whereas syntactic evidence will further the interpretations of semantically odd utterances.

---

[4] A value $\langle nil, 0\rangle$ serves as an indicator for the root node of the dependency tree.

[5] A constraint consists of a variable declaration, a constraint identifier, a constraint class identifier, the weight and the condition delimited by colons. A variable (e. g. X) can be bound to a dependency relation. The selector functions X.label, X↑ ⟨feature⟩ and X↓ ⟨feature⟩ provide access to the label of the relation and lexical properties or positional information of the upper or lower node, respectively.

[6] The self–defined error degree counts normal errors as single faults and assigns higher counts to more serious deviations.

A repetition of the above experiment with supporting and contradicting semantic evidence confirmed this characteristic behavior. Confirming support from a complementary level increased the robustness of structural identification considerably, whereas contradicting evidence caused a corresponding reduction (cf. Figure 3).

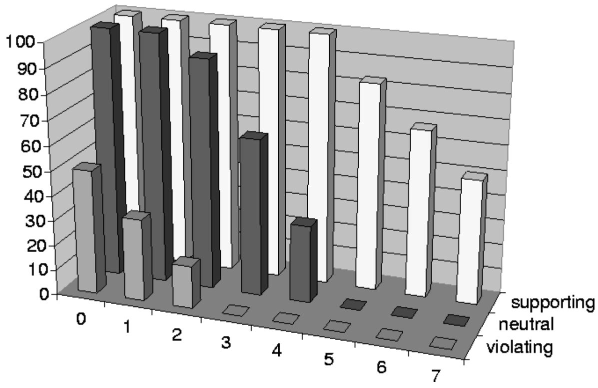

Figure 3. Percentage of correct analyses depending on the error degree for syntactic faults and additionally parameterized by the kind of semantic and domain–specific support: The better the semantic and domain–specific support, the more errors can be compensated for. Neutral support means no support, e. g. due to missing information.

A closer analysis indeed reveals the desired symmetrical relationship between adjacent layers. On the syntactic level even an error degree of five can be tolerated (and diagnosed) as long as sufficient semantic support is available. Severe semantic deviations (e.g. a simultaneous violation of semantic sortal restrictions and domain specific preferences) will be identified properly if the syntactic background is reliable enough.

4 ROBUST PARSING AS DIAGNOSIS

Robust parsing as described in the previous chapter generates a structural description for any utterance irrespective of it being grammatically or semantically sound. Acceptable and deviant utterances can be distinguished from each other only by the number and kind of constraints violated by their structural interpretation. If conflicts with preference constraints only have been detected the utterance can be accepted as being a grammatical one, while the violation of well–formedness conditions indicates more or less serious error configurations.

More detailed diagnostic results can easily be extracted from the output of the parser, since in most cases there is a straightforward mapping between constraint violations and possible explanations. This close correspondence has already been noted by Ohlsson [7] who proposed a constraint–based approach to the diagnosis of student errors but restricted its application to ("lexically" and structurally) unambiguous problems like algebraic calculations or the construction of molecular structures of chemical compounds.

Thanks to the use of graded constraints, a good deal of agreement violations, linear ordering problems, missing obligatory constituents, violated government requirements, and ignored sortal restrictions can be diagnosed even in highly ambiguous situations [6]. Note that this diagnostic capability is restricted to the invisible, i. e. structurally dependent components.

Thus, the application of graded constraints facilitates an integrated treatment of parsing *and* diagnosis for a wide range of (possibly ill–formed) utterances. The same mechanism of arbitration between conflicting evidence allows to tolerate *and* identify constraint violations. Being designed to determine the least corrupted interpretation among a vast variety of competing structural hypotheses, the procedure is able to perform the desired structural identification for diagnostic objects of unknown underlying structure producing diagnostic results as a byproduct.

However, if compared to the structure–based approach described in Section 2, the quality of diagnostic results as produced by the robust parser suffers from a number of shortcomings:

- The constraint satisfaction procedure is only able to passively test for lexical features of word forms, but not to actively assign or update them. Therefore, the transitive closure of local well–formedness constraints (e. g. agreement and linear precedence conditions) can only be approximated, producing imprecise diagnostic results in the general case. Combinations of several agreement requirements for morpho–syntactically ambiguous word forms (as in **This fish are good*) can only be modeled by defining specifically tailored constraints which establish pairwise agreement between all participating word forms.
- The focus on constraint violations leaves no room for the introduction of different error views. In particular the diagnosis of noisy lexical assignments (i. e. visible components) does not seem feasible, since this would remove the last source of reliable clues which are required to guide the search for a solution. In this way, however, an important class of students' misconceptions is excluded from the space of possible explanations. For the example of Section 2 only the diagnosis $D_1=\{a_1,a_2\}$ is directly available, but not the considerably more concise description $D_2=\{v_2\}$.
- As any other lexical assignment part–of–speech categories are considered to be non–defeasible. In combination with the presence of hard constraints on admissible structures this might cause the impossibility to establish a parse tree for the complete utterance (cf. Figure 4 for an example).

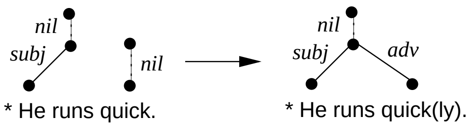

*(3) * He runs quick. (*cf. German: *Er rennt schnell.)*

Figure 4. Partial parses: If only the part–of–speech category was changed the sentence could be analyzed completely.

Possible remedies to these problems will be discussed in the following section.

5 TWO STAGE DIAGNOSIS

So far, two approaches to the diagnosis of grammatical faults have been presented:

1. model–based techniques which provide high quality results for utterances with no or little structural variation
2. a method for robust parsing which offers free form input but generates less ambitious diagnostic results

Both procedures are based on constraint satisfaction techniques with the important difference that the first case is defined to assign grammatical features to CSP variables while in the second partial structural descriptions are considered. This difference leads to complementary diagnostic abilities. Therefore, even a simple sequential combination allows to create a solution which shares the particular advantages. By using the robust parser as a component for structural identification a structural hypothesis becomes available, which serves a subsequent model–based diagnosis to guide the simulation of different behavior modes for well–formedness constraints and lexical assignments.

Thus, the desired high precision becomes also available for transitive well–formedness conditions. Additionally, alternative error interpretations can be provided if the simulation of lexical assignment faults for morpho–syntactic features is integrated into the model–based diagnosis of the second stage as described in Section 2. Of course, such a subsequent diagnosis does not allow to compensate deficiencies of the robust parser. Determining the wrong structural interpretation in the first stage will ultimately lead to useless diagnostic results.

There is, however, a class of erroneous part-of-speech assignments which can be diagnosed by trying to integrate partial parsing results through simulated category substitutions. This is rather straightforward if a missing constituent matches a superfluous partial tree, under the condition of relaxed categorial constraints. For instance, the two fragments for the utterance in Figure 4 can easily be integrated if the category of *quick* is changed from adjective to adverb.

*(4) * The man looks the picture (*cf. German: *Der Mann betrachtet das Bild.)*

More difficult cases will require the introduction of additional nodes into the resulting structure as in example 4 where, due to the available background knowledge, a high preference to consider *a picture* as the theme of the verb is given. Only by introducing a suitable preposition an appropriate syntactic structure can be created.

The class of interference errors is more difficult to attack. They result from the frequently observed transfer of first language regularities into the yet developing grammar of the foreign language where they can lead to systematic error situations. From a diagnostic perspective the identification of such cases would require to prove the unjustified existence of an invisible component in the diagnostic problem! In German, for instance, the finite verb has to be placed in the second position in a main clause. If this word order regularity is transferred into English, in many cases no immediate consequences can be observed, because accidentally a well–formed utterance has been produced despite using inappropriate constraints:

*(5) The man is tall. (*cf. German: *Der Mann ist groß.)*

The problem becomes only visible if constraints of the target language are violated as in

*(6) * Obviously is the man tall.*
*(*cf. German: *Offensichtlich ist der Mann groß.)*

Certainly, all cases of interference which cause the violation of target language constraints, can be diagnosed in terms of target language regularities. The true nature of the problem, however, can only be identified by carefully comparing a number of structurally equivalent configurations both wellformed and deviant. Specifically tailored constraints can be introduced, which cover the suspicious cases and trigger appropriate error messages. Unfortunately, this no longer can be considered a model-based approach, because it maps anticipated symptom collections directly onto diagnostic categories.

The difficulty to find a more general solution is obviously related to the fact that evidence for the existence of a first language constraint can not be derived from the structure of a second language utterance. Instead it depends on the structure of a suitable first language equivalent. Therefore, it remains unclear, how a truly model-based approach to the diagnosis of interference phenomena can be designed at all.

6 CONCLUSIONS

A model–based approach to the diagnosis of grammatical faults has been complemented by a robust parsing procedure which serves as a component for structural identification and is based on constraint satisfaction techniques. This combination of methods for the first time facilitates an error diagnosis in free form utterances. Robustness is achieved by using graded constraints to arbitrate between conflicting evidence. It has been enhanced, furthermore, by the implementation of a multi–level architecture combining the defeasibility of constraints on all representational levels with a bidirectional cross–level information flow. Thus, the approach facilitates mutual support between representational levels without the necessity to specify a pre–defined anchor point consisting of certain knowledge. Due to strong (but defeasible) semantic support a wide variety of syntactic constraint violations can be diagnosed. On the other hand, (defeasible) syntactic evidence is used to detect semantic anomalies, which might be caused e.g. by prior comprehension difficulties.

Since highly specific information from the exercise environment (the topic to be talked about, the expectations raised from the preceding context of interaction, etc.) can easily be included into the multi–layer representation, the approach particularly encourages the design of pragmatically ambitious exercises. In strong contrast to previous approaches such a biased diagnosis does not need to focus solely on the syntactic form, but will be able to account for the communicative relevance of a student's contribution as well. This opens up completely new opportunities for the development of intelligent systems for computer–assisted language learning, which put emphasis primarily on language communication instead of language reproduction.

ACKNOWLEDGEMENTS

This research has partly been funded by the German Research Agency "Deutsche Forschungsgemeinschaft" under grant no. Me 1472/1-1.

REFERENCES

[1] R. Davis, 'Diagnostic reasoning based on structure and behaviour', *Artificial Intelligence*, **24**(1–3), 347–410, (1984).

[2] Mary P. Harper and Randall A. Helzerman, 'Managing multiple knowledge sources in constraint-based parsing of spoken language', Technical Report EE-94-16, School of Electrical Engineering, Purdue University, West Lafayette, (1994).

[3] Hiroshi Maruyama, 'Constraint dependency grammar', Technical Report RT0044, IBM Research, Tokyo Research Laboratory, (1990).

[4] Wolfgang Menzel, 'Error diagnosing and selection in a training system for second language learning', in *Proceedings 12th International Conference on Computational Linguistics, Coling '88*, pp. 414–419, Budapest, (1988).

[5] Wolfgang Menzel, *Modellbasierte Fehlerdiagnose in Sprachlehrsystemen*, volume 24 of *Sprache und Information*, Niemeyer, Tübingen, 1992.

[6] Wolfgang Menzel and Ingo Schröder, 'Constraint-based diagnosis for intelligent language tutoring systems', in *Proceedings of the IT&KNOWS Conference at the IFIP '98 Congress*, Wien/Budapest, (1998).

[7] Stellan Ohlsson, 'Constraint-based student modelling', in *Student Modelling: The Key To Individualized Knowledge-Based Instruction*, eds., Jim E. Greer and Gordon I. McColla, 167–189, Springer-Verlag, Berlin, (1994).

[8] John Self, 'Cognitive diagnosis for tutoring systems', in *Proc. of the 10th ECAI*, ed., Bernd Neumann, pp. 699–703, Vienna, Austria, (1992).

[9] Peter Struß, 'Knowledge-based diagnosis — an important challenge and touchstone for AI', in *Proceedings 10th European Conference on Artificial Intelligence*, pp. 863–874, Budapest, (1992).