Robust Processing of Natural Language

Wolfgang Menzel

Fachbereich Informatik, Universität Hamburg
Vogt-Kölln-Straße 30, 22527 Hamburg, Germany

Abstract. Previous approaches to robustness in natural language processing usually treat deviant input by relaxing grammatical constraints whenever a successful analysis cannot be provided by "normal" means. This schema implies, that error detection always comes prior to error handling, a behaviour which hardly can compete with its human model, where many erroneous situations are treated without even noticing them. The paper analyses the necessary preconditions for achieving a higher degree of robustness in natural language processing and suggests a quite different approach based on a procedure for structural disambiguation. It not only offers the possibility to cope with robustness issues in a more natural way but eventually might be suited to accommodate quite different aspects of robust behaviour within a single framework.

1 Robustness in Natural Language Processing

The notion of robustness in natural language processing is a rather broad one and lacks a precise definition. Usually, it is taken to describe a kind of monotonic behaviour, which should be guaranteed whenever a system is exposed to some sort of non-standard input data: A comparatively small deviation from a predefined ideal should lead to no or only minor disturbances in the system's response, whereas a total failure might only be accepted for sufficiently distorted input.

Under this informal notion robustness may well be interpreted as a system's indifference to a wide range of external disruptive factors including

- the inherent uncertainty of real world input, e.g. speech or hand writing,
- noisy environments,
- the variance between speakers, for instance idiolectal, dialectal or sociolectal,
- "erroneous" input with respect to some normative standard,
- an insufficient competence of the processing system, if e.g. exposed to a non-native language or new terminology,
- highly varying speech rates and
- resource limitations due to the parallel execution of several mental activities.

One of the most impressive features of human language processing is the ability to retain its basic capabilities even if it is exposed to a combination of adverse factors. Technical solutions, on the other hand, are likely to have serious problems if confronted with only a single type of distortion, apart from the fundamental difficulties to supply the desired monotonic behaviour at all.

Accordingly, problems of robustness in NLP have almost never been considered from a unifying perspective so far. A number of very specific techniques for some of those different aspects has been developed, which hardly can be related to each other.

Robustness, for instance, is a key issue in speech recognition, where reliable recognition results for a variety of speakers and speaking conditions are desired. Two basic technologies attempt to support this goal

- robust stochastic modelling techniques which are able to capture generalizations across the individual variety[1] and
- sophisticated search procedures which select among huge amounts of competing recognition hypotheses by comparing probability estimations for signal segments of increasing length.

Special signal enhancement techniques are used to suppress stationary environmental noise. There are other aspects of robustness which even have not been treated at all, including the flexible adaptation to external time constraints or internal resource limitations.

Traditionally the notion of robustness has been strongly connected to the processing of ill-formed input[2], where ill-formedness can be defined both, in terms of human standards of grammaticality or in terms of unexpected input. Most of the work has been concerned with the problem from a *purely syntactic point of view* and usually relied on two basic techniques: error anticipation and constraint relaxation.

Error anticipation identifies a number of common mistakes and tries to integrate them into the existing grammar by devising dedicated extensions to its coverage. Therefore, the method is limited to a few selected types of deviant constructions which are notorious and therefore predictable, namely

- stereotypical spelling mistakes (**comittee, *rigth*, etc.),
- performance phenomena in spoken language, like restarts (cf. [6]) and
- interference-based competence errors in early phases of second language learning (cf. [1]).

Obviously, the complete "innovative" potential and the individual creativity for producing ill-formed input cannot be adequately captured by such means alone.

On the other hand, *constraint relaxation techniques* rely on a systematic variation of existing grammar rules written for standard input. Initially, the idea was restricted to the stepwise retraction of e.g. agreement conditions in syntactic rules. It can easily be extended to incorporate arbitrary rule transformations in order to allow for the insertion, deletion, substitution and transposition of

[1] The difficulties with a straightforward generalization of this approach to e.g. syntactic or semantic anomalies are obvious: It would require huge amounts of sufficiently deviant utterances being available as training data. This renders the approach technically infeasible and cognitively implausible.
For similar reasons connectionist approaches are not considered here: At the moment they seem to be limited to approximate solutions for flat representations (cf. [27]).

[2] For a good overview see [25].

elements. The difference vanishes completely within modern constraint-based formalisms [26] [2], where a transposition of constituents can be interpreted equally well as a relaxation of linear precedence constraints. Furthermore, constraints can be annotated by their degree of vulnerability, hence allowing to include aspects of error anticipation into the relaxation framework.

Since both, error anticipation and constraint relaxation considerably enlarge the generative capacity of the original grammar they will lead to spurious ambiguities and serious search problems. This restricts their application to a kind of *post mortem* analysis[3]. Only if a failure of the standard analysis procedure indicates the presence of non-standard input, error rules or relaxation techniques are activated to integrate the fragmentary results obtained so far.

Even a superficial comparison with human processing principles shows the fundamental deficit of these approaches. A human reader or listener accepts ill-formed input to a wide degree, often without noticing an error at all. This is particularly true if strong expectations concerning the content of the utterance are involved or if heavy time constraints restrict the processing depth.

Obviously, there is a fundamental parallelism between robustness issues and time considerations, which syntactically oriented solutions lack so far. Robustness in human language processing does not amount to an additional effort, but instead facilitates both, insensitivity to ill-formed input as well as a flexible adaptation to temporal restrictions.

This basic pattern is much better modelled by *semantically oriented approaches* based on the slot-and-filler-principle. Here, highly domain specific expectations are coded by means of frame-like structures and checked against the input for satisfaction. The schema can be successfully extended to a kind of skimming understanding bringing together the question of robustness against syntactically ill-formed input and some simple considerations concerning resource limitations.

This advantage of a semantically guided analysis, however, is won by the cost of excluding another important robustness feature, namely the ability to cope with unexpected input (e.g. a change of topic beyond the narrow limitations of the domain or the violation of selectional restrictions in metaphorical expressions).

2 Observations from Human Language Processing

Psycholinguistic evidence provides a contradictory picture of human language processing. Some observations clearly support a rather strong modular organization with processing units of great autonomy like syntax and semantics [4] [5]. On the other hand there is a considerable semantic influence on the assignment of syntactic structure [20] which suggests a highly integrated processing architecture.

[3] There are exceptions to every rule: For language learning purposes [17] propose an initial analysis based on a moderately weak grammar and followed by a more rigid second pass.

Robust behaviour in natural language understanding seems to require both,

- the autonomy between parallel lines of processing which embodies redundancy and allows to compensate partial insufficiencies and
- the interactive nature of informational exchange which allows to relate partial structures on different levels of granularity.

Functional autonomy undoubtedly is of fundamental importance for robustness. It allows to yield an at least vague interpretation even in cases of extremely distorted input:

1. A semantically almost empty sentence can be analysed quite well by syntactic means alone, delivering a hypothetical interpretation in terms of a possible world with highly underspecified referential object descriptions and possibly ambiguous thematic roles.

   "... und grausig gutzt der Golz."[23] (1)

2. Syntactically ill-formed utterances are interpreted based on semantic and background knowledge even if subcategorization regularities or other grammatical constraints are violated.

Although both processing units are – at least partially – able to generate some useful interpretation independently of the other one, best results, of course, are to be expected if they combine their efforts in a systematic way.

Parallel and autonomous structures in language processing have not only evolved between syntactic and semantic aspects of language. They can be observed equally well at the level of speech comprehension where auditory (hearing) and visual (lip-reading) clues are usually combined to achieve a reliable recognition result. Again, both systems – in principle – are able to work independently, but synergy occurs if both are activated concurrently.

A second group of observations related to the question of robustness concerns the *expectation-driven* nature of human language understanding. Here, expectations come to play at two different dimensions:

- Syntactic, semantic and pragmatic predictions about future input derived from previous parts of the utterance or dialogue.
- Expectations exchanged between parallel and autonomous processing structures for syntax and semantics.

The role of dynamic expectations has mostly been investigated from the viewpoint of a possible search space reduction in prediction based parsing strategies (namely left or head corner algorithms). If used to select between competing hypotheses in speech recognition the predictive capacity of a grammar can contribute additionally to an enhanced robustness of the overall system [10] [7].

Although the importance of predictions for robustness is beyond question, here the second type of expectations shall be examined as a matter of priority, since they are expected to establish the attempted informational coupling between parallel processing units. As the simple examples above have shown, no predefined direction for this exchange of information can be assumed. Certain syntactic constructions may trigger specific semantic interpretations, a view which is strongly supported by the traditional perspective on the relation between syntax and semantics. In the opposite direction, semantic relations, e.g. derived from background knowledge, can not only be used to disambiguate between preestablished syntactic readings, but moreover are able to actively propose suitable syntactic structures. This bidirectionality of interaction seems to be of great importance for the ability to provide the mutual compensation necessary to treat deviant constructions of different kind.

Of course, the expectation-based nature of natural language processing cannot guarantee a failure-proof performance under all circumstances[4]. There certainly are situations in which strong expectations may override even sensory data. Such a situation can easily be studied in everyday conversation whenever e.g. pragmatic expectations are predominant. A similar problem occurs in experimental settings using intentionally desynchronised video input, where lip reading information sometimes overrides even the auditory stimulus. The problem is witnessed as well by the difficulties usually encountered in proof-reading one's own text: Extremely strong expectations concerning the content usually cause minor mistakes to be passed unnoticed.

Typically, expectations are contradictory and will be of different impact on the progress of the analysis procedure. Hence, there is a third principle of robust language processing upon which the human model builds. It concerns the *preference-based selection* between both, competing interpretations as well as different expectations [12]. Expectations have to be ranked according to their particular strength and weighted against each other.

Recently linguistic research has shown a remarkable trend towards the development of integrated models of language structure. One of the more popular examples surely is Head-Driven Phrase Structure Grammar (HPSG [24]), where syntactic and semantic descriptions are uniquely related to each other by coreferential pointers within the framework of typed feature structures. The strong coupling on the level of representation *and* on the level of processing (i.e. within unification) completely lacks autonomy. The construction of a logical form is always mediated by syntactic descriptions taken e.g. from subcategorization information. Since syntactic and semantic restrictions are conjunctively combined the overall vulnerability against arbitrary impairment of the input utterances even increases: An analysis may now fail due to syntactic as well as due to semantic reasons.

A quite similar conclusion can be drawn for construction grammar [3], another integrated approach. It combines syntactic, semantic and even pragmatic information in a single representation named construction. Again, autonomy of individual description levels is missing and even if constructions are supplied with preferential weightings derived from their frequency of use (as realized in

[4] Note that perfect performance is not necessarily covered by the informal notion of robustness introduced earlier.

SAL [13]) robustness does not increase.

A clearcut separation of representational levels has actually been realized in the cognitively motivated parser COMPERE [11] [18]. The system aims at modelling error recovery techniques for garden-path sentences. It uses an arbitration mechanism to decide in case of a conflict situation which alternative reading should be backed up. This allows to combine early commitment decisions with the possibility to switch to another interpretation if necessary later on. Although the parser is guided in its decisions by different kinds of preferences, the mapping between syntactic and semantic representations seems to be a strict one. Accordingly, it does not provide the necessary means for conflict resolution in all those cases of non-standard input for which no interpretation can be established. In particular, three different cases can be distinguished

1. failure on a single level (syntax or semantics)
2. failure on both levels (syntax and semantics)
3. no consistent mapping between levels

Whereas the first case might be easily accommodated by the arbitration mechanism the latter two require the abandonment of the strict mapping and its replacement by a preference-based module interaction.

3 Disambiguation by Constraint Propagation

A suitable combination of the three principles discussed above might in fact provide the foundation for an effective use of redundancy in parallel processing structures

- *autonomy* guarantees a fall-back behaviour for failures of a single module
- *expectancy-oriented analysis* facilitates the informational exchange and
- *preference-based processing* guides the analysis towards a promising interpretation and establishes a loose coupling between modules.

These principles, even if taken together, do not explain the almost unconscious treatment of errors in everyday communication. To simulate a similar behaviour a *selective constraint invocation* strategy will become necessary. Then, parsing is understood as a disambiguation procedure, which activates only specific parts of the grammar, if this is deemed to be unavoidable for solving a particular disambiguation problem. The procedure can be terminated if a sufficiently reliable disambiguation has been achieved even if certain conditions of the grammar have never been checked so far. Robustness is not introduced by a *post mortem* retraction of constraints but rather by their careful invocation.

Along these lines a rudimentary kind of robustness has been achieved in the Constraint Grammar framework [15], a system for parsing large amounts of unrestricted text. Constraint Grammar (CG) attempts to establish a dependency description which is underspecified with respect to the precise identity of modifiees. Initially, it assigns a set of morphologically justified syntactic labels to each word form in the input sentence. Possible labels among others are

@+FMAINV	the finite verb of a clause
@SUBJ	a grammatical subject
@OBJ	a direct object
@DN>	a determiner modifying a noun to the right
@NN>	a noun modifying a noun to the right

The initial set of labels is successively reduced by applying compatibility and surface ordering constraints until a unique interpretation has been reached or the set of available constraints is exhausted. In the latter case, a total disambiguation cannot be achieved by purely syntactic means, as in the following attachment example:

Bill	*saw*	*the*	*little*	*dog*	*in*	*the*	*park*
@SUBJ	@+FMAINV	@DN>	@AN>	@OBJ	@<NOM	@DN>	@<P
					@<ADVL		

In contrast to traditional grammars of the phrase structure type which license well-formed structures according to their rule system, constraint grammar rather happens to be an *eliminative approach*[5]. Instead of imposing a normative description on the input data it takes them as starting point and tries to find a plausible interpretation for them.

This proceeding is motivated by the finding that language is an open-ended system and so grammar formalisms based on a "rigid and idealized conception of grammatical correctness are bound to leak" [14, p. 37]. Parsing, if understood as a disambiguation procedure, is put down to the principle of parsimony: The more effort is spent the better disambiguation results can be expected and [14, p. 39] points to the important psycholinguistic parallel:

> "Mental effort is needed for achieving clarity, precision and maximal information. Less efforts imply (retention of) unclarity and ambiguity, i.e. information decrease. In several types of parsers, rule applications create rather than discard ambiguities: the more processing, the less unambiguous information."

Parsing as disambiguation can well be extended to deal with fully specified dependency structures without loosing its promising characteristics. A complete disambiguation of structural descriptions has first been described for Constraint Dependency Grammar (CDG [21]) and simply requires to replace the monadic categories of CG by pairs consisting of the relation name and the exactly specified modifiee. The mutual compatibility of modifying relations is checked against a set of constraints and thus the set of possible modifications is successively reduced by a constraint propagation mechanism. Further extensions of the approach concern the inclusion of feature descriptions, valency specifications and valency saturation conditions [8].

[5] With this respect a strong parallel between the eliminative nature of disambiguation and cohort modelling ideas for spoken word recognition [19] becomes visible.

Though, a closer inspection of the kind of robustness feature introduced by the eliminative mode of operation reveals that its nature is quite accidental so far. Which types of deviation can be tolerated indeed, strongly depends on the rather arbitrary sequence of constraint applications. This shortcoming seems to be closely connected to the fact that both formalisms lack the notion of preference so far and therefore do not have the possibility to model the "quality" of a constraint[6]. Hence, adding a preference-based selection strategy will be one of the most pressing needs for further improvement. Such an extension will be proposed in section 5. Before we turn to this topic section 4 introduces a modular representation schema along the traditional syntax-semantics distinction. It supports the desired functional autonomy as well as a highly interactive exchange of expectations between the two layers.

4 Representation Layers

Whereas Constraint Grammar restricts itself to purely syntactic means, an integration of simple semantic criteria into Constraint Dependency Grammar has been proposed recently [9]. It takes into account sortal restrictions only, attaching them to surface syntactic relations without aiming at modularity and autonomous behavior. In order to facilitate functional independence it will become necessary to establish separate layers for structural description *and* constraint propagation

- a *syntactic layer* relating word forms according to functional surface structure notions e.g (SUBJECT-OF, DIR-OBJECT-OF, PREP-MODIFIER-OF, etc.) and using constraints on ordering, agreement, valency, and valency saturation to select among competing structural configuration and
- a *semantic layer* building sentence structures by means of thematic roles (like AGENT-OF, INSTRUMENT-OF, TIME-OF, etc.) thereby relying upon the argument structure of semantic predicates and their corresponding selectional restrictions.[7]

The following small and rather rigid sample grammar illustrates the different types of constraints needed:

1. licensing conditions for modification relations

   `sy1: cat(dep(X))=N`
   $\quad\quad\rightarrow$ `cat(synmod(X))=V` $\wedge$ `synlab(X)`$\in\{$`SUBJ,OBJ`$\}$ [8]

   A noun can modify a verb either as a subject or as a direct object.

[6] CG at least includes heuristic constraints, which may be activated at a particular stage of the disambiguation procedure

[7] The proposed separation quite closely corresponds with the one chosen in [11]

[8] `dep(X)` refers to the modifier of a relation, `syndom(X)` and `semdom(X)` to its modifiees. `synlab(X)` and `semlab(X)` are the respective relation names. `cat(X)`, `num(X)`, `semprop(X)`, ... denote properties of the corresponding node.

2. agreement conditions

   `sy2: synlab(X)=SUBJ → num(dep(X))=num(syndom(X))`

   A subject agrees with its modifiee with respect to number.

3. linear ordering constraints

   `sy3: synlab(X)=SUBJ → pos(dep(X))<pos(syndom(X))`

   The subject precedes the finite verb.

4. compatibility constraints[9]

   `sy4: syndom(X)=syndom(Y) → synlab(X) ≠ synlab(Y)`

   A word form cannot be modified twice by the same relation.

`sy1` through `sy3` are unary constraints, `sy4` is a binary one. Note that constraints refer to modifying relations instead of word forms. Therefore they are able to express admissibility conditions on local configurations consisting of up to three nodes. Note as well that `sy3` – even for German main clauses – has a strong heuristic appearance and simply states a preference condition which additionally requires a suitable exception handling mechanism.

In a very similar fashion semantic constraints comprise

1. licensing conditions

   `se1: cat(dep(X))=N → cat(semdom(X))=V ∧ semlab(X)∈{AG,PAT}`

2. selectional restrictions

   `se2: word(semdom(X))=fressen ∧ semprop(mod(X))=animal`
   `→ semlab(X)=AG`

   Animals do eat.

   `se3: word(semdom(X))=fressen ∧ semprop(mod(X))=plant`
   `→ semlab(X)=PAT`

   Plants are to be eaten.

3. compatibility constraints[10]

   `se4: semdom(X)=semdom(Y) → semlab(X)≠semlab(Y)`

Adhering to the principle of autonomy both layers are designed in a way which allows them to propagate constraints in a completely independent manner. Each modifier is specified for two possibly different modifiees and no cross-reference between the layers has been used so far.

In order to finally mediate the interaction between layers, a set of mapping constraints has to be provided which sets up bidirectional correspondences

---

[9] In fact there is another general compatibility constraint implicitly built into the decision procedure and excluding ambiguous modifying relations from being consistent:
`sygen: ¬ (syndom(X)=syndom(Y) ↔ dep(X)=dep(Y))`

[10] Again, supplemented by a general semantic uniqueness constraint
`segen: ¬ (semdom(X)=semdom(Y) ↔ dep(X)=dep(Y))`

ss1: `syndom(X)=semdom(X)` $\rightarrow$ `(synlab(X)=SUBJ` $\leftrightarrow$ `semlab=AG)`

The subject of a verb is always identical to its agens.

ss2: `syndom(X)=semdom(X)` $\rightarrow$ `(synlab(X)=OBJ` $\leftrightarrow$ `semlab=PAT)`

The direct object of a verb is always identical to its patiens.

It should have become obvious that the selectional restrictions as well as the mapping constraints at best can be taken to stand for a preferential interpretation. They surely are much to rigid to be sensibly used within a framework of strict reasoning.

Semantic constraints need not be restricted to linguistically motivated (i.e. universally valid) ones. In particular, domain-specific restrictions play a crucial role in semantic disambiguation and should urgently be incorporated whenever possible. Here the semantic layer offers a convenient interface to a knowledge representation component which (on demand) can contribute constraints from e.g. specialized ontologies, referential instantiations or temporal reasoning.

5 Weakening Constraints

So far, one of the most striking shortcomings has been the strictly binary nature of constraint satisfaction. Not surprisingly, it turned out to be most inappropriate within the area of semantic modelling where hardly a constraint can be formulated without restricting oneself to a particular, preferential reading.

In what follows, preferences are not modelled in the usual direct manner by emphasizing particular well-formed interpretations but rather indirectly by putting a *penalty* on all remaining alternatives which violate a constraint. For this purpose each constraint gets a penalty factor `pf` assigned reducing the confidence score in negative cases. Penalty factors may range from zero to one where

`pf=0`	specifies a strict constraint in the classical sense and
`0<pf<1`	indicates a soft constraint accepting contradictory cases with a confidence value proportional to `pf`

Obviously, a value of one is meaningless because it neutralizes the constraint. Penalty factors are combined multiplicatively, i.e. compatibility matrices within the constraint satisfaction problem no longer contain binary categories but *confidence scores* also ranging from zero (for impossible combinations) up to one (for combinations not even violating a single constraint).

The indirect treatment of preference by penalty factors offers a consistent extension to the basic paradigm of constraint satisfaction. It does not sacrifice the eliminative nature of constraints but simply softens it. Inappropriate readings are excluded only if they violate strict constraints. In all other cases they are downgraded to a certain degree.

In particular, the penalty-based approach helps to tackle some normalization problems otherwise inherently connected with the constraint satisfaction approach: Most modifying relations (or combinations of them) will pass a constraint simply because it is irrelevant for that particular configuration. An increase of goodness estimates for these cases would yield a highly undesirable, since unjustified reinforcement.

By assigning the penalty factors `pf(sy1)=pf(sy4)=0` to the constraints `sy1` and `sy4` from section 3 both are declared to be strict ones, a fact obviously being valid for the toy-size sample grammar which does not take into account coordinative structures. Using `pf(sy2)=0.1` the agreement condition is treated as a rather strong one which allows exceptions only occasionally. `pf(sy3)=0.3` on the other hand results in a much more permissive constraint justified by the fact that `sy2` is meant to exclude ungrammatical utterances but `sy3` only to disfavour a marked ordering.

On the semantic layer only the licensing constraint `se1` is declared as a strict one. The compatibility constraint `se4` is weakened considerably in order to account for double modification as in the case of anaphoric reference. The two selectional constraints receive penalty factors of different strength in order to model the lower probability of a plant eating something (`se2`) compared to an animal being eaten (`se3`):

```
pf(se1)=0.0
pf(se2)=0.1
pf(se3)=0.7
pf(se4)=0.5
```

The mapping constraints, finally, are weighted in a way which strongly favours the SUBJECT-AGENS and OBJECT-PATIENS pairings nevertheless allowing alternative interpretations e.g. in passive sentences.

```
pf(ss1)=0.2
pf(ss2)=0.3
```

Alternative readings can but need not be specified explicitly. In more realistic applications though it is recommended to aim at a considerably richer modelling, otherwise an unbalanced penalty factor as between `se2` and `se3` may create a sometimes undesired strong bias by default: If both constraints appear to be of no relevance the patience reading is clearly prioritized. In the example chosen this corresponds to the acceptable interpretation that an arbitrary thing is more likely to be eaten than to eat.

After having introduced penalty factors as a means of modelling preferences constraint propagation can be extended from the classical case of strictly binary decisions to the handling of confidence scores. The application of penalty-weighted constraints to a disambiguation problem now consists of two steps:

1. the calculation of initial confidence scores for all combinations of syntactic and semantic modification relations and
2. a selection procedure pruning the search space by sorting out unlikely interpretations

The selection procedure is based on a local assessment function heuristically identifying relations to be pruned. In order to not select promising hypotheses assessment first of all should only take into account modification relations characterized by the following three criteria

– being close to the global minimum for all modification relations, combined with
– an as possible as high contrast to alternative relations and
– a low contrast between all the confidence scores supporting the relation in question.

For experimental purposes a selection procedure based on the sum of quadratic errors for setting scores to zero has been used. Hence, structural interpretations violating a high number of rather strong constraints are pruned first.

Using the toy grammar specified above together with its penalty scores the arbitration process between syntactic and semantic evidence in simple disambiguation problems can be studied. Thus in a sentence like

Pferde fressen Gras. *(Horses eat grass.)* (2a)

both layers uniformly support a single interpretation

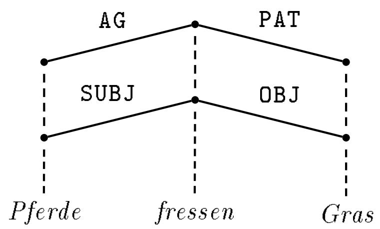

Due to the strong semantic support the interpretation remains unchanged if a marked ordering (topicalization of the direct object) is chosen (2b), an agreement error is introduced (2c) and both deviations are combined finally (2d).

$Gras_{PAT}$ *fressen* $Pferde_{AG}$. (2b)
$Pferd_{AG}$ *fressen* $Gras_{PAT}$. (2c)
$Gras_{PAT}$ *fressen* $Pferd_{AG}$. (2d)

The interpretation is retained even if its semantic support is neutralized as in the following utterance, containing a twofold type shift.

$Autos_{AG}$ *fressen* $Geld_{PAT}$. *(Cars eat money.)* (3a)
$Geld_{PAT}$ *fressen* $Autos_{AG}$. (3b)
$Auto_{AG}$ *fressen* $Geld_{PAT}$. (3c)

It switches to the alternative interpretation only in the case of combined syntactic distortions

$Geld_{AG}$ *fressen* $Auto_{PAT}$. (3d)

Even for the counterintuitive example

$Gräser_{AG}$ *fressen* $Pferd_{PAT}$. (4a)

which, if desired, could be taken as a headline-style utterance, syntactic evidence will gain the upper hand against the violation of two selectional constraints. This interpretation, however, happens to be a rather fragile one and breaks immediately under arbitrary syntactic variation.

Since the selection procedure operates on a global assessment of local structural configurations it cannot guarantee to find an optimal and globally consistent interpretation. The partially local mode of operation, on the other hand, can be expected to provide a quite natural explanation for human garden-path phenomena. Within the framework of preference-based disambiguation they turn out to be a special case of contradictory situations which manifest themselves as *expectation violations*: The consequences of a pruning decision may not coincide with local confidence estimations elsewhere in the constraint network.

Expectation violations not necessarily do indicate an erroneous situation. They are frequently used as a speaker's intentionally chosen means to attract the attention of the audience. This happens for instance by deviating from an unmarked ordering to emphasize a topicalized constituent (c.f. (3b)) or by otherwise producing unexpected utterances.

On the other hand there are the typical erroneous situations which in case of

- internal difficulties (e.g. due to early commitment strategies in garden path situations) might offer the possibility to initiate a reanalysis and
- external reasons (e.g. ill-formed input) can be used to track down the error to find a possible remedy for it.

Note that in the latter case the situation coincides with basic observations for the human model: Finding an interpretation for erroneous utterances will be easier – in terms of effort to spent – than detecting the error, which in turn will be less demanding than localizing or even correcting it.

As with human language processing there will be no predefined direction for the general flow of expectations during arbitration. Whether syntactic evidence is propagated from the syntactic to the semantic layer or vice versa depends only on the available information. This seems to be in accordance with recent psycholinguistic findings which contest the existence of purely structural disambiguation principles [16].

6 Preference-based Reasoning

Eliminating implausible interpretations by locally pruning less favoured modification relations represents only one, though fundamental method for the disambiguation of natural language utterances. By selecting among modifying relations according to negative evidence from maximally dispreferred hypotheses, the technique fits quite well into the constraint satisfaction approach and achieves its robust behaviour by avoiding extremely risky decisions on a locally topmost reading. Taking this as a starting point the basic way of reasoning can well be complemented by a second propagation principle based on preference-induced constraints. These are activated only in situations where enough positive evidence can be derived from almost uniquely determined preferences. Since the existence of convincing preferences in realistic disambiguation tasks represents rather the exception than the rule the nature of this propagation principle is secondary.

Preference-induced constraints consist of implications $P \rightarrow_p C$ which, given enough evidence for the unary precondition P, require the possibly binary constraint C to hold. Constraints of this type can be used to model e.g. the higher-order conditions which in many mapping situations involve more than two relations.

pss1: $\texttt{word(syndom(X))=im} \wedge \texttt{synlab(X)=PHEAD}$
$\quad\quad \wedge\ \texttt{semprop(dep(X))} \in \{\texttt{TEMP,LOC}\}$
$\quad \rightarrow_p \texttt{synlab(Y)=PMOD} \wedge \texttt{dep(Y)=syndom(X)}$
$\quad\quad \wedge\ \texttt{semdom(Z)=sydom(X)} \rightarrow \texttt{semlab(Z)=PART-OF}$

A prepositional phrase headed by the word form "im" fills a semantic PART-OF *slot*

In postnominal positions like

Dann nehmen wir die erste Woche im Mai. (5)
(Let's take the first week in may.)

this constraint puts a preference on the lower attachment since a part-of relation usually is not licensed as an argument position for verbs.

Preference-induced constraints can also be used to modify value assignments at certain nodes in the constraint network without the necessity to copy them. By applying this technique, phrasal feature projections can be modelled in order to build up descriptions for partial dependency trees

pss2: $\texttt{synlab(X)=DET}$
$\quad \rightarrow_p \texttt{case(syndom(X)):=case(syndom(X))} \cap \texttt{case(dep(X))}$

A noun group carries the intersection of possible case features found at its members.

Preference-induced constraints introduce a kind of inhibitory mechanism to the disambiguation procedure: Already preferred interpretations is given the chance to propagate their consequences over the network thus possibly leading to further suppression of alternative readings.

7 Conclusion

Combining the eliminative nature of a disambiguation procedure with a system architecture supporting bidirectional arbitration between syntactic and semantic evidence has turned out to be a key factor for achieving a higher level of robustness in language understanding. While the disambiguation paradigm provided the basic fall-back behaviour (an arbitrary utterance will get a description assigned) and the possibility to prune the search space towards a least disfavoured reading, the parallel arrangement of modules allows to interactively exchange expectations and thus bypassing local interpretation difficulties. By modelling a preference distribution based on penalty factors, the desired robust behaviour can be demonstrated at least for very simple sample utterances. Although no conclusive judgement about the feasibility of the approach can be given until the experimental setting has been scaled up to a fairly realistic problem size, a remarkable qualitative advance over comparable approaches becomes evident even on this elementary level:

- The approach departs from a predefined sequential arrangement of modules in favour of a strictly symmetrical architecture consisting of autonomous components for syntax and semantics.
- It allows to treat syntactic ill-formedness and semantic deviations by providing a mechanism for mutual compensation. Syntactically anomalous utterances can be understood as long as there is enough semantic and/or pragmatic evidence. In order to communicate novel or unusual content a sufficiently high degree of syntactic support is required.
- Insufficient modelling information on any one of the processing layers might well result in the selection of an odd interpretation but will not cause the language processing unit to break down entirely.
- Robustness is not an add-on feature of an otherwise temperamental procedure but falls out from the basic properties of the processing mechanism.

Since structural disambiguation by constraint satisfaction likewise lends itself to the creation of time sensitive parsing procedures [22], in the long run it might provide a unifying foundation to build language processing systems upon which embody aspects of robustness against such different disruptive factors as syntactically ill-formed input, metaphorical use and dynamic time constraints.

References

1. M. E. Catt. Intelligent diagnosis of ungrammaticality in computer-assisted language instruction. Technical Report CSRI-218, Computer Systems Research Institute, University of Toronto, 1988.
2. G. Erbach. Towards a theory of degrees of grammaticality. Report 34, Universität des Saarlandes, Computerlinguistik, Saarbrücken, 1993.
3. C. J. Fillmore, P. Kay, and M. C. O'Connor. Regularity and idiomaticity in grammatical constructions: The case of let alone. *Language*, 64:501–538, 1988.
4. K. I. Forster. Levels of processing and the structure of the of the language processor. In W. E. Cooper and E. C. T. Walker, eds., *Sentence Processing: Psycholinguistic studies presented to Merrill Garret*, p. 27–85. Lawrence Erlbaum, Hillsdale, NJ, 1979.
5. L. Frazier. Theories of sentence processing. In J. L. Garfield, ed., *Modularity in Knowledge Representation and Natural Language Understanding*, p. 291–307. MIT-Press, Cambridge MA, 1987.
6. S. Goeser. Chart parsing of robust grammars. In *Proc. 14th Int. Conf. on Computational Linguistics, Coling '92*, p. 120–126, Nancy, France, 1992.
7. G. Hanrieder and G. Görz. Robust parsing of spoken dialogue using contextual knowledge and recognition probabilities. In *Proc. of the ESCA Workshop on Spoken Dialogue Systems*, Denmark, 1995.

8. M. P. Harper and R. 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.
9. M. P. Harper, L. H. Jamieson, C. B. Zoltowski, L. L. McPheters, and R. A. Helzerman. Semantics and constraint parsing. In *Proc. of the Int. Conf. on Acoustics, Speech, and Signal Processing, ICASSP-92*, p. II–63–66, 1992.
10. A. Hauenstein and H. Weber. An investigation of tightly coupled time synchronous speech language interfaces using a unification grammar. Verbmobil Report 9, 1994.
11. J. K. Holbrook, K. P. Eiselt, and K. Mahesh. A unified process model of syntactic and semantic error recovery in sentence understanding. In *Proc. of the 14th Annual Conf. of the Cognitive Science Society*, p. 195–200, Bloomington, IN, 1992.
12. R. S. Jackendoff. *Semantics and Cognition*. MIT Press, Cambridge, 1983.
13. D. Jurafsky. A cognitive model of sentence interpretation: The construction grammar approach. Technical Report TR-93-077, International Computer Science Institute, Berkeley, 1993.
14. F. Karlsson. Designing a parser for unrestricted text. In F. Karlsson, A. Voutilainen, J. Heikkilä, and A. Anttila, eds., *Constraint Grammar – A Language-Independent System for Parsing Unrestricted Text*, p. 1–40. Mouton de Gruyter, Berlin, New York, 1995.
15. F. Karlsson, A. Voutilainen, J. Heikkilä, and A. Anttila, eds. *Constraint Grammar – A Language-Independent System for Parsing Unrestricted Text*. Mouton de Gruyter, Berlin, New York, 1995.
16. L. Konieczny, B. Hemforth, and N. Voelker. The impact of context and semantic bias on constituent attachment in reading. In J. J. Quantz and B. Schmitz, eds., *Ambiguity and Strategies of Disambiguation*, p. 105–126. KIT-Report 120, Berlin, 1994.
17. I. Kudo, H. Koshino, M. Chung, and T. Morimoto. Schema method: A framework for correcting grammatically ill-formed input. In *Proc. 12th Int. Conf. on Computational Linguistics, Coling '88*, p. 341–347, Budapest, 1988.
18. K. Mahesh and K. P. Eiselt. Uniform representations for syntax-semantic arbitration. In *Proc. of the 16th Annual Conf. of the Cognitive Science Society*, p. 589–594, 1994.
19. W. Marslen-Wilson. Functional parallelism in spoken word-recognition. *Cognition*, 25:71–102, 1987.
20. W. Marslen-Wilson and L. K. Tyler. Against modularity. In J. L. Garfield, ed., *Modularity in Knowledge Representation and Natural Language Understanding*, p. 37–62. MIT-Press, Cambridge, MA, 1987.
21. H. Maruyama. Structural disambiguation with constraint propagation. In *Proc. 28th Annual Meeting of the ACL*, p. 31–38, 1990.
22. W. Menzel. Parsing of spoken language under time constraints. In T. Cohn, ed., *Proc. 11th Europ. Conf. on Artificial Intelligence*, p. 560–564, Amsterdam, 1994.
23. C. Morgenstern. *Alle Galgenlieder*. Cassirer, Berlin, 1933.
24. C. Pollard and I. A. Sag. *Head-Driven Phrase Structure Grammar*. The University of Chicago Press, Chicago, 1994.
25. M. Stede. The search for robustness in natural language understanding. *Artificial Intelligence Review*, 6(4):383–414, 1992.
26. H. Uszkoreit. Strategies for adding control information to declarative grammars. Computerlinguistik Report 10, Universität des Saarlandes, Saarbrücken, 1991.
27. S. Wermter and V. Weber. Learning Fault-tolerant Speech Parsing with SCREEN. In *Proc. of the 12th Nat. Conf. on Artificial Intelligence*, p. 670–675, Seattle, 1994.