ABDUCTIVE REASONING FOR SYNTACTIC REALIZATION*

Ralf Klabunde
University of Heidelberg
Centre for Computational Linguistics
klabunde@novell1.gs.uni-heidelberg.de

Martin Jansche
The Ohio State University
Department of Linguistics
Jansche.1@osu.edu

Abstract

Abductive reasoning is used in a bidirectional framework for syntactic realization and semantic interpretation. The use of the framework is illustrated in a case study of sentence generation, where different syntactic forms are generated depending on the status of discourse information. Examples are given involving three different syntactic constructions in German root clauses.

1 Pragmatics in Natural Language Generation

The computational treatment of pragmatics in natural language generation is often—directly or indirectly—oriented around the Gricean maxims [Grice 75]. Their effects emerge from the pragmatic model of the generation system so that the generated texts satisfy these maxims. The texts should be a true characterization of a state of affairs, they should be as informative as possible, relevant, and perspicuous. While the first three maxims are related to what is said, the last maxim is related to how it is said. The category of perspicuity principles includes constraints on avoiding obscurity and ambiguity, or being brief and orderly. It is anything but clear how these principles should be interpreted precisely. Several attempts have been made to remedy this in computational work on generating texts that best satisfy these maxims, especially with respect to the generation of referring expressions (e.g. [Dale et al. 95]).

However, there is more to pragmatics than satisfying Gricean maxims. In particular, the category of perspicuity principles does not usually cover the important fact that texts are tailored to a specific addressee, not only in content, i.e., with respect to her or his informational needs, but also in the linguistic form, i.e., word order, syntactic constructions, the choice of lexical items, and eventually prosodic information. This tailoring of the linguistic form to the listener is termed "information structuring". In generating texts, information structuring requires, among other things, the use of some listener model, which may include information about the listener's knowledge, goals, properties, etc.

Linguistic approaches to describing the principles of information structuring have sometimes characterized information structure as an *instruction* to the listener about how to construct a model of the communicated state of affairs [Prince 81]. In AI and Computational Linguistics, tailoring the message to the user comprises very often solely content planning, which only indirectly determines the linguistic output. For example, systems tailor the information "density" to the user (e.g. [Paris 93]), or they drive the dialogue depending on an estimation of what the user might be interested in (e.g. [Jameson et al. 94]). Realizing texts by determining the information structure of the respective sentences, which again is a reflex of addressee orientation, has not yet received its due attention.

*The authors would like to thank Bob Kasper, Nathan Vaillette, Shravan Vasishth, and two anonymous referees for helpful comments and suggestions. All remaining mistakes are, of course, our own.

2 The Topic/Comment Structure in Information Structuring

The notion of information structure comprises at least two separate notions of how the information of a sentence may be structured, viz. the topic/comment structure and the focus/background structure [Vallduví 92, Lambrecht 94].[1] In order to motivate these structuring mechanisms, consider the following simple example. Suppose the purpose of a generation system is to describe a spatial scenario. One of the sentences might be

(1) Behind the town hall is a BAKERY.

with "bakery" the prosodically most prominent constituent (the *focus exponent*). In this sentence, the prepositional phrase "behind the town hall" functions as topic and the noun phrase "a bakery" is in focus. we will ignore aspects of focus and its role in language generation, especially since selecting the focus exponent is better understood as being part of utterance planning, and are limiting our attention to the topic/comment-structure only. The topic provides familiar discourse referents whose properties are further illuminated by the sentence; the relation between these discourse referents and the sentential predication is also referred to as an aboutness-relation. Many languages possess special topicalization constructions or morphological markers to single out the topic in a sentence. In German (and probably English as well), referring phrases provide topic referents, and the clause-initial position is their preferred position. Thus clause-initial positioning is the most important topic-relevant feature in generation.[2]

The same propositional content expressed by (1) can be realized with different information structures and, therefore, different sentence forms, as the following English examples demonstrate:

(2) A bakery is behind the town hall.

(3) Behind the town hall, there is a bakery.

(4) As for the town hall, behind it is a bakery.

Discourse referents functioning as topics must be identifiable for the listener. This is the reason why topics are usually packaged as definite noun phrases, or as prepositional phrases that contain definite noun phrases. Topic candidates will be selected from the set of discourse referents that the listener knows according to a topic acceptance scale. [Lambrecht 94] proposes the following scale:

(5) active > accessible (textually, situationally, or inferentially)
> unused > brand-new anchored > brand-new unanchored

Active referents are those that are currently lit up; they are in the center of attention. They are the most acceptable topics because the listener's mental effort needed for processing the respective sentence is minimal as compared with the effort needed to identify and anchor an unfamiliar or inactive topic referent. We consider the candidates below the accessible referents to be inappropriate as topics in most instances, and we limit our attention to a scale with three regions: active referents, accessible referents, and inaccessible referents.

To summarize, the first task in generating texts with sentences with appropriate topic/comment structures is to determine for each sentence the topic discourse referent. This referent should be identifiable for the listener and as high on the topic acceptance scale as possible. The phrase expressing the topic should

[1] Depending on one's theoretical background and/or affiliation with different schools, terminology differs considerably.

[2] [Altmann 81, 150] gives some counterexamples to this default. These examples are prosodically marked, however.

be placed in clause-initial position. However, these are only guidelines, not fixed rules. Hence, we need a mechanism to handle this kind of uncertainty.

This is of course not the whole story of topic-hood. In addition to selecting topic referents, we have to solve the problem of how one and the same topic/comment structure can be realized by different syntactic structures. German examples resembling the previous three ones are:

(6) Die Vitrine steht rechts von der Lampe.
'The showcase is standing to the right of the lamp.' (subject realization)

(7) Die Vitrine, die steht rechts von der Lampe.
'The showcase, it is standing to the right of the lamp.' (left dislocation)

(8) Was die Vitrine betrifft, die steht rechts von der Lampe.
'As for the showcase, it is standing to the right of the lamp.' (hanging topic)

In the first clause the topic is realized as the subject in clause-initial position. The second clause exhibits a left dislocation for the topic, and the third one uses a so-called hanging topic.

We assume that the functions of these three syntactic forms are more or less identical for German and English. All three examples express the same propositional content, viz. the localization of a uniquely identifiable showcase with respect to a uniquely identifiable lamp. Furthermore, all three examples exhibit the same topic/comment structure: "the showcase" functions as topic, i.e., the anchor for the proposition, and the rest of the clause comments on certain aspects of the showcase. However, these three forms are not mutually interchangeable in each imaginable context, because they invite different pragmatic inferences.

The *subject realization* is neutral with respect to topic accessibility. There is a strong correlation between the grammatical function of subject and the information structural notion of topic. The subject is the unmarked topic.

Left dislocation constructions, they can indicate a topic shift because the syntactically autonomous position of the detached noun phrase signals a change in the status of its discourse referent from being inactive to active [Lambrecht 94]. Additionally, left dislocations must satisfy a presupposition condition, namely to support the existence of another individual not having the property expressed by the matrix clause [Wiltschko 95]. The discourse referent is in some way related to a previously established set which the referent is a member of. This resembles the presuppositions restrictive relative clauses establish.

As for a *hanging topic*, it also indicates a topic shift. It introduces a new topic of the discourse from a set of discourse referents that have already been established in the discourse. The common property of shifting the discourse topic implies that hanging topics and left dislocations are not mutually exclusive. A distinction on pragmatic grounds is complicated by the fact that the various set phrases usable for the hanging topic can have different discourse functions and that left dislocations can be interpreted as special hanging topics. However, the main difference between left dislocations and hanging topics with the set phrase *was das X betrifft* ("as for the X") seems to be: left dislocations must satisfy the presupposition condition and they establish a topic shift by means of changing the status of a discourse referent, whereas hanging topics establish a topic shift by means of selecting a discourse referent from a previously established set of referents.[3] Despite their overlapping discourse functions, we confine ourselves to the distinctive pragmatic properties of both constructions for their generation.

Hence, the second problem that needs to be solved is to correlate the syntactic form with the status of discourse referents with respect to their being active or accessible, as well as with other discourse information and factual information pertaining to the presupposition conditions. How can we incorporate

[3] See the extensive descriptive analyses in [Altmann 81].

this informal characterization of topic, topic acceptability, and syntactic constructions into a unified and formally precise mechanism for a natural language generation system? We propose an abductive setting in the spirit of [Hobbs et al. 93] as a framework for integrating the diverse knowledge sources involved in the generation and interpretation of sentential information structure. The basic idea is to view generating a single proposition as finding the best proof for why a sentence and its information structure is congruent with the listener model. In the process of finding this proof, the sentence is generated by incrementally instantiating unbound variables.

Our basic scenario is the generation of spatial descriptions. The mechanism for content planning is not the subject of this paper (cf. [Jansche et al. 96, Meyer-Klabunde 96, Porzel et al. in press]). For spatial descriptions, content planning comprises for each proposition the selection of a reference object from the set of objects, the selection of a primary object, the selection of a point of view, and the computation of a spatial relation between both objects depending on the chosen point of view. For present purposes we assume that the propositional content of a sentence has already been established. What remains to be done is to construct a pragmatically appropriate sentence that conveys the new and informative part of this propositional information to the listener. It is for this syntactic/pragmatic realization process that we use the abductive framework. Ultimately, we aim to incorporate the abductive reasoning mechanism directly into the content planner so as to achieve a uniform framework.

3 Generation by Abduction

Abductive reasoning is reasoning about the best explanation for a given observation. To make precise what counts as a good explanation, one introduces a preference criterion by which alternative explanations can be compared. A preferred explanation for an observation might be the least specific one, the most specific one, the one with the lowest proof costs, etc. Abductive explanation is classically characterized as follows (cf. [Mayer et al. 96]): a knowledge base K, the usual consequence relation $\models$, and an observation E to be explained, such that $K \not\models E$, are given. A statement H is taken as the best explanation of E in K iff:

1. $K \cup \{H\} \models E$; and
2. H is "better" than any other statement in the set $\{H' \mid K \cup \{H'\} \models E\}$, according to the preference criterion.

We use a generalized version of what [Stickel 90, 236] calls *predicate specific abduction*, where only elements from a distinguished set of literals may be assumed. What counts as the best explanation will be based on the (preferably minimal) number of assumptions made.

Abduction has been used in natural language processing for interpretation tasks such as metonymy resolution, understanding vague expressions, or plan recognition. Recently, abductive reasoning has also been proposed for use in generation, partly for planning [Lascarides et al. 92, Thomason et al. 96], and as a framework for both the interpretation and the generation of discourse [Thomason et al. 97]. The basic idea behind these approaches is to find the best way to obtain a communicative goal state by modifying the conversational record, which roughly corresponds to our listener model, with applicable operators. The plan is the set of hypotheses discovered by an abductive proof of the proposition that the goal state has been achieved.

What remains open in these approaches is to make precise the relation between the planned propositional content of an utterance and an appropriate sentence form. Only very simple example sentences

could be generated because the local pragmatics of the sentence form does not play a role in the previous approaches. We are bridging this pragmatic gap between content planning and surface realization by abductive mechanisms.

3.1 The Abductive Component

For our purposes it is helpful to view abductive proofs as essentially *relational.* An abductive proof determines the relation between a knowledge base, an observation, a specification of what assumptions can be made, and proved and assumed literals that jointly provide an explanation for the initial observation. The prototypes we have implemented in Prolog make this relation available explicitly, and great care was taken to ensure that queries such as (9), where not all arguments are instantiated, are handled correctly by generating a manageable subset of all possible solutions.

(9) `?- abduce(Goal, Assumable, Proved, Assumed).`

In the above query, `Goal` is the observation to be proved by the abductive meta-interpreter, `Assumable` is a set of literals that may be assumed, and `Proved` and `Assumed` are multisets of literals that were used or assumed, respectively, during the abductive proof of `Goal`. Interpretation mode corresponds to queries where the goal is instantiated and everything can be assumed in principle, as in (10); during generation the meta-interpreter is invoked with the goal at least partially uninstantiated, while the set of assumable literals is specified, as in (11).

(10) `?- abduce(sentence([die,vitrine,steht,rechts,von,der,lampe]), _, Pr, As).`

(11) `?- abduce(sentence(S), [showcase(s),lamp(l),loc(s,r),rightof(r,l)], Pr, As).`

From the fact that queries like (11) are accepted it is clear that the abduction scheme we use is somewhat more general than predicate specific abduction: we supply information as to what *literals* may be assumed, whereas predicate specific abduction would only specify the functors and arities of those literals.

It is well worth noting that on our approach generation is not simply the inverse of interpretation. If that were the case, one would call the abductive meta-interpreter with the goal instantiated deriving the assumed literals during interpretation, while for generation the opposite instantiation pattern would be used. But for the latter case this amounts to requiring that all literals *must* be assumed in the proof, which is clearly too strong since some of them might be derivable from the knowledge base. Instead we only specify which literals *may* be assumed, leaving open the possibility that some of them are provable from the knowledge base.

Also note that since we do not use weighted abduction, the problem of assigning different assumption costs for generation and interpretation (cf. [Thomason et al. 97]) is avoided. On the other hand, what should we use as a preference criterion? A sequence of several criteria is used. First, proofs are preferred for the number of provable literals used, the more the better. In the cases we consider, there seems to be a loose correspondence between this criterion and the Gricean maxims of relevance and quality. Second, proofs are preferred compared to other proofs if they involve less assumptions. The number of assumptions made is determined by the cardinality of the set that is the reduction of the multiset found during an abductive proof. Third, everything else being equal, we prefer proofs with the highest amount of assumption re-use. This is determined by the difference between the cardinalities of the multiset of assumptions and of its corresponding set. The relevant idea—an assumption becomes more plausible if it is used to explain more than one thing—is essentially the same as the one behind the factoring rule of [Stickel 90].

Now we are in a position to consider some examples involving the interaction of discourse pragmatics and syntax in German root clauses.

3.2 Generating Phrases in Initial Position

In a language like German with relatively free word order, any argument of the verbal head of a sentence may appear first, depending on the relevance for the discourse. For the spatial scenarios we consider, there is almost always a choice between several noun phrases or prepositional phrases that can be arranged in almost any order. As seen before, elements referring to familiar entities usually precede phrases denoting things not mentioned before.

Consider the case of locative verbs such as *stehen* ('to stand'), *sich befinden* ('to be located'), etc. We use the conventional Prolog translations of extended phrase structure rules to generate sentences headed by these verbs:

```
(12) v2_sentence(P0,P) :-
          accessible(X),
          accessible(R),
          np(syn(nom,_), X, P0,P1),     % a nominative NP with discourse referent X
          loc_vp(X, R, P1,P2),          % verb locating X inside the region R
          pp(R, P2,P).                  % a PP denoting a region R
```

```
(13) v2_sentence(P0,P) :-
          accessible(R),
          pp(R, P0,P1),
          loc_vp(X, R, P1,P2),
          np(syn(nom,_), X, P2,P).
```

```
(14) v2_sentence(P0,P) :-
          np(syn(nom,_), X, P0,P1),
          loc_vp(X, R, P1,P2),
          pp(R, P2,P).
```

Suppose we want to generate the sentence *Rechts von der Lampe steht eine Vitrine*, a variation of (6). The propositional content of this sentence—`lamp(l)`, `rightof(r,l)`, `loc(s,r)`, `showcase(s)`—must be assumable, and `accessible(l)` must be derivable from the listener model. Backward-chaining on `v2_sentence(S,[])` is not possible using rule (12) since `accessible(s)` is not provable and cannot be assumed. Rule (13) is applicable, but so is the weaker (14). Proofs involving these two rules will be equivalent except for the presence or absence of sub-proofs of `accessible(r)`, which is derivable from `accessible(l)`. But since proofs using more proved literals are preferred, the best abductive proof will result in a sentence with an initial definite PP preceding both the verb and the indefinite subject NP.

3.3 Generating Hanging Topics

To model the "topic shift" signaled by hanging topics, we need some way to represent the currently active discourse referent. This is achieved by introducing two predicates, `active(X)` and `activate(X)`, which test whether a given discourse referent is active or declare a discourse referent as active, respectively. For reasons of simplicity, we present these predicates as though they depend on a state external to the rules; in the implemented prototypes, each predicate is actually equipped with two additional variables that are used to drag along, test, and update the discourse state, in order to ensure a simple declarative semantics. In any proof, the literal `activate(X)` cannot be proved and has to be assumed, whereas `active(X)` must be

resolved exactly once with the closest matching literal of the form `activate(X)`. Thus the active referent is identified with the last activated one.

The rule for a sentence with a hanging topic can be seen in (15). Here it is not sufficient that the discourse referent associated with the noun phrase be inferentially accessible, a stronger condition is imposed, the requirement that the discourse referent must be taken from a set of thematic referents presumably established in a superordinate planning stage.

```
(15) sentence(P0,P) :-
          active(A),
          thematic_referent(X),
          distinct_objects(A, X),
          'C'(P0, was, P1),
          np(syn(acc,Gender), X, P1,P2),
          'C'(P2, betrifft, P3),
          activate(X),
          v2_sentence(P3,P).
```

Since a sentence with a hanging topic is used to re-activate an inactive discourse referent, and since an NP may be realized as a pronoun if its discourse referent is active, sentences of this type usually contain a pronoun, rather than a full NP, that refers back to the hanging topic, as in (8) above.

3.4 Generating Left Dislocations

Left dislocation constructions involve a semantics beyond the first-order theories used so far. This construction type presupposes that some salient object other than the discourse referent of the dislocated constituent lacks the property predicated by the sentence. Since we are dealing with highly specific rules for sentences with locative verbs, it is possible to express these conditions without reference to negative properties. All we have to do is to find some salient object distinct from the discourse referent of the dislocated NP, and a region where it is located distinct from the region in which the head verb locates the object denoted by the NP. In addition to this, the familiar discourse referent of the dislocated NP is made active. The resulting rule is displayed in (16):

```
(16) sentence(P0,P) :-
          active(A),
          familiar(X),
          distinct_objects(A, X),
          accessible(Y),
          distinct_objects(Y, X),
          loc(Y, R2),
          distinct_objects(R1, R2),
          np(syn(nom,Gender), X, P0,P1),
          activate(X),
          pron(syn(nom,Gender), P1,P2),
          loc_vp(X, R1, P2,P3),
          pp(R1, P3,P).
```

3.5 An Example Proof

Finally, we consider in some detail an example proof that illustrates several of the techniques used in generating the syntactic forms discussed previously. The sentence we want to derive should express that a showcase is located to the right of the lamp; additionally, we know that the immediately preceding discourse was about a different object, and some time ago we had explicitly mentioned the lamp.

(17) facts: `active(c), couch(c), thematic_referent(l), lamp(l), etc.`
assumable: `showcase(s), lamp(l), loc(s,r), right_of(r,l), activate(_)`
to prove: `sentence(S,[])`

```
sentence(S,[])
   active(A)   % resolution with fact, A bound to c
   thematic_referent(Y)   % resolution with fact, Y bound to l
   distinct_objects(c, l)   % provable from knowledge base
   'C'(S, was, P1)   % provable from knowledge base
     % S bound to [was|P1]
   np(syn(acc,Gender), l, P1,P3)
      det(syn(acc,Gender), P1,P2)   % Gender bound to fem
         familiar(l)
            thematic_referent(l)   % resolution with fact
         'C'(P1, die, P2)   % provable from knowledge base
           % P1 bound to [die|P2]
      n(syn(acc,fem), l, P2,P3)
         lamp(l)   % resolution with fact
         'C'(P2, lampe, P3)   % provable from knowledge base
           % P2 bound to [lampe|P3]
   'C'(P3, betrifft, P4)   % provable from knowledge base
     % P3 bound to [betrifft|P4]
   activate(l)   % can only be assumed
   v2_sentence(P4,[])
      accessible(R)
         rel(R, Z)
            spat_rel(R, Z)
               right_of(R, Z)   % can only be assumed, R bound to r, Z bound to l
         familiar(l)
            thematic_referent(l)   % resolution with a fact
      pp(r, P4,P6)
         active(W)
            activate(W)   % factoring with a previous assumption, W bound to l
         right_of(r, l)   % factoring with a previous assumption
         'C'(P4, rechts, P5)   % provable from knowledge base
           % P4 bound to [rechts|P5]
         'C'(P5, davon, P6)   % provable from knowledge base
           % P5 bound to [davon|P6]
      loc_vp(X, r, P6,P7)
```

```
        'C'(P6, steht, P7)   % provable from knowledge base
          % P6 bound to [steht|P7]
        loc(X, r)   % can only be assumed, X bound to s
    np(syn(nom,_), s, P7,[])
        det(syn(nom,Gender), P7,P8)   % Gender bound to fem
            'C'(P7, eine, P8)   % provable from knowledge base
              % P7 bound to [eine|P8]
        n(syn(nom,fem), s, P8,[])
            showcase(s)   % can only be assumed
            'C'(P8, vitrine, [])   % provable from knowledge base
              % P8 bound to [vitrine]
```

result: S bound to `[was,die,lampe,betrifft,rechts,davon,steht,eine,vitrine]`
assumed: `activate(l), right_of(r,l), activate(l), right_of(r,l), loc(s,r), showcase(s)` including two factored assumptions

Our preference criteria favor this proof over competing ones, since many goals could be proved, few had to be assumed, and assumptions could be re-used.

4 Conclusion and Outlook

We have presented a specification of how information packaging and syntactic structure interact in German root clauses. This specification is purely declarative, and neutral with respect to the task at hand. The axioms we use provide that kind of "thorough integration of syntax, semantics, and pragmatics" in the sense of [Hobbs et al. 93, 104ff.] that makes it possible to use the same knowledge base for both interpretation and generation. The mechanism used for either task is essentially the same too, viz. a reversible generalization of predicate specific abduction where only the set of facts that may be assumed differs depending on the task. The rules covering different syntactic forms and their information structure have antecedents of different strengths, sometimes subsuming each other, and the abductive mechanism ensures that the best explanation is the one that uses a rule with the strongest abductively provable antecedent.

As a direction for further research we suggest that a logical framework is needed that extends traditional knowledge representation and reasoning. Existing prototypes like the one described in [Hobbs et al. 93] are dealing to a large extent with *static* data: when trying to interpret a discourse, information is simply accumulated. Once one tries to incorporate reasoning about information packaging, one is faced directly with the problem of having to update the conversational record several times while a single sentence is interpreted or generated. The relevant reasoning is not so much about facts, but about actions or resources. We suggest that it is this dynamic aspect of information flow that is constant across the different tasks, and that the difference in generation and interpretation lies in the different status of information sources and sinks. A uniform framework that permits explicit reasoning about these dynamic aspects is highly desirable.

A second line of investigation should be concerned with trying to integrate content planning and syntactic realization. This could be exploited to allow for incremental generation, or be used for the generation of idiomatic expressions and syntactic patterns that are not linked to any semantic content, but rather to certain discourse goals. If this task is carried out successfully, we might see a thorough integration of all aspects of natural language generation.

References

[Altmann 81] Altmann, H. 1981. *Formen der "Herausstellung" im Deutschen*. Tübingen: Niemeyer.

[Dale et al. 95] Dale, R., and E. Reiter. 1995. Computational interpretations of the Gricean maxims in the generation of referring expressions. *Cognitive Science* 18, 233–263.

[Grice 75] Grice, H.P. 1975. Logic and conversation. In P. Cole, and J.L. Morgan (eds), *Syntax and Semantics 3: Speech Acts*, New York: Academic Press, 41–58.

[Hobbs et al. 93] Hobbs, J.R., M.E. Stickel, D.E. Appelt, and P. Martin. 1993. Interpretation as abduction. *Artificial Intelligence* 63, 69–142.

[Jameson et al. 94] Jameson, A., B. Kipper, A. Ndiaye, R. Schäfer, J. Simons, T. Weis, and D. Zimmermann. 1994. Cooperating to be noncooperative: The dialog system PRACMA. In B. Nebel, and L. Dreschler-Fischer (eds), *Proceedings of the Eighteenth Annual German Conference on Artificial Intelligence*, Berlin: Springer, 106–117.

[Jansche et al. 96] Jansche, M., and R. Porzel. 1996. PAROLE*: A Cognitively Motivated NLG System for Spatial Descriptions*. Universitäten Heidelberg/Mannheim: Arbeiten aus dem Sonderforschungsbereich 245 "Sprache und Situation"; Report Nr. 93.

[Lambrecht 94] Lambrecht, K. 1994. *Information Structure and Sentence Form*. Cambridge: Cambridge University Press.

[Lascarides et al. 92] Lascarides, A., and J. Oberlander. 1992. Abducing temporal discourse. In R. Dale, E. Hovy, D. Rösner, and O. Stock (eds), *Aspects of Automated Natural Language Generation*, Berlin: Springer, 167–182.

[Mayer et al. 96] Mayer, M.C., and F. Pirri. 1996. A study on the logic of abduction. In W. Wahlster (ed), *Proceedings of the 12th European Conference on Artificial Intelligence (ECAI 96)*, Chichester: Wiley, 18–22.

[Meyer-Klabunde 96] Meyer-Klabunde, R. 1996. A case study on feedback in cognitive models of language production. In W. Hoeppner (ed), *Proceedings of the 6th European Workshop on Natural Language Generation*, Universität-GH Duisburg, Institut für Informatik; Report Nr. SI-17, 61–71.

[Paris 93] Paris, C. 1993. *User Modelling in Text Generation*. London: Pinter.

[Porzel et al. in press] Porzel, R., M. Jansche, and R. Meyer-Klabunde. In press. The generation of spatial descriptions from a cognitive point of view. In P. Olivier (ed), *Vision and Language*, Berlin: Springer.

[Prince 81] Prince, E.F. 1981. Toward a taxonomy of given–new information. In P. Cole (ed), *Radical Pragmatics*, New York: Academic Press, 223–255.

[Stickel 90] Stickel, M.E. 1990. Rationale and methods for abductive reasoning in natural language interpretation. In R. Studer (ed), *Natural Language and Logic*, Berlin: Springer, 233–252.

[Thomason et al. 96] Thomason, R.H., J.R. Hobbs, and J.D. Moore. 1996. Communicative goals. In K. Jokinen, M. Maybury, M. Zock, and I. Zukerman (eds), *Proceedings of the ECAI 96 Workshop Gaps and Bridges: New Directions in Planning and Natural Language Generation*, Budapest, 7–12.

[Thomason et al. 97] Thomason, R.H., and J.R. Hobbs. 1997. Interrelating interpretation and generation in an abductive framework. Paper presented at the AAAI Fall 1997 Symposium on Communicative Action in Humans and Machines. Cambridge, MA.

[Vallduví 92] Vallduví, E. 1992. *The Informational Component*. New York: Garland.

[Wiltschko 95] Wiltschko, M. 1995. Presuppositions in German dislocation constructions. *Folia Linguistica* XXIX, 265–295.