

Saliency and the Attentional State in Natural Language Generation

Chris Reed¹

Abstract. The importance of saliency and attention in natural language generation is often underestimated; this paper aims to demonstrate why it is necessary to view and operationalise saliency in a similar way as has happened with intention. The discussion draws on both theoretical and practical considerations, and describes a generation system which implements the model.

1 INTRODUCTION

This paper argues for the importance of recognising the operational role of saliency in natural language generation, and discusses a system which implements such a role. The premises for the argument fall into two categories: those based on aspects of the theory of discourse structure – particularly the work of Grosz and Sidner (1986) – and those based on empirical data. A small worked example then shows how the main elements of the theory play out in practice. Before considering these aspects, a brief summary of what is meant by saliency and its implementation is necessary.

2 PROPOSITIONAL SALIENCE

In early work on salience in natural language generation, the term 'salience' was employed to refer to objects which were prominent in a visual scene, and which were then selected in describing the scene (Conklin & McDonald, 1982). The notion of salience in that work was thus defined outside the domain of discourse, and was subsequently employed in determining content selection.

A decade later, such **object salience** was generalised, on the basis of psycholinguistic evidence, by considering 'canonical' and 'instantial' components (Pattabhiraman, 1993). The former reflected linguistic preferences, with some lexemes and structures simply being more salient than others in a given cultural milieu, whilst the latter incorporated features such as vividness and recency of mention to quantify salience in a given linguistic utterance. This domain independent approach to assessing salience was then employed as the foundation for a decision theoretic algorithm which viewed the generation process as salience maximisation. In both these works, salience is something intrinsic to a linguistic construct or the object to which such a construct refers.

In contrast, **propositional salience** refers to the (actual or anticipated) mental state of the intended audience, and considers the effect that an utterance fragment has upon the features to which the audience is attending. A linguistic construction can make some proposition salient (to some audience). Propositional salience, unlike object salience, can thus be operationalised in much the same way as intentional goals, endowing the generation system with the responsibility for determining what needs to be made salient, and what does not. Thus, generating text by

planning for propositional salience could adduce the results of algorithms that determine object salience. For example to make (propositionally) salient the fact that "John painted the house", a generator may employ canonical salience to determine that the active voice is the best alternative. Propositional salience thus refers to affecting the awareness of a hearer with respect to a belief.

2.1 Pairing Goals of Belief with Goals of Propositional Salience

The *Rhetorica* system presented in (Reed and Long, 1997, Reed, 1999) makes explicit use of goals of propositional salience, as distinct from intentional goals of belief. Typically, these goals are paired, so that the goals

```
BEL(H, P)
IS_SALIENT(H, P, C)
```

would together express the intention of bringing the hearer to believe and be aware of some proposition P (in a context C, as discussed in section 3.2). The role of the belief goal is to generate the structure of a text: *Rhetorica* is concerned with the construction of persuasive argument, in which goals of belief are typically supported by chains of reasoning. The planning process associated with these goals can be seen as constructing the skeleton of a discourse, to which the linguistic flesh can be attached. The goals of saliency are then responsible for this fleshing out.

Eventually, chains of reasoning must bottom out in assumptions made by a speaker about the hearer's knowledge – these assumptions represent immediate termination points for planning for goals of belief. Goals of saliency can also be satisfied in an immediate fashion: in some cases, linguistic realisation may not be required to make a fact salient. This typically occurs when the context can be guaranteed to make information salient to an interlocutor, and examples of this phenomenon are given in section 4.

This bipartite approach to high level abstract goals (and tripartite approach to language goals in general, if goals of linguistic realisation are included) is strongly supported on both theoretical and practical grounds.

3 THEORETICAL CONSIDERATIONS

3.1 Saliency and the Focus of Attention

Grosz and Sidner (1986) clearly state their views on the attentional component of communication: "Attention is an essential factor in explicating the processing of utterances in discourse." (p175). Most current researchers in NLP would probably agree. Grosz and Sidner go on to comment on the lack

¹ Applied Computing, University of Dundee, Dundee DD1 4HN Scotland, email chris@computing.dundee.ac.uk

of recognition of the difference between linguistic, intentional and attentional structure: "Most related work on discourse structure fails to distinguish among some (or, in some cases, all) of these components. As a result, significant generalizations are lost, and the computational mechanisms proposed are more complex than necessary." (p176). Remarkably, within natural language generation, at least, this comment still substantially holds true. Although most systems make a clear distinction between intentional ends and linguistic means, and gain significant flexibility and expressiveness as a result, few operationalise manipulation of salient facts. Without operationalising the attentional state in as comprehensive a way as the intentional and linguistic facets, NLG systems suffer the missed generalizations and unnecessary complexity described by Grosz and Sidner.

In the current work, the identification of goals of saliency as distinct from, but complementary to, goals of belief (and other mentalistic notions), retains Grosz and Sidner's distinction. Intentional structure is created by the fulfilment of goals of belief, and attentional structure by goals of saliency, and the rich interplay between the two aspects can be exploited.

One example of such exploitation is in the ordering of text spans. Analysis of natural language typically proceeds by building up tree structures – both RST and GST offer good examples of this approach. Mirroring the technique, modern natural language generation systems often start from some initial intention, and gradually introduce increasingly refined specifications resulting ultimately in linguistic structure. This is perhaps best demonstrated by operationalisations of RST such as (Hovy, 1990). In generating tree structures, however, there is an extra problem not encountered in text analysis: determining an ordering between children of a particular parent node. Solutions to the problem have involved either arbitrary choice (missing important generalizations) or complex heuristics (introducing unnecessary complexity). Using goals of saliency, utterances can be tied to particular focus spaces, by specifying, as part of the goal, which focus space is to be modified by the addition of a new salient proposition. Thus saliency is specified not only with respect to an interlocutor and a proposition, but also to a particular attentional state. As a discourse proceeds, and the focus of attention shifts, particular utterances become licensed. In this way, an ordering between spans can be resolved, with the added benefit that other propositions can be guaranteed to be salient – present in the focus space – with the utterance or presentation of any given span (Reed and Long, 1997).

3.2 The Focus of Attention and Context

For Grosz and Sidner, the role of the attentional state is in "recording the objects, properties, and relations that are salient at each point in the discourse" (p179). In the current work, the notion is extended to include not only salient features, but also (reference to) the *context* of an utterance. A single focus space provides access to a single context, which can be thought of as a specialised knowledge base of propositions true at this point in the discourse. This understanding of a context has been proposed by McCarthy and Buvac (Buvac *et al.*, 1995; McCarthy & Buvac, 1995), and relies upon the relation *ist* (which in Buvac *et al.* is formalised as a modal operator). A statement of the form

ist(c, p)

states that the proposition *p* is true in the context *c*, so that

ist(c1, at(jmc, Stanford))

is an assertion about McCarthy's affiliation in a context *c1* in which the symbol *jmc* refers to (the appropriate) John McCarthy, *Stanford* to Stanford University, and *at* to a relation of affiliation (McCarthy and Buvac, 1995, p2). The properties of this context logic, including quantification over contexts and the transfer of information between contexts, are described formally in (Buvac *et al.*, 1995), and although more recent work has developed logics of context further, the logic of *ist* suffices to demonstrate the interaction between context and focus.

Why include such context specification in the focus space? In the first place, it facilitates the characterisation of illocution-dependent responses (such as "yes" indicating the truth of a proposition included in a yes/no query) – this is one of the issues addressed by McCarthy and Buvac in their preliminary analysis of context in simplified, logical dialogue. More broadly, though, it provides a means of capturing a number of important features of discourse, such as common knowledge and, relatedly, ambiguity.

Grosz and Sidner mention the role of common knowledge in discourse, citing Grosz's earlier work, and caution against identifying this (in addition to intentionality) with attention (p180). It is left unclear, however, whether they would want to see common knowledge playing a role in the attentional state, though the claim that "the focusing structure is the central repository for the contextual information needed to process utterances at each point in the discourse" suggests that it would be consistent to adopt that assumption. McCarthy and Buvac's context logic provides a perfect means of characterising this knowledge – in a particular context, two interlocutors may share knowledge about a domain, but that common knowledge may change from one context to another, as the dialogue proceeds. Thus the context associated with a given focus space can include statements such as

ist(c99, MB("keys can undo screws"))

which says that in some context *c99*, connected to a focus space part way through a dialogue such as that given in Grosz and Sidner, p186, it is mutually believed that (allen) keys can undo (allen head) screws. (Following McCarthy and Buvac, the examples here use an English gloss to avoid commitment to any specific formalisation of knowledge). It is clear (i) that such information is of use in generating utterances such as "Are you sure you are using the right size key?", and (ii) that the fact that "keys can undo screws" is only true in some limited set of contexts – and may well be true in only a subset of those contexts referenced by focus spaces of a single dialogue.

The example above suggests an important generalisation: that exploiting knowledge available in a context can lead to significant abridging in the generation process. From an analysis point of view, a major stumbling block in understanding a sentence is to resolve ambiguities over, amongst other things, the meaning of predicates and terms. This stumbling block is at least in part crossed by McCarthy and Buvac. From a generation point of view, it is (in most cases) clearly undesirable to intentionally introduce ambiguity. It is, however, highly desirable to avoid redundancy, long-windedness and repetition. The contents of a given context, such as common knowledge, are precisely what are required to determine what information can be left implicit – so that the contextual mutual belief given above contributes to the utterance, "Are you sure you are using the right sized key?", rather than, "...the right sized allen key?". A similar contextual fact means that it would similarly be unnecessary to utter, "...the right sized bent piece of metal with a hexagonal cross-section?", or any other description which, though potentially appropriate in

some contexts, is inappropriate in this one. This approach to avoiding redundancy is a general technique which has very specific applications to the commonsense knowledge in a given context. This is in contrast to general techniques applicable to a wide range of situations such as Horacek's (1994) utilisation of conversational implicature in avoiding redundancy. The two approaches complement one another, and could together produce highly efficient text.

In this way, tying a goal of saliency to a particular focus space, and thereby, to a particular context, makes a range of additional information available to the realisation process, leading to more concise and context-specific utterances. In the limit, there may be information available in the given context which implies that the content of a saliency goal is already salient to a hearer, in which case no linguistic realisation is required at all. An example of such a situation is an argument form known as the enthymeme, and forms the first of a number of concrete examples which further support the need for explicit handling of saliency.

4 PRACTICAL CONSIDERATIONS

In (Reed, 1999) the importance of saliency is motivated by noting three features common in the domain of persuasive argumentation: enthymemes, repetition, and refutation. An enthymeme is a syllogistic argument in which one component is left implicit (e.g. *All men are mortal, so Socrates is mortal*, which leaves implicit the fact that *Socrates is a man*). In other situations, arguments with identical form may appropriately be rendered with every component explicit. To handle the need for this flexibility, goals of saliency are used and then, in some cases (such as the Socrates example) discharged as a result of the context, without need for any realisation. Repetition is used to great rhetorical effect, but poses a problem for traditional NLG since it fulfils an intention to bring an audience to believe something that the speaker believes they already know. The solution lies in posting a new goal of saliency, without any associated epistemic intention. Lastly, refutation often employs a technique of stating up front the conclusion which is to be disproved. This too poses problems for traditional NLG approaches, and can be handled by introducing goals of saliency with respect to a proposition p which are ultimately associated with intentions to bring an audience to believe *not(p)*. For these three problems, the saliency-based approach proves much more suitable than the main alternative, based on informationally redundant units (Walker, 1996).

Although it is possible to generalise from these examples in the domain of argumentation to broader natural language concerns (particularly in the case of the enthymeme, since contracting chains of reasoning is common in other forms of discussion such as task oriented dialogues), there are additional examples which demand consideration.

The first is the need to characterise a situation in which a hearer may be known to believe something, but is not, in the current (attentional) state, aware of it. In other words, the problem is one of how to remind someone. Though examples conceivably occur in all types of dialogue, pedagogical exchanges are a particularly rich source. Consider the following exchange:

1. T: OK, so to write a recursive function, you start with the terminating case, and then work out a way of nibbling off a part of the problem that you can solve

easily. Then you need to recurse on all but that nibbled off part, right?

2. S: Uhuh. Terminating case and then nibble.

3. T: OK, so how would you go about writing a recursive function to calculate the n -th fibonacci number?

4. S: Um, not sure.

5. T: Well, what do you start with when writing any recursive function?

6. S: The terminating case.

7. T: Right! ...

At (2), T is given some indication that S knows how to go about writing a recursive function. At (4), however, S, still with that knowledge, seems to need reminding. At (5), a reminder is supplied, and (on the basis of (6)) is successful. Unless (4) is taken to mean that S simply no longer has the requisite knowledge (which requires quite a stretch of the imagination, given its proximity to (2)), most standard intention-based NLG systems would have difficulty in generating an appropriate response at (5), since S's knowledge would be recorded as being present, and therefore any goal of the form `KNOW-ABOUT(S, terminating-case)` would be immediately and trivially satisfied. Separating out a notion of saliency, however, means that the initial `KNOW-ABOUT` goal is satisfied by (1), and (4) simply elicits the posting of a saliency goal such as `IS_SALIENT(S, terminating-case, c)`. If this were a non-didactic dialogue, it might be expected that this saliency goal would be realised as a straightforward imperative, "Start with the terminating case". In this case, however, the saliency is instead achieved through asking a question – the interesting problem posed by this example of determining the illocutionary force by which information is made salient is beyond the scope of the current work. It is clear, however, that (i) a goal of saliency alone could be fulfilled by assertion, imperative, question, or whatever, and that (ii) the introduction of an epistemic intention at (5) would, by itself, fail to generate appropriate text (quite apart from running counter to intuition about the role of (5) in the dialogue).

The second problem is related to the first, and relies on a construction of belief dependent upon the formalisation of context cited above. An interlocutor may believe something in one context, but not in the current one: how can a generation system produce text which allows the transfer of that belief to the current context? (In McCarthy and Buvac's terminology, what text will *lift* the relevant belief?) Consider a slightly different conclusion to the example in Grosz and Sidner:

1. E. Use the wheelpuller. Do you know how to use it?
2. A. No.
3. E. It's basically just a smaller version of the gizmo you used to take apart your gear assembly on your bicycle

Here, the locution at (3) might be based upon E's knowledge that A had recently renovated his bicycle, using a wheelpuller in the process. Without the notion of context, it would be necessary to assume that A actually already knows how to use a wheelpuller, and merely needs reminding at (3). This seems unsatisfactory, however. Using contexts the generation of (3) can be accounted for much more cleanly. First, E believes that A knows about wheelpullers in the context of bicycles (b)

`BEL(E, ist(b, KNOW-HOW(Use A wheelpullers)))`

and that A has demonstrated he does not know about them in the current context (*c*)

```
BEL(E, ist(c, ~KNOW-HOW(Use A wheelpullers)))
```

The goals that lead to (21) thus correspond to a pair of goals (one intentional, one attentional) that relate to the lifting of A's knowledge from context *b* to context *c*, plus a goal of saliency relating to A's knowledge in context *b*:

```
BEL(E, BEL(A, SUPPORTS(ist(b,
    KNOW-HOW(Use A wheelpullers),
    ist(c, KNOWS-HOW-TO(Use A wheelpullers))))
IS_SALIENT(A, ist(c, SUPPORTS(ist(b,
    KNOW-HOW(Use A wheelpullers),
    ist(c, KNOWS-HOW-TO(Use A wheelpullers),c)
IS_SALIENT(A, ist(b,
    KNOW-HOW(Use A wheelpullers)), c)
```

The last of these three goals, which explicitly involves the saliency of a belief from one context (*b*) in another (*c*) could reasonably be expected to usually be sufficient to guarantee (by the principle of relevance) the satisfaction of the second goal. As justification of inclusion of the first goal, it is conceivable that further support could be adduced for the lifting – e.g. "Wheelpullers all work on the same principle, you see". The fact that further discourse structure can be added at this point suggests that it involves an intention which can be planned for if necessary.

The third and final problem to be considered here concerns the problem of replanning segments of discourse in the case of communication failure. In a situation in which the perlocutionary effect of an utterance fails (a hearer is not persuaded, or does not know-about, or whatever), it is the speaker's intention, INTEND(*S*, BEL(*H*, *P*)) or INTEND(*S*, KNOW-ABOUT(*P*)), etc. which is the starting point for replanning. In the case where a hearer demonstrates a lack of understanding, it is goals of linguistic realisation which need replanning. For systems which do not handle saliency explicitly, however, it is difficult to handle a third, analogous case competently. If a goal of saliency fails, the result is a mismatch between what the speaker thinks the hearer is aware of, and what he actually is aware of. For example:

- (1) S: Radioactive elements decay and eventually turn into lead, right?
- (2) H: Yeah.
- (3) S: So, if matter had always existed, there'd be no radioactive elements left – but there are radioactive elements around.
- (4) H: Yes, I agree.
- (5) S: That means matter must have been created!
- (6) H: Whoah! How come?
- (7) S: Because matter can't have always existed.

At (3), S is uttering an enthymematic Modus Tollens argument, claiming first the major premise ($m \supset \sim r$), and then the minor premise (*r*), but leaving the conclusion ($\sim m$) implicit. The decision to leave the conclusion implicit is a result of the associated goal of saliency being fulfilled through contextual, rather than linguistic means (Reed, 1999). Thus the result of (3) is assumed by S to be that each of $m \supset \sim r$, *r* and $\sim m$ is salient to the hearer. After H concurring, S continues her line of reasoning, at which H balks. One possible problem with H's understanding – and the problem which S sets about rectifying – is that the implicit conclusion at (3) failed to be salient to the hearer. To rectify this, the saliency goal is reposted, and a new planning

process leads, this time, to the generation of explicit linguistic realisation. The process of detecting and attributing such failure is certainly very tough: the argument here is simply that saliency must be one of the building blocks of that process.

Finally, it is worth pointing out that Grosz and Sidner mention a number of specific phenomena based on the attentional state and its interplay with intentions, including pronominalisation, intention recognition, handling of interruptions, and generation of cue phrases. The architecture described in the current work carefully retains the distinctions proposed by Grosz and Sidner, and therefore has a solid base from which to attack the aforementioned problems, though, with the exception of cue phrases which have been discussed in (Reed and Long, 1997), these avenues are left for future work.

5 WORKED EXAMPLE

The core features of the approach proposed have been implemented in the *Rhetorica* system, described more fully elsewhere (Reed and Long, 1997). A small example will serve to demonstrate how the operationalisation of propositional salience is integrated in a wider plan-based generation system. *Rhetorica* is responsible for producing detailed, annotated discourse plans; here, as in earlier work, the plan is reduced to fully specified text manually (or, in fact, in a semi-automated fashion if canned text is associated with propositions). The scenario is a management decision support system, and involves trying to construct an argument for a particular user in favour of keeping the size of a project team to a minimum. To introduce issues discussed in the paper, let us assume that the system has a rich knowledge base that includes a sophisticated characterisation of the user in which cultural common knowledge such as that expressed in proverbs, is available.

The starting situation is one in which the system, S, maintains the following beliefs:

```
BEL(S, ist(projects, small_team))
BEL(S, ist(projects,
    SUPPORTS(too_many, small_team))
BEL(S, ist(projects, too_many))
BEL(S, BEL(H, ist(cooking, too_many)))
```

These carry the following gloss (or canned representation):

(In the context of software development projects) a small project team is best

(In the context of software development projects) (the fact that) having too many contributors can jeopardise product quality (supports the fact that) a small project team is best

(In the context of projects) too many cooks spoil the broth

(The hearer believes that) too many cooks spoil the broth

The first two are specific to the task in hand, of producing an argument in favour of a small project team, and the third a representation of the common proverb, that (at least in the UK) might reasonably be assumed to be believed by any interlocutor. The planning process starts with the system goals to bring the hearer to believe the conclusion and be aware of the fact:

```
BEL(H, ist(projects, small_team))
IS_SALIENT(H, ist(projects, small_team), null)
```

In the first round of planning, a straightforward Modus Ponens argument can be adduced from the *too_many* premise. This produces the following partial plan (in which some system output has been simplified in the interests of brevity and clarity):

```

MAKE_SALIENT(H, ist(projects, small_team), null)
PUSH_TOPIC(projects)
BEL (H, too_many)
IS_SALIENT (H, too_many, projects)
BEL (H, supports(ist(projects, too_many),
                 ist(projects, small_team)))
IS_SALIENT (H, supports(ist(projects, too_many),
                        ist(projects, small_team)), projects)
POP_TOPIC(projects)

```

Thus, at this stage, the hearer is to be convinced of two things: that the proposition *too_many* is true, and that it supports the claim *small_team*. For each of these there is both a goal of belief, used to generate further argumentative structure in the planning process, and a goal of saliency, to be used to introduce linguistic structure. Notice that goals of saliency can always be fulfilled, if nothing else is available, by the operation *MAKE_SALIENT*. This represents a demand for linguistic material to be introduced.

The next round of planning fulfils the second *BEL* goal heuristically (rhetorical evidence suggests that, other things being equal, an audience will question claims and premises rather than inference steps). The first *BEL* goal is satisfied through the application of a **lifting argument**, similar to that used in section 4, above. The function of this planning operator is to introduce appropriate goals (of belief and saliency) corresponding to the action of lifting a proposition from one context to another. The resulting plan is as follows:

```

MAKE_SALIENT(H, ist(projects, small_team), null)
PUSH_TOPIC(projects)
  PUSH_TOPIC(cooking)
  BEL(H, ist(cooking, too_many))
  IS_SALIENT (H, ist(cooking, too_many))
  BEL (H, SUPPORTS(ist(cooking, too_many),
                  ist(projects, too_many)))
  IS_SALIENT (H, SUPPORTS(ist(cooking, too_many),
                          ist(projects, too_many)), projects)
  POP_TOPIC(cooking)
  MAKE_SALIENT (H, too_many, projects)
  IS_SALIENT (H, supports(ist(projects, too_many),
                          ist(projects, small_team)), projects)
POP_TOPIC(projects)

```

The last two goals of belief are both satisfied trivially, the first because it matches one of the system's explicit initial assumptions about what the hearer believes, and the second for the same optimistic assumption of limited scepticism mentioned above. This leaves a final plan with the following components:

```

MAKE_SALIENT(H, ist(projects, small_team), null)
PUSH_TOPIC(projects)
  PUSH_TOPIC(cooking)
  MAKE_SALIENT(H, ist(cooking, too_many), cooking)
  POP_TOPIC(cooking)
  MAKE_SALIENT (H, too_many, projects)
POP_TOPIC(projects)

```

Relying on just the canned text described above, the *MAKE_SALIENT* directives in this plan would lead to the following text: *A small project team is best because too many cooks spoil the broth, and having too many contributors can jeopardise product quality.* Though *Rhetorica* is responsible for the two clue words (*and* and *because*) in this example, the structure of the lifting argument might, analogously to the Modus Tollens *but* (Reed, 1999), canonically involve a clue word such as *similarly* to be introduced between the last two phrases.

6 RELATED WORK

Object salience has been used for content selection in several generation systems, including Conklin and McDonald's (1982) seminal work. Propositional salience, however, has not previously played a central role, rendering operational uniformity between intention, attention and information.

Focus of attention structures have played a role in many venerable generation systems, with, for example, a focus algorithm in McKeown's (1985) *TEXT* system, context space suspensions and resumptions in Reichman's model (1987), and focus trees adumbrating RST-based planning in (Hovy and McCoy, 1989). These and similar approaches show how focus spaces can be manipulated in effecting anaphoric reference, content selection choices, and other phenomena related to text coherence. As with many other modern generation systems, *Rhetorica* too manipulates focus structures. Although Hovy and McCoy (1989) might argue that trees are the most appropriate focus structures, and Walker (1996) that a cache model is more appropriate, a strong case can still be made for the simpler, stack model of Reichman *inter alia* (Reed, 1999), and it is this that is adopted in *Rhetorica*. The implementation relies upon push and pop operations which form part of the discourse plan.

Thus, propositions are made salient with respect to contexts that are introduced and discharged through planned manipulations of the topic stack. For the implementation, this introduces a computational saving, by allowing the partial order of the discourse plan to remain unspecified until 'run time', when focus changes constrain utterance order. The important difference between the proposed model and those which do not treat saliency on a par with intentionality is that it is goals of saliency alone which lead to the introduction of linguistic material.

7 CONCLUSIONS

In this paper, the adoption of a distinction between goals of belief (and other mentalistic attitudes) which manipulate intentional structure, and goals of saliency, which manipulate attentional structure, is motivated through considerations of theoretical consistency and practical utility. The role of context logic in specifying the attentional state is sketched, and the model is demonstrated to be amenable to implementation by presentation of elements of the *Rhetorica* generation system.

8 REFERENCES

- Buvac, S., Buvac, V. & Mason, I.A. (1995) "Metamathematics of contexts" *Fundamenta Informaticae* **23** (3)
- Conklin, E.J. & McDonald, D.M. (1982) "Salience, the Key to the Selection Problem in NLG", in *Proceedings of the 20th ACL*, pp129-135
- Grosz, B.J. & Sidner, C.L. (1986) "Attention, Intention and the Structure of Discourse", *Computational Linguistics* **12** (3) pp175-204
- Horacek, H. (1994) "How to avoid explaining obvious things", in Cohn, A. (ed) *Proceedings of the 11th European Conference on Artificial Intelligence*, pp520-524
- Hovy, E.H. (1990) "Pragmatics and Natural Language Generation" *Artificial Intelligence* **43** pp153-197
- Mann, W.C. & Thompson, S.A. (1988) "Rhetorical Structure Theory: Toward a functional theory of text organization", *Text* **8** (3) pp243-281
- Marcu, D. (1999) "A Formal and Computational Synthesis of Grosz and Sidner's and Mann and Thompson's theories", *Proceedings of the Workshop on Levels of Representation in Discourse*, Edinburgh
- McCarthy, J. & Buvac, S. (1995) "Formalizing Context (Expanded Notes)" in *Working Notes of the AAAI Fall Symposium on Formalizing Context*
- Pattabhiraman, T. (1993) "Aspects of Salience in Natural Language Generation"
- Reed, C.A. & Long, D.P. (1997) "Content Ordering in the Generation of Persuasive Discourse" in *Proceedings of IJCAI'97*, Nagoya, Morgan Kaufmann, pp1022-1027
- Reed, C.A. (1999) "The Role of Salience in Generating Natural Language Arguments", in *Proceedings of IJCAI'99*, Stockholm, Morgan Kaufmann
- Walker (1996) "The effect of resource limits and task complexity on collaborative planning in dialogue", *Artificial Intelligence* **85** pp181-243