

Never Look Back: An Alternative to Centering

Michael Strube

IRCS – Institute for Research in Cognitive Science
University of Pennsylvania
3401 Walnut Street, Suite 400A
Philadelphia PA 19104
strube@linc.cis.upenn.edu

Abstract

I propose a model for determining the hearer’s attentional state which depends solely on a list of salient discourse entities (S-list). The ordering among the elements of the S-list covers also the function of the *backward-looking center* in the centering model. The ranking criteria for the S-list are based on the distinction between *hearer-old* and *hearer-new* discourse entities and incorporate preferences for inter- and intra-sentential anaphora. The model is the basis for an algorithm which operates incrementally, word by word.

1 Introduction

I propose a model for determining the hearer’s attentional state in understanding discourse. My proposal is inspired by the centering model (Grosz et al., 1983; 1995) and draws on the conclusions of Strube & Hahn’s (1996) approach for the ranking of the *forward-looking center* list for German. Their approach has been proven as the point of departure for a new model which is valid for English as well.

The use of the centering transitions in Brennan et al.’s (1987) algorithm prevents it from being applied incrementally (cf. Kehler (1997)). In my approach, I propose to replace the functions of the *backward-looking center* and the *centering transitions* by the order among the elements of the list of salient discourse entities (S-list). The S-list ranking criteria define a preference for *hearer-old* over *hearer-new* discourse entities (Prince, 1981) generalizing Strube & Hahn’s (1996) approach. Because of these ranking criteria, I can account for the difference in salience between definite NPs (mostly hearer-old) and indefinite NPs (mostly hearer-new).

The S-list is not a local data structure associated with individual utterances. The S-list rather describes the attentional state of the hearer at any given point in processing a discourse. The S-list is generated incrementally, word by word, and used

immediately. Therefore, the S-list integrates in the simplest manner preferences for inter- and intra-sentential anaphora, making further specifications for processing complex sentences unnecessary.

Section 2 describes the centering model as the relevant background for my proposal. In Section 3, I introduce my model, its only data structure, the S-list, and the accompanying algorithm. In Section 4, I compare the results of my algorithm with the results of the centering algorithm (Brennan et al., 1987) with and without specifications for complex sentences (Kameyama, 1998).

2 A Look Back: Centering

The centering model describes the relation between the focus of attention, the choices of referring expressions, and the perceived coherence of discourse. The model has been motivated with evidence from preferences for the antecedents of pronouns (Grosz et al., 1983; 1995) and has been applied to pronoun resolution (Brennan et al. (1987), inter alia, whose interpretation differs from the original model).

The centering model itself consists of two constructs, the *backward-looking center* and the list of *forward-looking centers*, and a few rules and constraints. Each utterance U_i is assigned a list of *forward-looking centers*, $Cf(U_i)$, and a unique *backward-looking center*, $Cb(U_i)$. A ranking imposed on the elements of the Cf reflects the assumption that the most highly ranked element of $Cf(U_i)$ (the *preferred center* $Cp(U_i)$) is most likely to be the $Cb(U_{i+1})$. The most highly ranked element of $Cf(U_i)$ that is *realized* in U_{i+1} (i.e., is associated with an expression that has a valid interpretation in the underlying semantic representation) is the $Cb(U_{i+1})$. Therefore, the ranking on the Cf plays a crucial role in the model. Grosz et al. (1995) and Brennan et al. (1987) use grammatical relations to rank the Cf (i.e., *subj* < *obj* < ...) but state that other factors might also play a role.

For their centering algorithm, Brennan et al. (1987, henceforth BFP-algorithm) extend the notion of centering transition relations, which hold across adjacent utterances, to differentiate types of shift (cf. Table 1 taken from Walker et al. (1994)).

	$Cb(U_i) = Cb(U_{i-1})$ OR no $Cb(U_{i-1})$	$Cb(U_i) \neq Cb(U_{i-1})$
$Cb(U_i) = Cp(U_i)$	CONTINUE	SMOOTH-SHIFT
$Cb(U_i) \neq Cp(U_i)$	RETAIN	ROUGH-SHIFT

Table 1: Transition Types

Brennan et al. (1987) modify the second of two rules on center movement and realization which were defined by Grosz et al. (1983; 1995):

Rule 1: If some element of $Cf(U_{i-1})$ is realized as a pronoun in U_i , then so is $Cb(U_i)$.

Rule 2: Transition states are ordered. CONTINUE is preferred to RETAIN is preferred to SMOOTH-SHIFT is preferred to ROUGH-SHIFT.

The BFP-algorithm (cf. Walker et al. (1994)) consists of three basic steps:

1. GENERATE possible $Cb-Cf$ combinations.
2. FILTER by constraints, e.g., contra-indexing, sortal predicates, centering rules and constraints.
3. RANK by transition orderings.

To illustrate this algorithm, we consider example (1) (Brennan et al., 1987) which has two different final utterances (1d) and (1d'). Utterance (1d) contains one pronoun, utterance (1d') two pronouns. We look at the interpretation of (1d) and (1d'). After step 2, the algorithm has produced two readings for each variant which are rated by the corresponding transitions in step 3. In (1d), the pronoun “she” is resolved to “her” (= Brennan) because the CONTINUE transition is ranked higher than SMOOTH-SHIFT in the second reading. In (1d'), the pronoun “she” is resolved to “Friedman” because SMOOTH-SHIFT is preferred over ROUGH-SHIFT.

- (1) a. Brennan drives an Alfa Romeo.
- b. **She** drives too fast.
- c. Friedman races **her** on weekends.
- d. **She** goes to Laguna Seca.
- d'. **She** often beats **her**.

3 An Alternative to Centering

3.1 The Model

The realization and the structure of my model departs significantly from the centering model:

- The model consists of one construct with one operation: the list of salient discourse entities (S-list) with an insertion operation.
- The S-list describes the attentional state of the hearer at any given point in processing a discourse.
- The S-list contains some (not necessarily all) discourse entities which are realized in the current and the previous utterance.
- The elements of the S-list are ranked according to their information status. The order among the elements provides directly the preference for the interpretation of anaphoric expressions.

In contrast to the centering model, my model does not need a construct which looks back; it does not need transitions and transition ranking criteria. Instead of using the Cb to account for local coherence, in my model this is achieved by comparing the first element of the S-list with the preceding state.

3.2 S-List Ranking

Strube & Hahn (1996) rank the Cf according to the information status of discourse entities. I here generalize these ranking criteria by redefining them in Prince’s (1981; 1992) terms. I distinguish between three different sets of expressions, *hearer-old discourse entities* (OLD), *mediated discourse entities* (MED), and *hearer-new discourse entities* (NEW). These sets consist of the elements of Prince’s *familiarity scale* (Prince, 1981, p.245). OLD consists of *evoked* (E) and *unused* (U) discourse entities while NEW consists of *brand-new* (BN) discourse entities. MED consists of *infernables* (I), *containing infernables* (I^C) and *anchored brand-new* (BN^A) discourse entities. These discourse entities are *discourse-new* but *mediated* by some *hearer-old* discourse entity (cf. Figure 1). I do not assume any difference between the elements of each set with respect to their information status. E.g., evoked and unused discourse entities have the same information status because both belong to OLD.

For an operationalization of Prince’s terms, I stipulate that *evoked* discourse entities are co-referring expressions (pronominal and nominal anaphora, previously mentioned proper names, relative pronouns, appositives). *Unused* discourse entities are




Figure 1: S-list Ranking and Familiarity

proper names and titles. In texts, *brand-new* proper names are usually accompanied by a relative clause or an appositive which relates them to the hearer’s knowledge. The corresponding discourse entity is *evoked* only after this elaboration. Whenever these linguistic devices are missing, proper names are treated as *unused*¹. I restrict *inferrables* to the particular subset defined by Hahn et al. (1996). *Anchored brand-new* discourse entities require that the anchor is either *evoked* or *unused*.

I assume the following conventions for the ranking constraints on the elements of the S-list. The 3-tuple (x, utt_x, pos_x) denotes a discourse entity x which is *evoked* in utterance utt_x at the text position pos_x . With respect to any two discourse entities (x, utt_x, pos_x) and (y, utt_y, pos_y) , utt_x and utt_y specifying the current utterance U_i or the preceding utterance U_{i-1} , I set up the following ordering constraints on elements in the S-list (Table 2)². For any state of the processor/hearer, the ordering of discourse entities in the S-list that can be derived from the ordering constraints (1) to (3) is denoted by the precedence relation \prec .

(1)	If $x \in \text{OLD}$ and $y \in \text{MED}$, then $x \prec y$. If $x \in \text{OLD}$ and $y \in \text{NEW}$, then $x \prec y$. If $x \in \text{MED}$ and $y \in \text{NEW}$, then $x \prec y$.
(2)	If $x, y \in \text{OLD}$, or $x, y \in \text{MED}$, or $x, y \in \text{NEW}$, then if $utt_x \succ utt_y$, then $x \prec y$, if $utt_x = utt_y$ and $pos_x < pos_y$, then $x \prec y$.

Table 2: Ranking Constraints on the S-list

Summarizing Table 2, I state the following preference ranking for discourse entities in U_i and U_{i-1} : *hearer-old* discourse entities in U_i , *hearer-old* discourse entities in U_{i-1} , *mediated* discourse entities in U_i , *mediated* discourse entities in U_{i-1} , *hearer-new* discourse entities in U_i , *hearer-new* discourse entities in U_{i-1} . By making the distinction in (2)

¹For examples of brand-new proper names and their introduction cf., e.g., the “obituaries” section of the *New York Times*.

²The relations \succ and $=$ indicate that the utterance containing x follows (\succ) the utterance containing y or that x and y are elements of the same utterance ($=$).

between discourse entities in U_i and discourse entities in U_{i-1} , I am able to deal with intra-sentential anaphora. There is no need for further specifications for complex sentences. A finer grained ordering is achieved by ranking discourse entities within each of the sets according to their text position.

3.3 The Algorithm

Anaphora resolution is performed with a simple look-up in the S-list³. The elements of the S-list are tested in the given order until one test succeeds. Just after an anaphoric expression is resolved, the S-list is updated. The algorithm processes a text from left to right (the unit of processing is the word):

1. If a referring expression is encountered,
 - (a) if it is a pronoun, test the elements of the S-list in the given order until the test succeeds⁴;
 - (b) update S-list; the position of the referring expression under consideration is determined by the S-list-ranking criteria which are used as an insertion algorithm.
2. If the analysis of utterance U^b is finished, remove all discourse entities from the S-list, which are not *realized* in U .

The analysis for example (1) is given in Table 3⁵. I show only these steps which are of interest for the computation of the S-list and the pronoun resolution. The preferences for pronouns (in bold font) are given by the S-list immediately above them. The pronoun “*she*” in (1b) is resolved to the first element of the S-list. When the pronoun “*her*” in (1c) is encountered, FRIEDMAN is the first element of the S-list since FRIEDMAN is *unused* and in the current utterance. Because of binding restrictions, “*her*” cannot be resolved to FRIEDMAN but to the second element, BRENNAN. In both (1d) and (1d) the pronoun “*she*” is resolved to FRIEDMAN.

³The S-list consists of referring expressions which are specified for text position, agreement, sortal information, and information status. Coordinated NPs are collected in a set. The S-list does not contain predicative NPs, pleonastic “*it*”, and any elements of direct speech enclosed in double quotes.

⁴The test for pronominal anaphora involves checking agreement criteria, binding and sortal constraints.

⁵I here define that an utterance is a sentence.

⁶In the following Tables, discourse entities are represented by SMALLCAPS, while the corresponding surface expression appears on the right side of the colon. Discourse entities are annotated with their information status. An “*e*” indicates an elliptical NP.

(1a)	Brennan drives an Alfa Romeo S: [BRENNAN _U : Brennan, ALFA ROMEO _{BN} : Alfa Romeo]
(1b)	She drives too fast. S: [BRENNAN _E : she]
(1c)	Friedman S: [FRIEDMAN _U : Friedman, BRENNAN _E : she] races her on weekends. S: [FRIEDMAN _U : Friedman, BRENNAN _E : her]
(1d)	She drives to Laguna Seca. S: [FRIEDMAN _E : she, LAGUNA SECA _U : Laguna Seca]
(1d')	She S: [FRIEDMAN _E : she, BRENNAN _E : her] often beats her . S: [FRIEDMAN _E : she, BRENNAN _E : her]

Table 3: Analysis for (1)

The difference between my algorithm and the BFP-algorithm becomes clearer when the *unused* discourse entity “Friedman” is replaced by a *brand-new* discourse entity, e.g., “a professional driver”⁷ (cf. example (2)). In the BFP-algorithm, the ranking of the *Cf*-list depends on grammatical roles. Hence, DRIVER is ranked higher than BRENNAN in the *Cf*(2c). In (2d), the pronoun “she” is resolved to BRENNAN because of the preference for CONTINUE over RETAIN. In (2d’), “she” is resolved to DRIVER because SMOOTH-SHIFT is preferred over ROUGH-SHIFT. In my algorithm, at the end of (2c) the *evoked* phrase “her” is ranked higher than the *brand-new* phrase “a professional driver” (cf. Table 4). In both (2d) and (2d’) the pronoun “she” is resolved to BRENNAN.

- (2) a. Brennan drives an Alfa Romeo.
b. **She** drives too fast.
c. A professional driver races **her** on weekends.
d. **She** goes to Laguna Seca.
d'. **She** often beats **her**.

Example (3)⁸ illustrates how the preferences for intra- and inter-sentential anaphora interact with the information status of discourse entities (Table 5). Sentence (3a) starts a new discourse segment. The phrase “a judge” is *brand-new*. “Mr. Curtis” is mentioned several times before in the text, Hence,

⁷I owe this variant Andrew Kehler. – This example can misdirect readers because the phrase “a professional driver” is assigned the “default” gender masculine. Anyway, this example – like the original example – seems not to be felicitous English and has only illustrative character.

⁸In: *The New York Times*. Dec. 7, 1997, p.A48 (“Shot in head, suspect goes free, then to college”).

(2a)	Brennan drives an Alfa Romeo S: [BRENNAN _U : Brennan, ALFA ROMEO _{BN} : Alfa Romeo]
(2b)	She drives too fast. S: [BRENNAN _E : she]
(2c)	A professional driver S: [BRENNAN _E : she, DRIVER _{BN} : Driver] races her on weekends. S: [BRENNAN _E : her, DRIVER _{BN} : Driver]
(2d)	She drives to Laguna Seca. S: [BRENNAN _E : she, LAGUNA SECA _U : Laguna Seca]
(2d')	She S: [BRENNAN _E : she, DRIVER _{BN} : Driver] often beats her . S: [BRENNAN _E : she, DRIVER _E : her]

Table 4: Analysis for (2)

the discourse entity CURTIS is *evoked* and ranked higher than the discourse entity JUDGE. In the next step, the ellipsis refers to JUDGE which is evoked then. The nouns “request” and “prosecutors” are *brand-new*⁹. The pronoun “he” and the possessive pronoun “his” are resolved to CURTIS. “Condition” is *brand-new* but anchored by the possessive pronoun. For (3b) and (3c) I show only the steps immediately before the pronouns are resolved. In (3b) both “Mr. Curtis” and “the judge” are *evoked*. However, “Mr. Curtis” is the left-most *evoked* phrase in this sentence and therefore the most preferred antecedent for the pronoun “him”. For my experiments I restricted the length of the S-list to five elements. Therefore “prosecutors” in (3b) is not contained in the S-list. The discourse entity SMIRGA is introduced in (3c). It becomes *evoked* after the appositive. Hence SMIRGA is the most preferred antecedent for the pronoun “he”.

- (3) a. A judge ordered that Mr. Curtis be released, but ϵ agreed with a request from prosecutors that **he** be re-examined each year to see if **his** condition has improved.
b. But authorities lost contact with Mr. Curtis after the Connecticut Supreme Court ruled in 1990 that the judge had erred, and that prosecutors had no right to re-examine **him**.
c. John Smirga, the assistant state’s attorney in charge of the original case, said last week that **he** always had doubts about the psychiatric reports that said Mr. Curtis would never improve.

⁹I restrict *inferrables* to the cases specified by Hahn et al. (1996). Therefore “prosecutors” is *brand-new* (cf. Prince (1992) for a discussion of the form of inferrables).

(3a)	<p>A judge S: [JUDGE_{BN}: judge] ordered that Mr. Curtis S: [CURTIS_E: Mr. Curtis, JUDGE_{BN}: judge] be released, but ϵ S: [CURTIS_E: Mr. Curtis, JUDGE_E: ϵ] agreed with a request S: [CURTIS_E: Mr. Curtis, JUDGE_E: ϵ, REQUEST_{BN}: request] from prosecutors S: [CURTIS_E: Mr. Curtis, JUDGE_E: ϵ, REQUEST_{BN}: request, PROSECUTORS_{BN}: prosecutors] that he S: [CURTIS_E: he, JUDGE_E: ϵ, REQUEST_{BN}: request, PROSECUTORS_{BN}: prosecutors] be re-examined each year S: [CURTIS_E: he, JUDGE_E: ϵ, REQUEST_{BN}: request, PROSECUTORS_{BN}: prosecutors, YEAR_{BN}: year] to see if his S: [CURTIS_E: his, JUDGE_E: ϵ, REQUEST_{BN}: request, PROSECUTORS_{BN}: prosecutors, YEAR_{BN}: year] condition S: [CURTIS_E: his, JUDGE_E: ϵ, CONDITION_{BN_A}: condition, REQUEST_{BN}: request, PROSECUTORS_{BN}: prosec.] has improved. S: [CURTIS_E: his, JUDGE_E: ϵ, CONDITION_{BN_A}: condition, REQUEST_{BN}: request, PROSECUTORS_{BN}: prosec.]</p>
(3b)	<p>But authorities lost contact with Mr. Curtis after the Connecticut Supreme Court ruled in 1990 that the judge had erred, and that prosecutors had no right S: [CURTIS_E: his, CS COURT_U: CS Court, JUDGE_E: judge, CONDITION_{BN_A}: condition, AUTH._{BN}: auth.] to re-examine him. S: [CURTIS_E: him, CS COURT_U: CS Court, JUDGE_E: judge, CONDITION_{BN_A}: condition, AUTH._{BN}: auth.]</p>
(3c)	<p>John Smirga, the assistant state's attorney in charge of the original case, said last week S: [SMIRGA_E: attorney, CASE_E: case, CURTIS_E: him, CS COURT_U: CS Court, JUDGE_E: judge] that he had doubts about the psychiatric reports that said Mr. Curtis would never improve. S: [SMIRGA_E: he, CASE_E: case, REPORTS_E: reports, CURTIS_E: Mr. Curtis, DOUBTS_{BN}: doubts]</p>

Table 5: Analysis for (3)

4 Some Empirical Data

In the first experiment, I compare my algorithm with the BFP-algorithm which was in a second experiment extended by the constraints for complex sentences as described by Kameyama (1998).

Method. I use the following guidelines for the hand-simulated analysis (Walker, 1989). I do not assume any world knowledge as part of the anaphora resolution process. Only agreement criteria, binding and sortal constraints are applied. I do not account for false positives and error chains. Following Walker (1989), a segment is defined as a paragraph unless its first sentence has a pronoun in subject position or a pronoun where none of the preceding sentence-internal noun phrases matches its syntactic features. At the beginning of a segment, anaphora resolution is preferentially performed within the same utterance. My algorithm starts with an empty S-list at the beginning of a segment.

The basic unit for which the centering data structures are generated is the utterance U . For the BFP-algorithm, I define U as a simple sentence, a complex sentence, or each full clause of a compound sentence. Kameyama's (1998) intra-sentential centering operates at the clause level. While tensed

clauses are defined as utterances on their own, untensed clauses are processed with the main clause, so that the Cf -list of the main clause contains the elements of the untensed embedded clause. Kameyama distinguishes for tensed clauses further between sequential and hierarchical centering. Except for reported speech (embedded and inaccessible to the superordinate level), non-report complements, and relative clauses (both embedded but accessible to the superordinate level; less salient than the higher levels), all other types of tensed clauses build a chain of utterances on the same level.

According to the preference for inter-sentential candidates in the centering model, I define the following anaphora resolution strategy for the BFP-algorithm: (1) Test elements of U_{i-1} . (2) Test elements of U_i left-to-right. (3) Test elements of $Cf(U_{i-2}), Cf(U_{i-3}), \dots$ In my algorithm steps (1) and (2) fall together. (3) is performed using previous states of the system.

Results. The test set consisted of the beginnings of three short stories by Hemingway (2785 words, 153 sentences) and three articles from the *New York Times* (4546 words, 233 sentences). The results of my experiments are given in Table 6. The

first row gives the number of personal and possessive pronouns. The remainder of the Table shows the results for the BFP-algorithm, for the BFP-algorithm extended by Kameyama’s intra-sentential specifications, and for my algorithm. The overall error rate of each approach is given in the rows marked with *wrong*. The rows marked with *wrong (strat.)* give the numbers of errors directly produced by the algorithms’ strategy, the rows marked with *wrong (ambig.)* the number of analyses with ambiguities generated by the BFP-algorithm (my approach does not generate ambiguities). The rows marked with *wrong (intra)* give the number of errors caused by (missing) specifications for intra-sentential anaphora. Since my algorithm integrates the specifications for intra-sentential anaphora, I count these errors as strategic errors. The rows marked with *wrong (chain)* give the numbers of errors contained in error chains. The rows marked with *wrong (other)* give the numbers of the remaining errors (consisting of pronouns with split antecedents, errors because of segment boundaries, and missing specifications for event anaphora).

		Hem.	NYT	Σ
Pron. and Poss. Pron.		274	302	576
BFP-Algo.	Correct	189	231	420
	Wrong	85	71	156
	Wrong (strat.)	14	2	16
	Wrong (ambig.)	9	15	24
	Wrong (intra)	17	13	30
	Wrong (chain)	29	32	61
BFP/Kam.	Correct	193	245	438
	Wrong	81	57	138
	Wrong (strat.)	3	0	3
	Wrong (ambig.)	17	8	25
	Wrong (intra)	17	27	44
	Wrong (chain)	29	15	44
My Algo.	Correct	217	275	492
	Wrong	57	27	84
	Wrong (strat.)	21	12	33
	Wrong (chain)	22	9	31
	Wrong (other)	14	6	20

Table 6: Evaluation Results

Interpretation. The results of my experiments showed not only that my algorithm performed better than the centering approaches but also revealed insight in the interaction between inter- and intra-sentential preferences for anaphoric antecedents. Kameyama’s specifications reduce the complexity in that the C_f -lists in general are shorter after splitting up a sentence into clauses. Therefore, the

BFP-algorithm combined with her specifications has almost no strategic errors while the number of ambiguities remains constant. But this benefit is achieved at the expense of more errors caused by the intra-sentential specifications. These errors occur in cases like example (3), in which Kameyama’s intra-sentential strategy makes the correct antecedent less salient, indicating that a clause-based approach is too fine-grained and that the hierarchical syntactical structure as assumed by Kameyama does not have a great impact on anaphora resolution.

I noted, too, that the BFP-algorithm can generate ambiguous readings for U_i when the pronoun in U_i does not co-specify the $Cb(U_{i-1})$. In cases, where the $C_f(U_{i-1})$ contains more than one possible antecedent for the pronoun, several ambiguous readings with the same transitions are generated. An example¹⁰: There is no $Cb(4a)$ because no element of the preceding utterance is realized in (4a). The pronoun “them” in (4b) co-specifies “deer” but the BFP-algorithm generates two readings both of which are marked by a RETAIN transition.

- (4) a. Jim pulled the burlap sacks off the deer
b. and Liz looked at **them**.

In general, the strength of the centering model is that it is possible to use the $Cb(U_{i-1})$ as the most preferred antecedent for a pronoun in U_i . In my model this effect is achieved by the preference for hearer-old discourse entities. Whenever this preference is misleading both approaches give wrong results. Since the Cb is defined strictly local while hearer-old discourse entities are defined global, my model produces less errors. In my model the preference is available immediately while the BFP-algorithm can use its preference not before the second utterance has been processed. The more global definition of hearer-old discourse entities leads also to shorter error chains. – However, the test set is too small to draw final conclusions, but at least for the texts analyzed the preference for *hearer-old* discourse entities is more appropriate than the preference given by the BFP-algorithm.

5 Comparison to Related Approaches

Kameyama’s (1998) version of centering also omits the *centering transitions*. But she uses the Cb and a ranking over simplified transitions preventing the incremental application of her model.

¹⁰In: Ernest Hemingway. *Up in Michigan*. In. *The Complete Short Stories of Ernest Hemingway*. New York: Charles Scribner’s Sons, 1987, p.60.

The focus model (Sidner, 1983; Suri & McCoy, 1994) accounts for *evoked* discourse entities explicitly because it uses the discourse focus, which is determined by a successful anaphora resolution. Incremental processing is not a topic of these papers.

Even models which use salience measures for determining the antecedents of pronoun use the concept of *evoked* discourse entities. Hajičová et al. (1992) assign the highest value to an evoked discourse entity. Also Lappin & Leass (1994), who give the subject of the current sentence the highest weight, have an implicit notion of *evokedness*. The salience weight degrades from one sentence to another by a factor of two which implies that a repeatedly mentioned discourse entity gets a higher weight than a *brand-new* subject.

6 Conclusions

In this paper, I proposed a model for determining the hearer's attentional state which is based on the distinction between *hearer-old* and *hearer-new* discourse entities. I showed that my model, though it omits the *backward-looking center* and the *centering transitions*, does not lose any of the predictive power of the centering model with respect to anaphora resolution. In contrast to the centering model, my model includes a treatment for intrasentential anaphora and is sufficiently well specified to be applied to real texts. Its incremental character seems to be an answer to the question Kehler (1997) recently raised. Furthermore, it neither has the problem of inconsistency Kehler mentioned with respect to the BFP-algorithm nor does it generate unnecessary ambiguities.

Future work will address whether the text position, which is the weakest grammatical concept, is sufficient for the order of the elements of the S-list at the second layer of my ranking constraints. I will also try to extend my model for the analysis of definite noun phrases for which it is necessary to integrate it into a more global model of discourse processing.

Acknowledgments: This work has been funded by a post-doctoral grant from DFG (Str 545/1-1) and is supported by a post-doctoral fellowship award from IRCS. I would like to thank Nobo Komagata, Rashmi Prasad, and Matthew Stone who commented on earlier drafts of this paper. I am grateful for valuable comments by Barbara Grosz, Udo Hahn, Aravind Joshi, Lauri Karttunen, Andrew Kehler, Ellen Prince, and Bonnie Webber.

References

- Brennan, S. E., M. W. Friedman & C. J. Pollard (1987). A centering approach to pronouns. In *Proc. of the 25th Annual Meeting of the Association for Computational Linguistics; Stanford, Cal., 6–9 July 1987*, pp. 155–162.
- Grosz, B. J., A. K. Joshi & S. Weinstein (1983). Providing a unified account of definite noun phrases in discourse. In *Proc. of the 21st Annual Meeting of the Association for Computational Linguistics; Cambridge, Mass., 15–17 June 1983*, pp. 44–50.
- Grosz, B. J., A. K. Joshi & S. Weinstein (1995). Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21(2):203–225.
- Hahn, U., K. Markert & M. Strube (1996). A conceptual reasoning approach to textual ellipsis. In *Proc. of the 12th European Conference on Artificial Intelligence (ECAI '96); Budapest, Hungary, 12–16 August 1996*, pp. 572–576. Chichester: John Wiley.
- Hajičová, E., V. Kuboň & P. Kuboň (1992). Stock of shared knowledge: A tool for solving pronominal anaphora. In *Proc. of the 14th Int. Conference on Computational Linguistics; Nantes, France, 23–28 August 1992*, Vol. 1, pp. 127–133.
- Kameyama, M. (1998). Intrasentential centering: A case study. In M. Walker, A. Joshi & E. Prince (Eds.), *Centering Theory in Discourse*, pp. 89–112. Oxford, U.K.: Oxford Univ. Pr.
- Kehler, A. (1997). Current theories of centering for pronoun interpretation: A critical evaluation. *Computational Linguistics*, 23(3):467–475.
- Lappin, S. & H. J. Leass (1994). An algorithm for pronominal anaphora resolution. *Computational Linguistics*, 20(4):535–561.
- Prince, E. F. (1981). Toward a taxonomy of given-new information. In P. Cole (Ed.), *Radical Pragmatics*, pp. 223–255. New York, N.Y.: Academic Press.
- Prince, E. F. (1992). The ZPG letter: Subjects, definiteness, and information-status. In W. Mann & S. Thompson (Eds.), *Discourse Description. Diverse Linguistic Analyses of a Fund-Raising Text*, pp. 295–325. Amsterdam: John Benjamins.
- Sidner, C. L. (1983). Focusing in the comprehension of definite anaphora. In M. Brady & R. Berwick (Eds.), *Computational Models of Discourse*, pp. 267–330. Cambridge, Mass.: MIT Press.
- Strube, M. & U. Hahn (1996). Functional centering. In *Proc. of the 34th Annual Meeting of the Association for Computational Linguistics; Santa Cruz, Cal., 23–28 June 1996*, pp. 270–277.
- Suri, L. Z. & K. F. McCoy (1994). RAFT/RAPR and centering: A comparison and discussion of problems related to processing complex sentences. *Computational Linguistics*, 20(2):301–317.
- Walker, M. A. (1989). Evaluating discourse processing algorithms. In *Proc. of the 27th Annual Meeting of the Association for Computational Linguistics; Vancouver, B.C., Canada, 26–29 June 1989*, pp. 251–261.
- Walker, M. A., M. Iida & S. Cote (1994). Japanese discourse and the process of centering. *Computational Linguistics*, 20(2):193–233.