DEACCENTUATION AND PERSISTENCE OF GRAMMATICAL FUNCTION AND SURFACE POSITION

Draft: 10 Jan 1994

Jacques Terken
Institute for Perception Research
Eindhoven, The Netherlands

Julia Hirschberg
AT&T Bell Laboratories
Murray Hill, New Jersey

Mail to:
Jacques Terken
Institute for Perception Research
P.O.Box 513, 5600 MB Eindhoven, The Netherlands
tel. +31.40 773 844
e-mail: terken@prl.philips.nl

Short title:
Deaccentuation, Grammatical Function and Surface Position

Subject Index:
Prosody, assignment of sentence accents, grammatical function and surface position

Abstract

The absence of intonational prominence on a referring expression (DEACCENTUATION) is commonly explained as a consequence of the GIVENNESS of the discourse entity referred to, or, the fact that it represents 'old' information in the discourse. However, speakers sometimes use **accented** expressions to refer to such GIVEN entities, so that GIVENNESS is no sufficient explanation for DEACCENTUATION and a more refined explanation is required. It has also been suggested that speakers tend to express GIVEN entities as grammatical subjects and to put them in early in the utterance. The present work investigates the contribution of grammatical role and surface position to the occurrence of deaccentuation in English. An experiment is reported in which speakers produced descriptions of visual materials, where the content of the materials was manipulated so that successive descriptions contained coreferential expressions in which grammatical role and surface position were systematically varied. The results indicated that persistence of grammatical role and surface position both contributed to deaccentuation. Some implications for the way in which listeners link incoming expressions to entities which are already available from the context are discussed.

1 Introduction

The distribution of ACCENTS in a discourse has often been explained with respect to the INFORMATION STATUS of the entities being mentioned in the discourse: Words or word groups expressing GIVEN information tend to be DEACCENTED, whereas words (or word groups) expressing NEW information tend to be accented [Chafe, 1976, Clark and Clark, 1977, Prince, 1981, Prince, 1992], where GIVEN means that the information has been mentioned in the preceding discourse or is inferrable from such information or from the situation. However, speakers do not always deaccent GIVEN information, which implies that GIVENNESS is not a sufficient condition for the occurrence of an unaccented expression. In this paper, we investigate some conditions under which speakers use accented and deaccented expressions to refer to given information, in order to get a better understanding of the underlying factors governing the speakers' deaccentuation behavior. The term ACCENT is used here to mean "intonational prominence". In English, it is characterized by a conspicuous pitch change in or near the lexically stressable syllable of the word [1] [Pierrehumbert, 1980, 't Hart et al., 1990]. Accented words also usually exhibit longer duration than their unaccented counterparts, as well as greater amplitude [Fry, 1955, Fry, 1958, Lehiste, 1970].

The association of information status with accent status is often illustrated via the QUESTION TEST: The utterance of (1a) with *book* accented, represents a felicitous response to (1b), but not to (1c).

(1) a. John is reading a BOOK.

b. What is John reading?

c. Who is reading a book?

[1] Barring cases of accent placement due to metalinguistic contrast (e.g., *I didn't say DEgrade, I said UPgrade.*) or stress shift in complex nominals (e.g., *FIFteen CHInese JUGglers*). Here and below, capitals mark the location of prominent syllables.

The accenting of *book* appears odd when *book* has just been mentioned, because an expression which refers to an item that has been mentioned in the previous discourse should remain unaccented and expressions referring to an item which has not yet been mentioned should be accented.

However, turning from such simple examples to naturally occurring discourse it becomes immediately obvious that there is no simple one-to-one mapping between GIVENNESS and the use of an unaccented expression, for at least two reasons.

First, the current utterance may impose requirements which overrule the discourse effects, such as in the following example.

(2) President Clinton has arrived for a three-day visit in London.
The president, who also visited London last year, ...

The second occurrence of 'president' must be accented, even though it clearly refers to GIVEN information, due to the phonological constraint that every prosodic phrase must contain at least one accent. The question of how the properties of the current utterance affect the assignment of accents is dealt with by FOCUS THEORY [Ladd, 1980, Gussenhoven, 1983, Baart, 1987, Dirksen, 1992]. Since our present concern is with discourse effects, we will leave this issue out of consideration.

Second, the definition of GIVEN as "mentioned in the previous discourse or inferrable from this information or from the situation" [Haviland and Clark, 1974, Prince, 1981] is very broad, and in addition to accentuation it has been applied to a variety of linguistic phenomena which are also sensitive to discourse influences, such as the form of expressions and their position in the utterance. Empirical research has shown that such a broad definition is not appropriate for accentuation [Brown, 1983, Terken, 1985]. For instance, Brown found that all expressions referring to items which had been mentioned in the previous discourse were unaccented, but that expressions referring to inferrable items were usually accented. She adds a caveat that the discourses were quite short and concerned relatively few entities, so that all entities may have been highly accessible to the hearer. Terken found that, in task-oriented monologues, speakers used unaccented, pronominal expressions to refer to the local topic of discourse, and accented, full NPs otherwise, even though many of these NPs referred to entities which had already been mentioned in the previous discourse.

Thus, a more stringent criterion than "mentioned in or inferrable from the previous discourse" is needed to determine which items may be referred to with deaccented expressions. Several such refinements have been proposed. For example, Halliday has argued that an expression may be left unaccented if the information conveyed by the expression is situationally or anaphorically recoverable on the basis of the prior discourse or by being salient in the situation [Halliday, 1967]. Chafe, taking a more cognitive approach, suggested that an expression may be unaccented if the information is in the listener's consciousness [Chafe, 1974, Chafe, 1976]. It seems reasonable to assume that not all items which have been mentioned previously in a discourse of some length are recoverable anaphorically or are in the listener's consciousness.

These attempts to further delineate the relationship between accenting behavior and the information status of mentioned discourse entities share a common intuition, that the use of an unaccented expression may be understood in terms of the relative ACCESSIBILITY of entities in the discourse model. Although the notion of ACCESSIBILITY is not clearly defined, it is intuitively clear that not all information in the discourse model is equally accessible at a particular time — to serve as the referent of a pronoun, for instance – but that particular entities will represent much more likely referents than others.

In order to explore the possibility that the accessibility of entities in the discourse model is relevant to accentuation, we have to link the notion of accessibility to observable properties of utterances, since it cannot be measured directly. The relation between the relative accessibility of information in a discourse and a number of observable properties of utterances has been broadly explored in theories of COMMUNICATIVE DYNAMISM, ATTENTIONAL FOCUSSING and CENTERING in discourse [Grosz, 1977, Sidner, 1979, Grosz, 1981, Sidner, 1983, Grosz et al., 1983, Grosz and Sidner, 1986, Chafe, 1974, Kameyama, 1986, Brennan et al., 1987, Asher and Wada, 1988, Hajicova et al., 1990, Gordon et al., 1993, Gundel et al., 1993], and in models of sentence production [Bock and Irwin, 1980, Bock, 1982, Bock and Warren, 1985]. The available evidence supports the notion that the relative accessibility of entities in the discourse model is a major factor in the assignment to grammatical role and surface position, and in the choice of the form of referring expressions: Highly accessible entities tend to be realized as the grammatical subject, to occur early in the utterance and to be pronominalized. Furthermore, available evidence from studies on comprehension shows that accessibility is also an important factor in the way the listener processes the incoming message [Kameyama, 1986, Gordon et al., 1993, McKoon et al., 1993]. Kameyama proposes a "property-sharing constraint" governing the interpretation of pronominal expressions in adjacent utterances: If two pronominal expressions in adjacent utterances share a certain common grammatical property, e.g. subject, and the first one refers to the most accessible discourse entity, the second one will be interpreted as referring to that entity also [Kameyama, 1986].

If accessibility is what really matters for deaccentuation, one may expect that these observable properties also bear a meaningful relationship to the use of deaccentuation. Indeed, Terken (1984) found a strong correlation between deaccentuation and the use of pronominal expressions. In addition, it was found that deaccented expressions tend to be associated preferably with the function of grammatical subject and to occur in initial position in the surface structure. Due to the limited set-up of the study, however, the relevance of the individual factors could not be assessed independently. In the current study, we aim to further explore the possibility that accessibility is a major factor governing the use of deaccentuation, by systematically manipulating the assignment of discourse entities to specific grammatical roles and surface positions in adjacent utterances and investigating the effects on speakers' decisions to accent or deaccent references to GIVEN items in a discourse.

2 The Experiment

For the purpose of the current experiment, we will investigate items that have been evoked by explicit mention in the immediately preceding discourse, confining our definition of GIVENNESS to this constrained sense. While our definition may not encompass all categories of this phenomenon, it certainly captures the core sense of the term, which all definitions of GIVENNESS must include.

In order to test the relevance of grammatical function and surface position assignment in adjacent utterances to deaccentuation, we compare speakers' accentuation behavior with respect to an expression **x** referring to an entity **X** in a target utterance **i** under the following scenarios: a) **X**, the referent of **x**, has **not** been mentioned in the immediate context; b) **X has** been mentioned in the immediate context, with the **same** grammatical function and/or surface position it exhibits in **i**; and c) **X** has been mentioned in the immediate context, but with a **different** grammatical function and/or surface position than those it exhibits in **i**. By manipulating grammatical function and surface position independently, we will be able to assess the relative effects of these factors.

To investigate the contribution of grammatical function and surface position while controlling for other factors which may also influence accent assignment, and to maintain some degree of spontaneity, we chose an elicitation technique for our study. We elicited spontaneous productions by presenting speakers with visual materials and asking them to describe changes in the observed scene. We manipulated properties of the visual materials in such a way that we obtained the same descriptions in different contexts, thus allowing for comparison of items with the same lexical contents both within and between speakers. In this way, we overcome the problems inherent to the analysis of corpora of spontaneous speech, such as the lack of control over the context in which utterances are produced. Despite the degree of predictability this type of elicitation produces, the task exhibits the desirable property that the speaker must construct each utterance himself. Although the content of the utterance is dependent upon the materials displayed and the task, the actual formulation process is under the control of the speaker. The technique has been applied before in the study of prosody by [Terken, 1985] and by [Swerts and Collier, 1992].

In this way, we collected sets of utterances for a number of speakers and investigated how the assignment of accents in these productions was influenced by properties of utterances and their prior context. In particular, we wanted to answer the following questions.

- To what extent does plain GIVENNESS (defined as simple prior mention in the preceding context) account for the occurrence of an unaccented expression in an utterance?
- To what extent does the grammatical role of an expression referring to a GIVEN entity affect its accent status, and to what extent is this influence dependent on the grammatical role of the antecedent expression in the context?
- To what extent does the surface position of an expression referring to a GIVEN entity affect its accent status, and to what extent is this influence dependent on the surface position of the antecedent expression in the context?

2.1 Method

2.1.1 Task

Speech was elicited from each subject in the following manner: The speaker was seated in a recording booth in front of a workstation screen displaying a simple configuration of geometrical shapes such as a circle and a triangle. The position of objects in the display could be changed in ways that had been predefined, such that one or more objects appeared to move to different positions. The timing of such events was under the control of the experimenter. The speaker's task was to describe these events to a listener watching a second screen containing the same lay-out. Since the movements were few in type and very simple, they could be described in a simple, straightforward manner. By an appropriate selection of the objects involved in each event and the role each object played in the event, and by manipulating the way in which events were sequenced, we obtained utterances in which grammatical function, surface position, and information status (GIVEN vs. NEW) were systematically varied.

2.1.2 Materials

Each subject described 75 SCENARIOS, each consisting of a sequence of four events; each event consisted of a change in the position of one or two objects in the visual display relative to one of the other objects in the display. By systematically varying the role of the objects involved in the events, the descriptions of these events given by our speakers could be manipulated indirectly to vary the grammatical function of the referring expressions, in the following way. The visual display always contained nine objects — a ball, a cone, a cross, a star, a line, a box, a diamond, a rectangle, and a triangle. For ease of reference — and to encourage lexical uniformity — each object was explicitly labelled with its name. These names remained visible throughout the experiment.

Three types of events were simulated: (a) object_1 could approach and touch object_2, (b) object_1 could approach and cover object_2, and (c) object_1 could approach object_2 and push it against object_3. So, for instance, the ball could approach and touch the cone (a), the ball could approach and cover the cone (b), or the ball could approach the cone and push it against the cross (c). Objects performing particular roles in the events associate naturally with particular grammatical functions in the descriptions. For instance, the moving or pushing object is naturally assigned to the function of grammatical subject, Thus, typical descriptions of these events were: (a) *The ball touches the cone*, (b) *The ball covers the cone*, and (c) *The ball pushes the cone against the cross.* Note that all three sample descriptions contain a grammatical subject and a direct object; description (c) also contains the object of a preposition. Only the utterance describing the fourth and last event in each scenario was used for analysis. We term this utterance the TARGET UTTERANCE for our analysis, and the first three utterances in each discourse, the CONTEXT. The screen was cleared between scenarios, so that the sequence of descriptions for each individual scenario would constitute a separate discourse.

The information status of items mentioned in the target utterances was manipulated by varying the objects and the roles they played in successive events, thereby influencing the contents of successive utterances. Table 1 shows some sample context event sequences with typical accompanying context utterances, and three sample target events and utterances, which might follow any of the three contexts. Table 1 shows three alternative Context Event sequences, a, b, and c. Each context consists of three successive changes (context events CE1–CE3) to an initial display (I) and the corresponding descriptions of the changes (CD1–CD3) [2]. Three alterna-

[2]In this study, the context events and their descriptions served as the non-linguistic and linguistic context for the target event and its description, the target utterance. Since the context descriptions were accurate representations of the context events, we assume

tive target events (TEa–c) are also presented, with accompanying target descriptions (TDa–c). The referring expressions in the target utterances can be classified as either GIVEN or NEW, relative to the context. That is, regardless of how we combine targets and contexts in Table 1, the ball is mentioned in the context in all three Context Event sequences, and thus will be GIVEN (under our definition) in the target utterance. However, the other items mentioned in the target utterances (e.g., '*the star*', '*the box*') are **not** mentioned in the context and thus represent NEW information in the target utterance [3]. Note further that the GIVEN items in the target utterances are mentioned in all context utterances for the current scenario. Furthermore, as the GIVEN item plays the same role in all three context events, the corresponding expression has the same grammatical function and surface position throughout the context [4]. Thus, we assume that these items will indeed be considered "highly accessible" by the speaker by the time the target utterance is produced [Terken and Nooteboom, 1987]. The GIVEN items in the target utterances are the TARGET EXPRESSIONS for our analysis. That is, *the ball* is the TARGET EXPRESSION in all of the target utterances in Table 1.

Orthogonal combination of targets and contexts in Table 1 produces nine different scenarios. The scenarios differ with respect to the grammatical function of the GIVEN item in the target and context. For example, the expression *the ball* has the function of grammatical subject, object and prepositional object in contexts a, b and c, respectively. Likewise, *the ball* fulfills these three functions in targets a, b and c, respectively. We thus obtained a 3x3 matrix with grammatical function (SUBJ, DIRECT OBJ and

that the events and the descriptions were treated by our subjects as merged in the discourse model. Therefore, we do not distinguish in our discussion between the textual and situational context.

[3] Although these items may have been mentioned in preceding scenarios, we assume that each scenario defines a separate discourse.

[4] Alternatively, speakers might have chosen to assign GIVEN items to the role of grammatical subject and/or to initial position, so that grammatical role and/or surface position would vary in the context descriptions, but this did not occur in our materials.

PP-OBJ) in the context and in the target description as categorization factors. (In the following discussion we will identify these conditions in the form 'A→B', where 'A' gives the function in the prior context and 'B' gives the function in the target utterance. So 'SUBJ→DIRECT OBJ' identifies target expressions that are subjects in a prior context and direct objects in the target utterance.) In this way, we can study the effect of a change or persistence of grammatical function and the contribution of the grammatical functions involved, e.g. by pairwise comparison of SUBJ→SUBJ and DIRECT OBJ→SUBJ, and DIRECT OBJ→SUBJ and SUBJ→DIRECT OBJ. The possible contribution of properties of the current utterance to accentuation is kept constant across conditions, since each target utterance is combined with three different context types.

To a certain extent, we were also able to separate the effects of a change or persistence of grammatical function from the effects of change or persistence of surface position in the following way. Table 1 illustrates two of the types of target utterances we elicited, one containing only SUBJ and DIR OBJ (TDa and TDb), and the other containing a PP-OBJ as well (TDc). TDc was elicited only when the PP-OBJ was the target expression. Similarly, for the context utterances: Context c was elicited only when the GIVEN discourse entity figured as PP-OBJ in the context. I.e, for Z→SUBJ and Z→DIRECT OBJ (where Z is any grammatical function) only target utterances of types TDa and TDb were used. Likewise, for SUBJ→PP-OBJ and DIRECT OBJ→PP-OBJ only context utterances of types CDa and CDb were used. Hence, in conditions DIRECT OBJ→PP-OBJ (context b + target c) and PP-OBJ→DIRECT OBJ (context c + target b), the grammatical functions of the GIVEN expression in context and target utterance were different, but their positions in the surface structure were the same — both represented the last (non-empty) constituent in their utterance. By comparing these conditions with conditions in which either both grammatical function and surface position remain the same between context and target, and with conditions in which both grammatical function and surface position differ between target and context, we can assess the effect of change or persistence of surface position separately from the effect of grammatical function.

For each cell of the 3x3 design there were five scenarios, for a total of 45 target expressions elicited per subject. In addition, target descriptions a and b were combined with different contexts, in order to create two control conditions. One control was an ALL NEW condition, in which the expressions in the target description were not mentioned in the prior context, as illustrated in Example (3).

(3) CDb1: The cone touches the ball.
CDb2: The cross touches the ball.
CDb3: The diamond touches the ball.
TD: The rectangle touches the square.

The other control was an ALL GIVEN condition, in which the subject and direct object of the last context description (CD3) were also subject and direct object in the target description, so that only the verb represented NEW information, as illustrated in Example (4).

(4) CDb1: The cone touches the ball.
CDb2: The cross touches the ball.
CDb3: The diamond touches the ball.
TD: The diamond covers the ball.

No controls were included for target expressions containing a PP-OBJ. There were fifteen scenarios for each control condition.

2.1.3 Subjects

Ten male native speakers of standard American English, employed at AT&T Bell Laboratories in Murray Hill, New Jersey, participated as volunteers.

2.1.4 Procedure

Each subject was treated individually and received oral and written instructions. He was instructed that he was participating in an experiment aimed at improving the way computers provide spoken information to humans. He was informed about the basic structure of the experiment, about the possible events he would see in the display, and about the format of the scenarios. He was then given examples of possible descriptions, in order to inform him about the degree of precision needed in the descriptions. He was informed that there would be someone listening to him to verify his descriptions relative to the events in another display, and was introduced to the listener, but he was told that there would be no interaction or feedback, either auditory or visual. He was told that the best speaking mode would be as if telling a story or describing a movie he had seen, to the listener.

During the experiment, both subject and experimenter were in a soundproof booth. The subject was seated before the screen of a SUN 361 workstation. The actual events and their order were predefined and read from file. The pacing of events was controlled by the experimenter: following the speaker's description of an event in the display, the experimenter pressed the return key of the workstation keyboard, triggering the display of the next event. The events all involved a change of the relevant object's or objects' position in a number of steps to a new position. Between scenarios, the screen was cleared. The actual experiment was preceded by ten practice scenarios, to familiarize the subject with the task. Upon request, the experimenter made clear that object color and size and relative orientation of an object relative to another object were not relevant. No subject reported difficulty with the task.

The actual experiment took about twenty minutes for each subject. Halfway through the scenarios there was a short break. Subjects were also told that they could take a break whenever they wished, but few subjects did so. Some subjects tended to speak in a very low voice, possibly due to the artificiality of the communicative situation or initial hesitancy about the recording situation. In such cases, they were encouraged to speak up a bit and they were reminded that they were actually speaking to a listener.

The scenarios were presented in a different random order for each subject, to balance for order effects. The speech was recorded by means of a Bruel and Kjaer microphone 2231. The descriptions were stored on DAT (sampling frequency 48 kHz).

3 Results and Discussion

3.1 Prosodic Analysis

The descriptions obtained had the same format as in Table 1: simple active sentences of the form *the* NOUN VERB*s the* NOUN or *the* NOUN VERB*s the* NOUN PREP *the* NOUN. Formal devices which are commonly employed by speakers to establish discourse coherence, such as the occurrence of pronouns or passive constructions, did not occur in our materials. We deliberately discouraged the use of pronominal expressions by explictly labelling objects to be described in the experiment with their names. One might object that this artificiality reduces the usefulness of the results as data relevant to the question of why speakers use DE-ACCENTUATION to mark discourse coherence. However, it was necessary to attempt to limit the number of factors involved in the experiment in order to study the relationships among a manageable few, which could be manipulated in a principled way.

Each subject's utterances were stored on disk in separate files and down-sampled to 12 KHz. The target utterances were then excised from the context for further analysis. All target utterances were analysed by ear by the two authors independently, to determine the location of pitch accents, in such a way that they had no information about the context of a target description, the condition to which a target item belonged, or the judgment of the other judge. For each referring expression in a target utterance each judge indicated whether the expression was accented or not on a three-point scale: 0 *unaccented*, 1 *uncertain whether accented or unaccented*, and 2 *accented*. Summing the scores for the two judges gave a score between 0 and 4 for each referring expression. These scores will be termed **accentedness scores**. This procedure combines accentedness judgments and confidence reatings, and appears to be relatively free of theoretical and response biases.

3.1.1 Agreement between Judges

Table 2 shows the data for agreement between judges on the accentedness of referring expressions in target utterances. The marginals in Table 2 show that both judges judged about the same percentage of expressions to be accented; the two judges differed, however, in the percentage of expressions they judged deaccented. These facts suggest that the judges' thresholds for deciding that an expression was accented are about the same, but that judge A is more cautious than judge B in classifying an expression as unaccented. Next, consider the percentages for full agreement (summing cells 00, 11 and 22), moderate disagreement (01, 10, 12 and 21) and full disagreement (02 and 20); the percentages for these are 81.3%, 11.6% and 7.1%, respectively (barring small rounding errors). Since cell 11 contains only 0.5% of the data, it follows that the agreement between judges is not due to shared uncertainty: If we subtract this value from the percentage for full agreement, we get 80.8% agreement about the absence or presence of an accent (cells 00 and 22), which appears acceptable. Inspection of the agreement data for the nine different main conditions (SUBJ→SUBJ, SUBJ→DIRECT OBJ, DIRECT OBJ→SUBJ, ...) gives about the same results (we do not represent these data here). This means that our results are not due to unequal distributions of clear and less clear cases of (de-)accentuation in the different main conditions (although there seems to be less agreement between judges about the accentuation of grammatical subjects than for direct objects and pp-objects).

3.2 Main Results

The main results of our experiment are presented in Tables 3–5. The cells in these tables represent mean accentedness scores (cf. Section 3.1), averaging across replications and speakers. An average score of 4.0, for example, would indicate that both judges identified all expressions in a particular cell as accented, an average score of 0.0 would mean that both judges identified all expressions in a cell as unaccented[5]. Table 3 shows accentedness scores for (GIVEN) target expressions as a function of grammatical role and surface position in context vs. target utterance. Table 4 shows scores for non-target nps in target utterances, i.e., expressions referring to items which had **not** been mentioned in the preceding context utterances and are NEW in the current scenario. Note that Table 4 contains data for SUBJ and DIRECT OBJ only, since PP-OBJs in target descriptions were always GIVEN. Table 5 contains data for two control conditions. Target utterances in these conditions also contain only SUBJ and DIRECT OBJ expressions. In the ALL GIVEN control, both expressions were GIVEN and shared grammatical function and surface position with the previous utterance (cf. Example (4)). In the ALL NEW control, both expressions were NEW, i.e., not mentioned in the context utterances of the scenario (cf. Example (3)). Note from Table 5 that in the ALL NEW control, both the subject and the object tend to be accented. This suggests that speakers treated subject and object as separate focal domains, each with its own prosodic head. It therefore seems appropriate to treat an **unaccented** subject in this experiment as a case of true deaccentuation — rather than just as an "unaccented expression"

[5]Since agreement between judges about the presence or absence of an accent was high and did not vary strongly across conditions (see Section 3.1.1.), the average accentedness scores may be converted with some caution to percentages. For example, average accentedness scores of 0.0, 1.0, 2.0, 3.0 and 4.0 would correspond roughly to percentages of 0, 20, 40, 60, 80 and 100 % +ACCENT expressions.

within the (metrically defined) scope of an accent on the DIRECT OBJ (cf. [Gussenhoven, 1983, Dirksen, 1992]).

The important results for our study appear in Table 3. This table shows that average accentedness scores for target expressions vary considerably across context conditions. To evaluate these differences, the data were analysed with an analysis of variance procedure. For each cell in the 3x3 matrix, average accentedness scores were computed for each condition for each speaker, and these data were subjected to an analysis of variance, with grammatical function in the target utterance and grammatical function in the prior context as fixed factors and speakers as the replication factor (an analysis with items as the replication factor was not performed, given the homogeneity of the items). The two main effects and their interaction were all found to be significant: For function in prior context, $F(2,18) = 13.87$, $p=.0002$, for function in target utterance $F(2,18) = 21.97$, $p \leq .00005$, and for the interaction: $F(4,36) = 28.5$, $p \leq .00005$.

The significant interaction seems to have several components.[6] First, we note that the values on the diagonal in Table 3 are lower than the off-diagonal values. This means that expressions for which the grammatical function in the current utterance (hereafter GF-TARGET) and the grammatical function in the context (hereafter GF-CONTEXT) are the same, i.e., where there is persistence of grammatical function, are more likely to remain unaccented than expressions for which there is a change in grammatical function. Thus, persistence of grammatical function appears to be a relevant factor in triggering deaccentuation. It may also be noted that some of the off-diagonal values in Table 3 are as high as those for the ALL NEW conditions in Table 5. This is consistent with the observation that simple GIVENNESS in the sense defined here as "previous mention in the context" does not itself suffice to trigger deaccentuation.

[6]The strong disordinal interaction makes it meaningless to make general statements about the effects of the main factors. Therefore, we will not discuss them further.

Second, the off-diagonal values for DIRECT OBJ→PP-OBJ and PP-OBJ→DIR-OBJ (1.6 and 1.4, respectively) are lower than the other off-diagonal values (3.6, 3.2, 3.3 and 3.0). Note that DIRECT OBJ→PP-OBJ and PP-OBJ→DIR-OBJ have different GF-CONTEXT and GF-TARGET but the same surface position in context and target (hereafter SP-CONTEXT and SP-TARGET, respectively), while the other conditions have both different GF-CONTEXT and GF-TARGET **and** different SP-CONTEXT and SP-TARGET. So, the difference in average accentedness scores must be due to persistence vs. change of surface position. At the same time, the scores for conditions with different GF-TARGET and GF-CONTEXT but same SP-TARGET and SP-CONTEXT (1.6 for DIRECT OBJ→PP-OBJ and 1.4 for PP-OBJ→DIR-OBJ) are higher than those for conditions with both same GF-TARGET and GF-CONTEXT **and** same SP-TARGET and SP-CONTEXT (0.6 for DIRECT OBJ→DIRECT OBJ and 0.7 for PP-OBJ→PP-OBJ.[7] This means that both persistence of grammatical function and persistence of surface position contribute to de-accentuation. Closer inspection of the data suggested that the effects of the two types of persistence were different across speakers. We will pursue the issue of speaker variability in Section 3.3.

Third, the values on the diagonal in Table 3 seem to fall into two classes: SUBJ→SUBJ has a higher accentedness score (2.1) than the other two conditions, DIRECT OBJ→DIRECT OBJ and PP-OBJ→PP-OBJ (0.6 and 0.7, respectively). This is surprising, since it has been argued that subjects are more likely to be deaccented than other expressions [Nooteboom and Terken, 1982]. Our findings here may reflect a tendency towards rhythmic alternation of accents in pre-nuclear position, i.e., before the final accent in a prosodic phrase [Horne, 1991a]. Comparison of the values for SUBJ→SUBJ in Table 3 with those for SUBJ→SUBJ in the ALL GIVEN condition in Table 5 lends support to this interpretation. While the SUBJ→SUBJ condition of Table 3 represents sequences such as (5), the SUBJects in the ALL GIVEN condition of Table 5 represent sequences such as (6).

[7]Post-hoc analyses show that these differences are significant for DIRECT OBJ $F(1,9) = 9.7$, p=.01, but not for PP-OBJ, $F(1,9) = 2.8$, p=.13.

(5) a. The ball touches the CONE.

b. The ball touches the BOX.

(6) a. The ball touches the CONE.

b. The ball COVERS the cone.

For pragmatic reasons, utterance (5b) has an accented object noun, and utterance (6b) has an accented verb. If we represent unaccented content words by *w* for WEAK and accented content words by *s* for STRONG, the metrical structure for (5b) would be *wws*, but for (6b) it would be *wsw*. The rhythmic alternation hypothesis holds that there is a tendency to alternate strong and weak nodes in a prosodic phrase, thus converting *wws* into *sws*; the result of this conversion would be the accenting of *ball* in (5b). But since (6b) already exhibits rhythmic alternation, no conversion would occur, and *ball* would remain weak, and thus, unaccented. The absence of a difference between the accentedness scores for DIRECT OBJ→DIRECT OBJ in Table 3 and the ALL GIVEN condition in Table 5 suggests that rhythmic alternation applies only before the last accent in the phrase.

An alternative explanation for our findings that subjects tend to be accented more than either direct objects or PP objects is that, in pre-nuclear position, the distinction between GIVEN and NEW is not signalled by the absence or presence of an accent but by varying the pitch range, i.e., by creating a difference in the phonetic realization of the accents [Kruyt, 1985, Horne, 1991b]. In order to test this possibility, we computed average pitch maxima for accented GIVEN and NEW grammatical subjects. To avoid contamination of the data by judges' uncertainty, we included only those cases where judges agreed that an expression was accented. To avoid contamination of the data by vowel intrinsic f0, we included only those cases where a set of utterances with the same target expression was available in all 6 relevant conditions: SUBJ→SUBJ, DIRECT OBJ→SUBJ and PP-OBJ→SUBJ for the GIVEN conditions; SUBJECTS following contexts with constant SUBJECTS, OBJECTS and PP-OBJ, respectively, for ALL NEW utterances. In all, there were only nine sets of utterances which had accented target expressions in all 6 relevant conditions, from six different speakers.

Pitch measurements for these utterances were obtained using a cross correlation algorithm [Talkin, 1989]. Pitch maxima were measured at the amplitude peak of the stressed syllable of the subject noun. Average pitch maxima for the different conditions are 123.3 Hz for SUBJ→SUBJ, 119.1 for OBJ→SUBJ, 119.9 for PP-OBJ→SUBJ and 120.8, 120.1 and 121.0 for the subjects in the different ALL NEW conditions. These results do not confirm Kruyt's and Horne's findings; there is very little difference in our data between accented GIVEN and NEW subjects. Instead, these data suggest that speakers made a simple binary decision to accent or to deaccent an expression. Once that decision was made, subjects realized it acoustically in very predictable ways. However, since the number of data points from which the means are computed is very small, this conclusion should be interpreted with some caution.

3.3 Speaker Variability

In Section 3.1.1 we noted that there was full agreement between judges about whether or not an expression was accented in 81% of cases. We also noted that the different context conditions showed either uni-modal or bi-modal frequency distributions of accentedness scores, with observations clustering at the extremes of the scale. Thus, an accentedness score near 2.0 for a particular condition corresponds to a situation in which approximately half of the relevant cases were judged accented and half unaccented. Such bi-modal distributions can be found in several cells of Table 3, including SUBJ→SUBJ, PP-OBJ→DIRECT OBJ and DIRECT OBJ→PP-OBJ. In order to determine whether this bi-modality can be explained as resulting from systematic differences between speakers, we inspected accentedness scores for individual speakers.

For the SUBJ→SUBJ condition, inspection of the data for individual speakers revealed no clear subgroups of speakers employing different strategies. Instead, there was a gradual transition from speakers with a high average accentedness score to those with a low score. For PP-OBJ→DIRECT OBJ and DIRECT OBJ→PP-OBJ conditions, however, there was evidence that groups of speakers pursued several common strategies. In order to trace these strategies, three different numbers were computed for each speaker: (a) an average score for conditions with same SP-TARGET and SP-CONTEXT **and** same GF-TARGET and GF-CONTEXT (DIRECT OBJ→DIRECT OBJ and PP-OBJ→PP-OBJ), (b) an average score for conditions with same SP-TARGET and SP-CONTEXT but **different** GF-TARGET and GF-CONTEXT (PP-OBJ→DIRECT OBJ and DIRECT OBJ→PP-OBJ), and (c) an average score for conditions in which both SP-TARGET and SP-CONTEXT and GF-TARGET and GF-CONTEXT differed (SUBJ→DIRECT OBJ and SUBJ→PP-OBJ). Again, these scores are computed on our five-point accentedness scale, as described in Section 3.1. These scores are shown in Table 6. The data shown in Table 6 suggest that speakers fall into three subgroups with respect to their accenting strategy. For speakers 3 and 4, conditions (b) and (c) have about the same average accentedness score and condition (a) has a lower score. Apparently, for these speakers, "same function" is a **necessary** condition for deaccentuation. For speakers 1, 2, 6, 7 and 10, however, there is a large difference between conditions (a) and (b), on the one hand, and condition (c), on the other. Apparently, for these speakers "same position" is a **sufficient** condition for deaccentuation. Speakers 5, 8 and 9 appear to differentiate between all three conditions. Apparently, for these speakers "same position" and "same function" are more or less independent conditions; the more the conditions influencing deaccentuation are fulfilled, the more deaccentuation occurs.

4 Discussion

Our findings show that simple GIVENNESS, as defined as mere mention of an item in the immediate context, is not a sufficient condition for deaccentuation. Instead, additional factors, the persistence of surface position and the persistence of grammatical function from the context to the current utterance, play an important role. If an expression has the **same** grammatical role as its antecedent expression in the immediate context ("persistence of grammatical function"), it is likely to be deaccented. If an expression has a **different** grammatical function than the antecedent expression but occupies the **same** surface position ("persistence of surface position"), it is also likely to be deaccented, but less likely than in the former condition. If there is a change in both grammatical function and surface position from one utterance to the next, a GIVEN expression is in fact likely to be accented — as likely in fact as a NEW expression. It should be kept in mind, though, that surface position and grammatical function were not fully orthogonal. In particular, grammatical subjects occurred always in initial position.

We also found that the effects of persistence of grammatical function and surface position were not constant across speakers. For some, persistence of grammatical role seemed to be a necessary condition for deaccentuation; for others, persistence of surface position seemed to be a sufficient condition; for still other speakers persistence of position and of grammatical role appeared to make independent and additive contributions.

The data we examined on the accentuation of grammatical SUBJECTS lend some support to the hypothesis that the occurrence of deaccentuation may be affected by the rhythmical properties of the utterance. Our speakers' accentuation of these expressions might in fact be explained as the application of some form of rhythmic alternation of accented and unaccented words before the final accent in a prosodic domain.

What answers do these results provide to the question of why expressions referring to GIVEN information are accented in some cases and unaccented in others, assuming that we can generalize from our elicited speech to spontaneous human discourse? Our results imply that expressions referring to GIVEN information are most likely DEACCENTED if they share grammatical and surface positional properties with their antecedents. If we assume that the distribution of unaccented expressions can be used to draw inferences about the function of deaccentuation to the listener, our findings suggest that grammatical function and surface position may serve as cues to the mapping of referring expressions onto entities in the mental representation of the preceding discourse. The grammatical and positional characteristics of unaccented expressions may be used as cues to access candidate antecedents in the discourse model. We may speculate that such antecedents are made accessible from the current expression by virtue of this property-sharing. When an expression does **not** share these properties with its intended antecedent — but with a non-antecedent expression or expressions, the intended antecedent is in fact less readily accessible, due to interference from these competing expressions. In such cases this distinction in accessibility may be signalled by the accenting of the target expression. In this way, the number of candidate antecedents may be pruned for further processing.

This interpretation suggests that accessibility is indeed what matters to deaccentuation, but not just the relative accessibility of entities in the discourse model, without taking into consideration the properties of the referring expression itself. Rather, accessibility is to be defined in terms of how easily an entity in the discourse model can be accessed from the referring expression, on the basis of the properties both of the referring expression and the entity in the discourse model.

Clearly, one would not want to argue at this stage of research that syntactic and surface position information represent the **only** relevant sources of information in the mapping processing. Other sources of information such as lexical and morphological information, are obviously important for antecedent identification, not only in cases of lexical repetition and morphological relationship, but also in matters of synonymy and gender/number information. The relative contribution of each of these factors remains to be explored.

Our findings are relevant to current cognitive aproaches to the mapping process. For instance, centering theory [Grosz, 1981, Grosz et al., 1983, Gordon et al., 1993], defines for each non-initial utterance $\mathbf{utt}_i$ in a discourse segment a BACKWARD-LOOKING CENTER ($\mathbf{Cb}_i$) which serves to link $\mathbf{utt}_i$ to the previous discourse, and a set of FORWARD-LOOKING CENTERS ($\mathbf{Cf}_i$), including $\mathbf{Cb}_i$, which provide potential links between $\mathbf{utt}_i$ and subsequent utterances in the discourse. Elements of $\mathbf{Cf}_i$ are ranked according to their prominence in the speaker's and listener's discourse model on the basis of properties such as grammatical role, surface order, surface realization of the reference (e.g., pronoun vs. full NP), and the accentuation of the expression. In this way, centering theory captures the notion that items in the discourse model may have different degrees of accessibility. It is assumed that speakers tend to realize the most prominent item in $\mathbf{Cf}_i$ as the $\mathbf{Cb}_{i+1}$, the backward looking center of the next utterance, if this item is in fact referred to in $\mathbf{utt}_{i+1}$; otherwise, the next most prominent element of $\mathbf{Cf}_i$ will be $\mathbf{Cb}_{i+1}$, and so on.

Centering theory has been applied primarily to the solution of pronominal expressions, by formulating constraints on the occurrence of pronouns. The solution of pronouns is governed by listeners' knowledge of such constraints [Gordon et al., 1993]. It may be assumed that the ranking of elements of $\mathbf{Cf}_i$ according to their prominence in the listener's discourse model provides an ordering for evaluating candidate interpretations of pronouns in $\mathbf{utt}_{i+1}$. Thus, centering theory assumes that surface position and grammatical function of expressions in the immediate context contribute to the accessibility of discourse antecedents in the discourse model, but are not directly involved in the mapping process. It has been suggested, however, that such formal properties as grammatical function may also have a direct effect on the interpretation of pronouns, by means of a "property-sharing" constraint [Kameyama, 1986].

The assumption that accessibility is a function both of the properties of entities in the discourse model and those of the current expression is also made elsewhere [McKoon et al., 1993]. The present findings provide further support for this assumption, for the following reason. The experimental scenarios in our study were constructed so that the context-setting utterances preceding the target would favor a single entity **X** as the most prominent item in the set $\mathbf{Cf}_i$ at the onset of the target utterance$_{i+1}$. In the absence of factors enhancing the prominence of other entities in the discourse model, one might expect that, at the onset of the target utterance, **X** should have higher prominence than any competing entity. If the prominence of **X** in the set $\mathbf{Cf}_i$ were the only property of **X** relevant to the mapping process, one would expect deaccentuation to occur in **all** cases in which the target utterance$_{i+1}$ in fact contained an expression $\mathbf{x}'$ referring to **X**. However, in our data this clearly did not occur. Instead, deaccentuation occurred most frequently when $\mathbf{x}'$ in utterance$_{i+1}$ shared grammatical function and/or surface position with **x** in the immediately preceding utterance$_i$. Unless we assume that these properties are represented explicitly in the representation of elements of the **Cf**, it is difficult to explain this difference.

A possible limitation to this conclusion might be that the strong effects of persistence of grammatical function and surface position may arise only due to the syntactic parallelism of successive utterances in our context and target utterances. In situations where there is more syntactic variation between consecutive utterances, these factors might in fact constitute less useful cues to the process of establishing a link between a referring expression and its antecedent, compared to other sources of information. For instance, in less constrained situations speakers employ other kinds of devices such as topicalization to maintain coherence, which may result in syntactic variation. The possibility that persistence of grammatical function and/or surface position is less relevant to deaccentuation in such situations should be investigated in future research.

It might also be suggested that our speakers produced a high percentage of accented expressions for GIVEN items because our experiment inadvertently fostered a CONTRASTIVE situation, in which speakers were always contrasting one geometric shape with other members of a small set of such shapes and thus may have accented their expressions to indicate CONTRAST. The possibility of reinterpreting our findings in this light does NOT seem useful in explaining the differences in accentuation behavior we observed, since the possibilities for using contrastive stress were available for all speakers in all conditions of the experiment. Alternatively, one might suggest that the CONTRAST was not between a specific object and all other elements of the set of objects used in the experiment, but between a specific object and the one that was mentioned in the previous utterance with the same grammatical role and/or in the same position. However, this interpretation is close to our own interpretation, as it appears to imply the same assumptions about the mechanisms involved in mapping a referring expression onto possible antecedents in the discourse model.

5 Conclusion

In the experiment described in this paper, we investigated the conditions under which speakers use unaccented expressions to refer to GIVEN entities. We started from the observation that GIVENNESS by itself is not a sufficient condition for deaccentuation. We found that persistence of grammatical function and persistence of surface position are additional constraints on deaccentuation. We hypothesize that speakers use deaccentuation when they believe that an intended interpretation is easily accessible in the discourse model. Therefore, we propose that the properties of grammatical function and surface position may be used by the listener as cues to access potential antecedents in the discourse model, by leading him or her to look for a candidate antecedent with the same properties. However, we leave open the possibility that the cue validity of grammatical function and surface position in relation to other sources of information may vary with changes in the characteristics of the communicative situation. This hypothesis remains to be tested in further investigations.

Acknowledgments

The research was conducted while the first author was visiting AT&T Bell Laboratories, Murray Hill NJ, during the spring of 1992. The visit was made possible by a grant from the Netherlands Organization for Scientific Research (NWO), Grant R 30-351, and additional funding from AT&T Bell Labs and the Institute for Perception Research. The authors wish to thank Jan van Santen, Ananth Sankar and René Collier for useful discussions.

References

[Asher and Wada, 1988] Asher, N. and Wada, H. (1988). A computational account of syntactic, semantic and discourse principles for anaphora resolution. *Journal of Semantics*, 6:309–344.

[Baart, 1987] Baart, J. L. G. (1987). *Focus, Syntax and Accent Placement*. PhD thesis, University of Leyden, Leyden.

[Bock, 1982] Bock, J. K. (1982). Toward a cognitive psychology of syntax: Information processing contributions to sentence formulation. *Psychological Review*, 89:1–45.

[Bock and Irwin, 1980] Bock, J. K. and Irwin, D. E. (1980). Syntactic effects of information availability in sentence production. *Journal of Verbal Learning and Verbal Behavior*, 19:467–484.

[Bock and Warren, 1985] Bock, J. K. and Warren, R. K. (1985). Conceptual accessibility and syntactic structure in sentence formulation. *Cognition*, 21:47–67.

[Brennan et al., 1987] Brennan, S. E., Friedman, M. W., and Pollard, C. J. (1987). A centering approach to pronouns. In *Proceedings of the 25th Annual Meeting*, pages 155–162, Stanford. Association for Computational Linguistics.

[Brown, 1983] Brown, G. (1983). Prosodic structure and the given/new distinction. In Ladd, D. R. and Cutler, A., editors, *Prosody: Models and Measurements*, pages 67–78. Springer Verlag, Berlin.

[Chafe, 1976] Chafe, W. (1976). Givenness, contrastiveness, definiteness, subjects, topics, and point of view. In Li, C., editor, *Subject and Topic*, pages 25–55. Academic Press, New York.

[Chafe, 1974] Chafe, W. L. (1974). Language and consciousness. *Language*, 50:111–133.

[Clark and Clark, 1977] Clark, H. H. and Clark, E. V. (1977). *Psychology and Language*. Harcourt, Brace, Jovanovich, Inc.

[Dirksen, 1992] Dirksen, A. (1992). Accenting and deaccenting: A declarative approach. In *Proceedings of COLING-92*, pages 865–869.

[Fry, 1955] Fry, D. B. (1955). Duration and intensity as physical correlates of linguistic stress. *Journal of the Acoustical Society of America*, 27:765–769.

[Fry, 1958] Fry, D. B. (1958). Experiments in the perception of stress. *Language and Speech*, 1:126–152.

[Gordon et al., 1993] Gordon, P. C., Grosz, B. J., and Gillion, L. A. (1993). Pronouns, names, and the centering of attention in discourse. *Cognitive Science*. To appear.

[Grosz et al., 1983] Grosz, B., Joshi, A., and Weinstein, S. (1983). Providing a unified account of definite noun phrases in discourse. In *Proceedings of the 21st Annual Meeting*, pages 44–50, Cambridge MA. Association for Computational Linguistics.

[Grosz, 1977] Grosz, B. J. (1977). The representation and use of focus in dialogue understanding. Technical Report 151, SRI International, Menlo Park Ca. University of California at Berkeley PhD Thesis.

[Grosz, 1981] Grosz, B. J. (1981). Focusing and description in natural language dialogues. In Joshi, A., Webber, B., and Sag, I., editors, *Elements of Discourse Understanding*. Cambridge University Press, Cambridge, England.

[Grosz and Sidner, 1986] Grosz, B. J. and Sidner, C. L. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204.

[Gundel et al., 1993] Gundel, J. K., Hedberg, N., and Zacharski, R. (1993). Cognitive status and the form of referring expressions in discourse. *Language*, 69:274–307.

[Gussenhoven, 1983] Gussenhoven, C. (1983). *On the Grammar and Semantics of Sentence Accents*. Foris Publications, Dordrecht, Neth.

[Hajicova et al., 1990] Hajicova, E., Kubon, P., and Kubon, V. (1990). Hierarchy of salience and discourse analysis and production. In *Papers Presented to the 13th International Conference on Computational Linguistics*, pages 144–148, Helsinki. International Conference on Computational Linguistics.

[Halliday, 1967] Halliday, M. A. K. (1967). Notes on transitivity and theme in English, part 2. *Journal of Linguistics*, 3:199–244.

[Haviland and Clark, 1974] Haviland, S. E. and Clark, H. H. (1974). What's new? acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior*, 13:512–521.

[Horne, 1991a] Horne, M. (1991a). Accentual patterning in 'new' vs 'given' subjects in English. Working Papers 36, Department of Linguistics, Lund University, Lund.

[Horne, 1991b] Horne, M. (1991b). Phonetic correlates of the new/given parameter. In *Proceedings of the XIIth International Phonetic Congress*, pages 230–233, Aix-en-Provence. International Phonetic Congress.

[Kameyama, 1986] Kameyama, M. (1986). A property-sharing constraint in centering. In *Proceedings of the 24th Annual Meeting of the American Association of Computational LInguistics*, pages 200–206, Cambridge, Mass. American Assocaition of Computational Linguistics.

[Kruyt, 1985] Kruyt, J. G. (1985). *Accents from Speakers to Listeners: An Experimental Study of the Production and Perception of Accent Patterns in Dutch*. PhD thesis, University of Leyden.

[Ladd, 1980] Ladd, D. R. (1980). *The Structure of Intonational Meaning*. Indiana University Press, Bloomington, Ind.

[Lehiste, 1970] Lehiste, I. (1970). *Suprasegmentals*. The MIT Press.

[McKoon et al., 1993] McKoon, G., Ward, G., Ratcliff, R., and Sproat, R. (1993). Morphosyntactic and pragmatic factors affecting the accessibility of discourse entities. *Journal of Memory and Language*, 32:56–75.

[Nooteboom and Terken, 1982] Nooteboom, S. G. and Terken, J. (1982). What makes speakers omit pitch accents?: An experiment. *Phonetica*, 39:317–336.

[Pierrehumbert, 1980] Pierrehumbert, J. B. (1980). *The Phonology and Phonetics of English Intonation*. PhD thesis, Massachusetts Institute of Technology. Distributed by the Indiana University Linguistics Club.

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

[Prince, 1992] Prince, E. F. (1992). The zpg letter: Subjects, definiteness, and information-status. In Thompson, S. and Mann, W., editors, *Discourse Description: Diverse Analyses of a Fund Raising Text*, pages 295–325. John Benjamins B. V., Philadelphia.

[Sidner, 1979] Sidner, C. (1979). *Towards a Computational Theory of Definite Anaphora Comprehension in English Discourse*. PhD thesis, The Massachusetts Institute of Technology.

[Sidner, 1983] Sidner, C. L. (1983). Focusing in the comprehension of definite anaphora. In Brady, M. and Berwick, R., editors, *Computational Models of Discourse*, pages 267–330. MIT Press, Cambridge MA.

[Swerts and Collier, 1992] Swerts, M. and Collier, R. (1992). On the controlled elicitation of spontaneous speech. *Speech Communication*, pages 463–468.

['t Hart et al., 1990] 't Hart, J., Collier, R., and Cohen, A. (1990). *A Perceptual Study of Intonation.* Cambridge University Press, Cambridge UK.

[Talkin, 1989] Talkin, D. (1989). Looking at speech. *Speech Technology*, 4:74–77.

[Terken and Nooteboom, 1987] Terken, J. and Nooteboom, S. G. (1987). Opposite effects of accentuation and deaccentuation on verification latencies for given and new information. *Language and Cognitive Processes*, 2(3/4):145–163.

[Terken, 1985] Terken, J. M. B. (1985). *Use and Function of Accentuation: Some Experiments.* PhD thesis, University of Leiden, Helmond, Neth.

	EVENT		DESCRIPTION
		Context a	
I	Initial Display		
CEa1	ball touches cone	CDa1	the ball touches the cone
CEa2	ball touches cross	CDa2	the ball touches the cross
CEa3	ball touches diamond	CDa3	the ball touches the diamond
		Context b	
I	Initial Display		
CEb1	cone touches ball	CDb1	the cone touches the ball
CEb2	cross touches ball	CDb2	the cross touches the ball
CEb3	diamond touches ball	CDb3	the diamond touches the ball
		Context c	
I	Initial Display		
CEc1	cone pushes rectangle against ball	CDc1	the cone pushes the rectangle against the ball
CEc2	cross pushes line against ball	CDc2	the cross pushes the line against the ball
CEc3	diamond pushes triangle against ball	CDc3	the diamond pushes the triangle against the ball
		Alternative Target Types	
TEa	ball touches star	TDa	the ball touches the star
TEb	star touches ball	TDb	the star touches the ball
TEc	box pushes star against ball	TDc	the box pushes the star against the ball

Table 1: Sample Context Event sequences (CE) and Corresponding Context Descriptions (CD), and Target Events (TE) and Target Descriptions (TD). Any target event may follow any sequence of context events, giving nine different scenarios.

		Judge B			
		0	1	2	Sum
Judge A	0	392 (23.8%)	16 (1.0%)	22 (1.3%)	430 (26.1%)
	1	95 (5.8%)	9 (0.5%)	63 (3.8%)	167 (10.2%)
	2	94 (5.7%)	17 (1.0%)	937 (57.0%)	1048 (63.7%)
	Sum	581 (35.3%)	42 (2.6%)	1022 (62.1%)	1645 (100.0%)

Table 2: Agreement between Judges on Absence (0), Presence (2) and Inability to Determine (1) Accent for Referring Expressions in Target Utterances

		Function in Target Utterance		
		SUBJ	DIRECT OBJ	PP-OBJ
Function	SUBJ	2.1	3.6	3.2
in	DIRECT OBJ	3.3	0.6	1.6
Context	PP-OBJ	3.0	1.4	0.7

Table 3: Average Accentedness Scores of GIVEN Target Expressions as a Function of Grammatical Role in Target and Context. Each cell represents 50 observations.

Function in Target Utterance	
SUBJ	DIRECT OBJ
3.7	3.8

Table 4: Average Accentedness Scores of NEW Non-Target Expressions in Target Utterances as a Function of Grammatical Role. Each cell represents 150 observations.

	Function in Current Utterance	
	SUBJ	DIRECT OBJ
ALL GIVEN	0.6	0.2
ALL NEW	3.3	3.6

Table 5: Average Accentedness Scores of SUBJ and DIRECT OBJ in Control Conditions. Each cell represents 150 observations.

Table 6: Average Accentedness of Target Expressions for Individual Speakers

		Condition		
		(a)	(b)	(c)
		same position same function	same position diff function	diff position diff function
Speaker	1	0.3	0.4	3.0
	2	1.1	1.2	3.4
	3	0.0	3.2	3.6
	4	1.8	2.6	2.8
	5	1.1	1.8	2.7
	6	1.0	1.5	3.9
	7	0.0	0.4	3.6
	8	0.5	2.0	3.7
	9	0.0	1.0	2.8
	10	0.5	1.0	4.0