

Semantically Enhanced Pronouns

Donna K. Byron

Department of Computer Science
University of Rochester
Rochester NY 14627, U.S.A
dbyron@cs.rochester.edu

Abstract

This paper revisits the relationship between the form of referring expressions and the assumed salience of their referents. Previous work suggests that pronouns must refer to highly salient discourse entities, while descriptive noun phrases should be used to refer to less salient items. Our claim is that the referring expression itself is only part of this relationship, since it is not the only constituent that contributes semantic constraints that bear on its referent. Semantic constraints can be contributed by additional constituents, and such Semantically Enhanced pronouns may coherently refer to less salient entities, even entities not previously mentioned in the discourse.

1. Introduction

For the past twenty years or so, a large body of work has developed to explore the relationship between the form a referring expression (RE) may take and the salience of its referent within the current discourse state. Entities are initially brought into a discourse with a descriptive expression, then re-mentioned via reduced forms such as pronouns. In order for a referring expression to be interpreted correctly by the addressee, the speaker must accurately judge the referent's status within the addressee's attentional state. For reduced referring expression forms, the speaker must accurately judge which entities are the current focus of attention.

Because of this relationship between reduced forms and high salience, a common belief among anaphora researchers, voiced below by Gundel et al. (1993, pg. 279), is that the referent of a pronominal referring expression must be in focus:

“[Focus] status is necessary for appropriate use of zero and unstressed pronominals.”

Although referring to highly salient entities is certainly the default behavior of unstressed pronouns, the above principle fails to explain a pronoun like the one in the final sentence of Example 1 (from the Penn Treebank (Marcus et al., 1993), pronoun of interest in boldface)¹:

Example (1)

Propaganda is just information to support a viewpoint, and the beauty of a democracy is that it enables you to hear or read every viewpoint and then make up your own mind on an issue. The restrictions on viewing and dissemination of [Voice of America] material were especially absurd: An agency in the information business was not being allowed to inform. In June 1988, I

wrote in this space about this issue. Assuming **it** wasn't one of those columns that you clipped and put on the refrigerator door, I'll review the facts.

In Example 1, the highlighted pronoun refers to the column about the Voice of America. The column was only implicitly present in the previous discourse, triggered by the verb phrase “I wrote” in the previous sentence, and it is definitely not the focused entity at the beginning of the final sentence.² Nonetheless, the sentence is easily interpretable as coherent English. The Treebank corpus contains many additional examples of unstressed pronouns referring to non-focused entities.

If unstressed pronouns are not limited to referring to focused entities, we are immediately led to wonder how salient the referent of a pronoun does need to be. In this paper, we explore that question within the context of Gundel et al.'s Givenness Hierarchy, which provides a continuum of possible salience statuses for the referents of referring expressions. After reviewing some examples that demonstrate that the referents of pronouns have quite a bit more mobility than previously believed, we propose a reformulation of the necessary conditions on referents of different cognitive status, which takes into account additional linguistic material rather than just the referring expression. We claim that the required salience of the referent of a referring expression is determined not just by the form of the referring expression, but by the total of all the semantic information about the referent contributed by the sentence.

The remainder of this paper proceeds as follows: in Section 2, we review theories of salience and discuss how referring expression forms and the salience of the referent are believed to relate. The Givenness Hierarchy and Centering Theory are described as examples of well-

¹ The reader is asked to please forgive the length of the examples used in this paper. When the example pronouns were previously shown to colleagues, some were uncomfortable with my suggested interpretation unless more context was provided.

² Interpretations of example sentences in this paper were verified by interviewing a number of English native speaker colleagues. For this example, I showed them the paragraph without the final sentence and asked what a sentence-subject ‘it’ in the final sentence would refer to, thus forcing them to rely on salience alone to compute an interpretation. My subjects gave a variety of responses, none of which was the column.

Status	In focus >	Activated >	Familiar >	Uniquely Identifiable >	Referential >	Type identifiable
Form	it	that this this N	that N	the N	this N (indef)	a N

Table 1: The Givenness Hierarchy

substantiated theories that explain this relationship. Next, Section 3 explores some heretofore problematic example pronouns that the original formulation of the Givenness Hierarchy could not explain. We use these examples to demonstrate that semantically enhanced pronouns may refer to the same range of referents as definite noun phrases. In Section 4 we lay out explicitly our reformulation of the Givenness Hierarchy, and discuss its implications for pronoun resolution.

2. Related Background

Each stretch of discourse triggers a set of *discourse entities* (Karttunen, 1976), alternately called *discourse referents* (Sidner, 1983), into the attentional states of the discourse participants. We will call this set the Discourse Context (DC). It is an open research question to operationalize exactly what entities should be included in the DC and when they should be added and deleted.

At any particular point in an ongoing discourse, all entities to which a speaker might want to refer are not equally salient for the addressee. Some are still totally unknown to the addressee, others are unrelated to the current discourse, some are already in the DC, and a small set of those already in the DC are highly salient by virtue of being the current topic of discussion. Making a prediction of the referent's current salience for the addressee, speakers follow Grice's maxim of quantity (Grice, 1975) to construct a felicitous reference - being as informative as they feel necessary to point the addressee to the correct referent. "Necessary" here involves judging how salient that entity is for the addressee. Being more informative than is required, for example, referring by name to a person who is currently the focus of a dialog, is confusing for the addressee. Being less informative than is required, for example using a reduced form to refer to a person you mistakenly thought the addressee was just looking at, results in a reference failure.

2.1. Use of Different Referring Expression Forms

The Givenness Hierarchy shown in Table 1 is a continuum developed by Gundel et al. (1993, 1998) to explain the relationship between different referring expression forms and the cognitive status of the referent within the addressee's attentional state (or more accurately, the speaker's assumption of the addressee's attentional state). According to this theory, each cognitive status is a necessary and sufficient condition for using the referring expression form indicated. Given a referring expression of a specific form to interpret, an addressee utilizes this pragmatic clue as a signal of the likely cognitive status of the referent.

In this continuum, unstressed pronouns and zero pronouns may be used to refer to entities with *focused*

status, defined as follows: "The entities in focus at a given point in the discourse will be that partially-ordered subset of activated entities which are likely to be continued as topics of subsequent utterances".[pg. 279] The preference of unstressed pronouns for highly salient entities has also been confirmed by many computational linguistics studies, which have found that utilizing salience in determining the most likely referent outperforms purely syntactic approaches (Brennan et al., 1987; Strube, 1998).

Demonstrative pronouns may be used to refer to entities with *activated* status, which is described as "represented in current short-term memory". Activated entities have been evoked into short-term memory by some trigger in either the linguistic context or in the extra-linguistic context (all entities in the DC, using our terminology from above, are *activated*). Items with *familiar* status are known to the addressee, either in long-term or short-term memory, and an RE with a demonstrative determiner can refer to them. The *uniquely identifiable* category includes items for which the addressee can build a referent solely with the content of the referring expression. Descriptive noun phrases using the definite determiner fall into this category. The remaining categories of reference rely only on the addressee's prior knowledge of types, and the referential and *type* statuses differ by whether the addressee interprets the RE as referring to an individual or simply to the type. Gundel et al. report compelling cross-linguistic evidence to support this theory.

Because the statuses are related in a continuum, referring expression forms may be used to refer to items that are more salient than the form requires. For example, a definite noun phrase may be used to refer to entities that are in focus, activated, or familiar in addition to uniquely identifiable entities. The different statuses are not discrete. Focused entities are a subset of activated entities, which are a subset of familiar entities, etc.

Several other studies have explored the tendency of personal pronouns to refer to focused items and for demonstrative pronouns to refer to activated items (c.f. Schuster, 1988; Webber, 1988; Schiffman, 1985). Speakers utilize the pragmatic contrast between demonstrative and definite pronouns as a clue to the salience of the referent. Alternating the pronoun's form creates a different default interpretation of its referent. An example (from Borthen, 1997) is:

Example (2)³

- A) There was a snake_i on my desk. **It**_i scared me.
 B) {There was a snake on my desk}_j. **That**_j scared me.

³ Braces are used when the referent is an entity triggered by a constituent other than a base noun phrase. It is not meant to indicate that the pronoun and the trigger constituent are coreferent.

In a similar study, Schuster (1988) found that alternating the pronoun in Example 3 changed the referent for her test subjects:

Example (3)

A) John thought about {becoming a street Person}_i. It_i would hurt his mother and **it**_i would make his father furious.

B) John thought about {becoming a street person}_j. It_j would {hurt his mother}_j and **that**_j would make his father furious.

In Example 3, the first *it* in the second sentence thrusts the referent described by “becoming a street person” into focus. Use of the personal pronoun further down in the sentence maintains the reference, while switching to a demonstrative pronoun forces the interpretation to an unfocused entity. These examples demonstrate the preference of unstressed pronouns to refer to focused entities, and demonstrative pronouns to refer to unfocused but activated entities.

The Centering Framework (Grosz et al, 1995) is another well-substantiated theory on the relationship between the form of referring expressions and the salience of their referents. In the centering model, discourse entities that are either *directly* realized or *indirectly realized* in a sentence are ranked in descending order of salience, according to ranking factors that are the subject of ongoing research. This ranking represents a prediction of the entity’s centralness for the subsequent unit of discourse. Gundel (1998) relates the *directly realized* status to the *focus* status in GH, and *indirectly realized* with *activated*. In centering terminology, the *Cb* is the entity considered to be most central to the current sequence of sentences, and is defined as the most highly ranked element of sentence N-1 that is realized in sentence N. Centering’s main claim (Centering Rule1) is that if any entity is pronominalized, the *Cb* should be. The second claim of centering is that a sequence of sentences in which the *Cb* changes from utterance to utterance is felt to be less coherent than sequences which maintain an entity as *Cb*.

Centering provides a looser constraint on the cognitive status of the central entity than does the Givenness Hierarchy, because centering does not predict how high up in the salience list the *Cb* must be. It is simply the highest ranked entity that appears. An activated entity that is pronominalized does not violate Centering Rule1 as long as no other more-salient entity is realized in the same sentence. Many psycholinguistic studies have confirmed the Centering Framework’s claims (cf. Gordon et al. 1993).

2.2. What entities are activated?

There are few comprehensive theories that provide a definition of which entities are *activated* by a stretch of discourse. The set of activated entities should include all referents triggered by the discourse that could receive subsequent re-mention. Any object, situation, fact, event, etc. explicitly mentioned in the discourse activates its referent unless it occurs in one of a handful of contexts that preclude subsequent mention (Karttunen, 1976). In

addition to mentioned entities, syntactically-recoverable omitted material, such as ellipsed required objects, triggers entities that are available for subsequent reference (Hobbs, 1986).

Besides mentioned entities, a variety of higher-order entities are also activated into the DC. Asher’s comprehensive study (Asher, 1993) defines specific constructions which trigger many higher order entities such as events, facts, propositions, and situations. Segments of the discourse itself and properties of segments, such as truth conditions and speech acts, are also activated in a manner dependent on the discourse structure (Webber, 1988; Stone, 1994).

The discourse participants themselves, the discourse setting, and any perceptually salient entities, such as an event the participants are watching, are activated and available for subsequent mention. *Inferrables* (Prince, 1992) are entities that are closely related to mentioned entities and that may be referred to anaphorically. Examples are subparts such as “the house ... the roof” and generic types such as in this example (from Sidner, 1983):

Example (4)

My neighbor has a monster Harley 1200. **They** are really huge but gas-efficient bikes.

“From any discourse entity, the speaker can presume that a listener is capable of deriving a discourse entity corresponding to one of a limited number of generic sets to which the referent of the original discourse entity belongs.” (Webber, 1983, pg. 398). Given the plural pronoun in Example 4, the addressee must decide which generic set containing the Harley 1200 is referred to. Reference to related entities is sometimes called *bridging descriptions* (Clark, 1977), and as the name implies, bridging references have been often studied as a phenomenon of definite descriptions (cf. Grosz, 1981; Poesio et al., 1998). The use of unstressed pronouns to accomplish bridging references is relatively unexplored.

Because there is no model that completely and precisely defines the criteria for activated entities, it is sometimes difficult to judge whether a reference should be considered anaphoric or whether the referent is new to the discourse. In Example 5, the atmosphere described in the final sentence is a result of the situation under discussion, which has been mentioned several times. It is unclear whether we should judge this pronoun’s referent to be activated because of its relationship to previously mentioned entities, or if it should be considered new to the discourse:

Example (5)

Above all, Mr. Oxnard noted, the situation is extremely confused. “Professional sugar people here who have strong contacts with the Brazilian sugar industry have been unable to confirm the reports or get enough information to clarify the situation,” he said. “**It**’s the type of nervous atmosphere in which a report can be put out, such as the one saying exports will be suspended, and no one can confirm it.”

As this section demonstrates, a variety of different entities may be activated by a particular bit of discourse, and they carry varying degrees of salience for subsequent

reference. Judging the relative salience of activated entities is a very active area of research.

2.3. What entities are in focus?

Determining which of the activated entities is in focus at any particular point is currently an open research problem. Many factors affect the focused status of a particular discourse referent, for example sentence structure (cf. Brennan, 1987; Baldwin, 1997), thematic roles (cf. Sidner, 1983), prosodic stress (cf. Kameyama, 1998), repeated mention (cf. Mitkov, 1998), and information status (Strube, 1998) have all been acknowledged as factors that affect the salience of discourse referents. Although the perfect model of local focus has not yet been developed, all current models concur that explicitly mentioned entities have higher salience than higher-order entities. Typically, the grammatical subject or object of the sentence is considered to be most salient. We will not commit to any particular definition of focus in this paper.

2.4. Semantic constraints on pronoun referents

While most researchers working on pronoun resolution algorithms have seized on the claim that pronouns must refer to the most salient entity, Kameyama (1986) reminds us that salience ordering in centering is

“...in order to identify the *grammatically-based default order that gives rise to preferred interpretations in neutral contexts*. Note that this default order alone does not *determine* interpretations of pronominal elements. Rather, its role in the centering framework is to give an ordered list of referents (centers) so that commonsense inferences can be controlled.”(pg. 201).

Although it is obvious that semantic and pragmatic constraints should be applied in determining the referent of a pronoun, many authors only seem to consider semantic evidence as a way to judge between highly salient entities. This is understandable given the amount of literature that exists that proves that pronouns tend to refer to focused items. An example is Carbonell and Brown (1988), which discusses the semantics that must be applied to correctly resolve the referents of the pronouns in Example 6, in which both the cake and the table have roughly equal salience:

Example (6)

- A) John took the cake_i from the table and ate it_i.
- B) John took the cake from the table_i and washed it_i.

Using semantics to search for the referent of a pronominal RE outside the set of most salient entities is a neglected area of research. While previous authors argue that focus is primary in disambiguating pronouns, we argue that semantics is primary. In the next section, we discuss some problematic examples from previous literature that demonstrate the problems with an over-reliance on focus, and that motivate our enhancement of

the Givenness Hierarchy to incorporate semantic information that comes from outside the referring expression.

3. Some Problematic Examples

Given the above discussion of previous work, we should expect that unstressed pronouns always refer to highly salient entities. This presumption is so strong in the research community that when faced with a naturally occurring unstressed pronoun that appears to refer to an unfocused entity, many authors will attempt to revise their definition of focus to accommodate the example. This section revisits some of those examples in an attempt to convince the reader that focus does not tell the whole story in these cases. A combination of salience and semantics is needed to explain how these pronouns work.

3.1. Pronominal reference to activated entities

Our first example comes from Borthen et al. (1997), which proposes that higher-order entities, such as sentences and events, are sometimes focused.⁴ Consider this example (pg. 90):

Example (7)

- A) Gro_i asked Anne_j if she_j was aware that {the Center Party's popularity was decreasing}_k.
- B) {She_i didn't get any answer}_i.
- C) It_k came as a surprise to her_j.
- C') That_i came as a surprise to her_j.

Subscripts in the example indicate the preferred interpretation of the pronouns for their test subjects. Because the unstressed pronoun *it* in sentence C is interpreted as referring to proposition *k* expressed by “the Center Party's popularity was decreasing”, the researchers interpret this result to indicate that proposition *k* is in focus at the conclusion of the B sentence. Altering the sentence to include a demonstrative pronoun causes the reader to prefer a different referent for the pronoun. The authors assert this alternation as further evidence that proposition *k* is in focus since it is unselected by using the demonstrative pronoun. However, if use of an unstressed pronoun is proof of a referent's focused status, the referents of the pronoun in all the following possible C sentences are also in focus at the same time:

⁴ We do not dispute this claim in general, just the authors' explanation of this particular example. In the same paper, the authors provide an example of a proposition becoming the most focused entity as a result of being included in a question:

Speaker A) “How do you feel about the fact that Computer Scientists make more money than Linguists?”
Speaker B) “It's terrible!”

We agree that in this example, the proposition is in focus at the end of A's question.

Possible C sentences for Example 7

1. It was the tenth annoying question that Gro had asked that day. (it = Gro's question)
2. It was the beginning of the end of their friendship. (it = episode in which this scene is unfolding)
3. It revealed her profound annoyance with the whole political process. (it = her silence)
4. **It** was Anne's way of showing that she didn't care. (it = the event of Anne being silent)
5. **It** was just at the moment when Bob walked up. (it = the moment when Anne was going to speak)
6. **It** was not a conversation that Anne felt like getting into at the moment. (it = a discussion of the Center party's decline)

These are just a few of the many possible continuations. If we define the focus to be a small set of items, it seems impossible that all these referents would be in focus. However, if we accept the proposition that subsequent reference via an unstressed pronoun proves they are all in focus, then the focus is actually a large set and the distinction between the small set of focused entities and the other activated entities becomes unclear. From this evidence, we conclude that unstressed pronouns can refer to activated entities that are not focused. The reader will undoubtedly have noticed that the pronouns in these possible C sentences appear in highly constrained constructions where the referent of the pronoun is made clear from the predicate complement information. Some semantic information is required to inform the addressee not to resolve the pronoun to the default referent. We call pronouns in such semantically-enriched contexts **Semantically Enhanced Pronouns**. They violate the traditional expectation that "all pronouns require the referent to be at least activated, which is no doubt related to the fact that the minimal descriptive content of a pronoun provides little if any basis for identifying the referent." (Gundel et al, 1993, pg. 285) Semantically enhanced pronouns can have as much descriptive content for identifying the referent as full noun phrases.

Our alternative explanation of Example 7 is that the verb phrase "came as a surprise" constrains the subject to be a proposition. A pronominal RE therefore instructs the addressee that its referent is the *most focused entity that is a proposition*. The demonstrative pronoun prefers a non-focused entity, so the demonstrative pronoun in 7C' instructs the listener to choose an unfocused entity that is a proposition.

Another example presented in Gundel et al. (1998) is repeated here as Example 8:

Example (8)

- (a) And the guy wrote little marks on his golf club as to where to put his hands
- (b) And he had marks as to where to put his feet
- (c) And he did it all, uh, very scientifically
- (d) And he got his golf score way down, you know
- (e) And George played with him, like, uh, once or twice
- (f) And each time he'll just, George'll just whack **it**.

The pronoun of interest in sentence 8f refers to the golf ball, which was brought into *activated* status by the

mention of the golf game. Gundel et al. claim that the pronominal reference to the golf ball is licensed because "it would be reasonable to assume that the addressee's attention is focused on the [golf ball]." (pg. 193), but we strongly disagree with this statement. We encounter the same difficulty trying to use focus to explain this example as we did with Example 7. The golf ball is inferred and therefore has activated status, but is not in focus. If it is in focus, then so are the referents of these continuations:

Possible F sentences for Example 8:

1. and each time, George just hates **it** (it = the situation of playing golf with his father)
2. and each time, he'll just, George'll just complain about **it** for the rest of the day (it = his father's annoyingly scientific method)
3. and each time, George'll just duff **it** (it = the shot)
4. and each time, George'll just end **it** early rather than play a whole game with him (it=the golf game)
5. And each time, George'll just come home from **it** all demoralized (it = the game)

The original sentence 8(f) doesn't achieve a felicitous reference to the golf ball because the golf ball is in focus, but instead the reference works because the golf ball is the most likely thing to be whacked. Using an unstressed pronoun here instructs the addressee to search for the most focused entity that can be the object of whacked. Once again, if we allow for all referents in our alternative f sentences to be considered *in focus*, we have lost the distinction between the *in-focus* referents and the other activated referents. We can demonstrate that the golf ball is not actually in focus at the end of sentence 8(e) by using a different verb in 8(f) with looser semantic restrictions to see if the golf ball is selected as the default referent. If we change the sentence to "And each time, George'll loose it", the pronoun becomes ambiguous because there are multiple things that could be lost, including the golf ball, the golf club, and George's temper. If the golf ball were in focus, it would be the preferred interpretation in this sentence.

Another example from Gundel (1998) is reprinted here as example 9. The authors claim that the bull mastiff is not in focus, therefore a pronominal reference to it is inappropriate in sentence B:

Example (9)

- A) Sears delivered new siding to my neighbors with the bull mastiff.
- B) ***It's** the same dog that bit Mary Ben last summer.

The semantic constraints created by the description "the same dog" indicate that the pronoun refers to the bull mastiff, the only entity from sentence 9A that could meet those constraints, yet the sentence is extremely awkward nonetheless. The problem with this sentence isn't that the pronoun is infelicitous, only that the shift to a new topic is very abrupt. Modifying the sentence slightly in 9B' to contain more overt markers that it is a side discussion helps, and modifying it even more in 9B" so that it makes a natural continuation of 9A seems to solve the problem. So the problem with 9B isn't the fact that the pronoun is misbehaving and referring to a non-focused entity, but that it includes a complete change of scene from sentence 9A.

B') **It's** not the same dog that bit Mary Ben last summer, you know. They got a new one.

B'') **It's** such a huge dog, I wonder if it will think the siding is for snacks.

We can also use Example 9 to show that the unfocused bull mastiff cannot be referred to with a pronoun without strong semantic clues. Dropping the predicate complement or using a verb with loose semantic restrictions makes the dog harder and eventually impossible to access. In sentence 10B, the dog is understood as the referent of 'it' because of the semantics of 'ate' and because the siding cannot eat itself. In sentence 10C, the pronoun refers to the dog because it is the only entity from the sentence that can eat, but this sentence is incoherent. Sentence 10D and 10E contain semantic constraints that can apply to the siding, which effectively blocks the mastiff from being referred to since the siding is more focused, the siding is selected as the referent.

Example (10)

- A) Sears delivered new siding to my neighbors with the bull mastiff.
- B) **It** ate the previous siding. (it = the dog)
- C) ***It** ate my flowers. (it = the dog)
- D) **It's** white. (it = the siding)
- E) **It** looks good. (it = the siding)

If the speaker misjudges his addressee's estimation of the most focused entity and doesn't provide enough semantic constraints to preclude reference to that entity, his RE will be incorrectly interpreted. This phenomenon is a major source of humor on television. Example 11 is from a BBC sitcom:

Example (11)

Speaker A: I have a new tattoo, here under my shirt.

Speaker B: Doesn't it, hurt?

Speaker A: No, it's, quite soft.

Gundel has a hard time explaining why sentence 9B is incoherent while sentence 12B works, since it mentions the wind energy project in a similar syntactic position as the bull mastiff in 9B. She explains that the wind energy project is very important in this context, and therefore brought into focus at the end of sentence 12A:

Example (12)

- A) However, the government of Barbados is looking for a project manager for a large wind energy project.
- B) I'm going to see the man in charge of **it** next week.

It appears to this author that semantic factors are at work for this example. There is only one plausible referent from sentence A that meets the constraints 1) a man is in charge of it and 2) the speaker is likely to visit him.⁵ Sentence

12B contains a smoother shift of topic, without the abrupt and unexplained shift of scene in Sentence 9B. We do not, therefore, need to resort to claims that the wind energy project has heightened salience due to pragmatic factors to explain why Example 12 works.

Our next example is reprinted in Hitzeman and Poesio (1998) from the ILEX corpus (Hitzeman et al., 1997) and is provided in that paper as an example of long-distance pronominal reference. The previous mention of the referent of 'him' in the final sentence appears too far back to be a candidate for local focus when the final pronominal reference occurs:

Example (13)

The jeweller_i who made these bangles was particularly interested in the idea of using intrinsically worthless material – material that has been thrown away, old junk – and he_i lavished on those materials an incredibly painstaking and time-consuming technique, so that the amount of time put into the labour of making these jewels bears absolutely no relation to the value of the materials that he_i's used. And if you look at, for instance, the bangle at the bottom – that's the blue and red one – what looks as though it's painted decoration is in fact inlaid; it's bits of cut-off razor-blade, biro, knitting needles, inlaid into layer after layer of resin, which is done in emulation of Japanese lacquer technique. And that particular bangle took **him**_i something like 120 hours of work.

The authors assert that the jeweller is the global topic, and that the pronoun works because it is allowed to refer to the global topic. But the jeweller is the only candidate referent at all here. The semantics of the sentence specify that the referent is 1) male and 2) worked 120 hours to make the bangle under discussion. Since we already know that the jeweller is the maker of the bangle, and since he is the only male entity that has been mentioned, the reference succeeds smoothly. An alternate final sentence shows that entities other than the global focus could also be accessed as long as they are the most salient candidate to match the stated semantic constraints:

Possible continuations for Example 13:

- A) And **it** was the longest period of time he'd ever spent on a piece – over 120 hours.

Two additional examples of inferred entities referenced with pronouns from the Treebank follow. These examples exploit verb semantic restrictions to constrain their referents.

Example (14)

While Vichy collaborated with the Germans during World War II in the deaths of thousands of Resistance fighters and Jews, its officials needed a diversionary symbolic traitor. Marie-Louise, a small-time abortionist, was their woman. She became an abortionist accidentally and continued because **it** enabled her to buy war-rationed goodies.

⁵ The only competitor referent is Barbados, and if the speaker were visiting the man in charge of Barbados, the

RE would probably be a name or title such as "the Governor General."

Example (15)

The U.S. and Soviet Union are holding technical talks about possible repayment by Moscow of \$188 million in pre-Communist Russian debts owed to the U.S. government, the State Department said. If the debts are repaid, **it** could clear the way for Soviet bonds....

3.2. Pronominal reference to new entities

The examples above show activated entities that are referred to with unstressed pronouns. Since the determination of focused entities from among the activated entities is still an open research question, the skeptical reader may still be wondering whether the referent of each of these example pronouns was actually in focus, even though we have argued to the contrary. As additional evidence, consider the following examples from the Treebank in which we find pronominal references to entities that are completely new to the discourse. These pronouns refer to entities that could also be specified by descriptive noun phrases, and are best classified as *uniquely identifiable*:

Example (16)

Rekindled hope that two New England states will allow broader interstate banking boosted Nasdaq's bank stocks, but the over-the-counter market was up only slightly in lackluster trading. The Nasdaq composite index added 1.01 to 456.64 on paltry volume of 118.6 million shares. In terms of volume, **it** was an inauspicious beginning for November.

Example (17)

Mr. Pratt remarked that he thinks steeper prices have come about because producers don't like to see a hit wine dramatically increase in price later on. Even if there is consumer resistance at first, a wine that wins high ratings from the critics will eventually move. "There may be sticker-shock reaction initially," said Mr. Pratt, "but as the wine is talked about and starts to sell, they eventually get excited and decide it's worth the astronomical price to add it to their collection. **It's** just a sort of one-upsmanship thing with some people."

In Example 16, the discussion of stock market status has not previously been tied to the beginning of November, but the referent of the pronoun is something like the start of the stock market for November. The statistics of the previous two sentences provide a description of the results of the previous day's trading, without implicitly evoking a discourse entity for the start of November. This event is new to the discourse at the time of the pronominal mention, but uniquely identifiable based on the predicate complement of this sentence.

Again in Example 17, we find a discourse-new referent mentioned by a pronoun. The referent of the pronoun in the final sentence of 16 is the cause of the behavior described by the above discourse. This entity is too abstract to be considered activated, and would typically be described with a full noun phrase. For

example "The cause of this behavior is the pervasive one-upsmanship among wine collectors."

Notice that the pronouns in Example 16 and 17 are referential. They cannot be analyzed as prop-it or extrapositions or any of the forms of expletive pronouns, such as in the sentence "It's true that President Clinton likes golf." They also don't fit the traditional format for cataphora, which usually occur in highly marked constructions such as "While he was fixing his_i roof, Bob_i noticed his_i neighbor's new pool." They are simply referential pronouns and their connection to the previous discourse must be inferred once the description is encountered. A similar process occurs when an addressee interprets a definite noun phrase. If she doesn't find an entity in the DC that the noun phrase could reasonably describe, a new entity is constructed.

There are also conventionalized contexts in which pronouns can refer to discourse-new entities. One that comes immediately to mind is with superlatives, often found in advertising and journalistic writing. For example, a television ad in the US might start with the phrase "It's the biggest sales event of the season" or a newspaper article might start with "It's the largest black pearl in existence, and it's coming to the City museum next month". These entities would be classified as Uniquely Identifiable on the Givenness scale.

3.3. What about other cognitive statuses?

The examples above provide evidence that pronouns can refer to activated and uniquely identifiable entities. Although this paper doesn't include examples of pronouns referring at every point in the GH, it is possible to construct instances of pronominal reference to two remaining statuses: Referential and Type Identifiable entities.

Pronouns can occur as the first mention of entities in a fictional work. One can imagine a novel that begins with "It haunted my memory..." or "He was the last one to leave..." This produces an effect of cognitive dissonance because the reader doesn't have the information necessary to resolve the reference, and the author exploits this situation to heighten our curiosity about the story. I would classify these referents as referential in the Givenness Hierarchy. They are understood as referring to a particular individual, although we have minimal information about that individual beyond its type. In a narrative format, we allow the pronoun's referent to be developed as the story proceeds and trust that the initially underspecified discourse entity will eventually be fleshed out, and that this process is part of the journey the author wishes to take us on during the course of the work.

At the extreme end of the Givenness Hierarchy, type identifiable entities are indefinite references that are successful as long as the addressee has previous knowledge of this type of entity. Although referring to a type with a pronoun is difficult while maintaining an interpretation that is not referential, one example of reference in this category is the vague *they* used in English to refer to some unspecified set of people, for example:

Example (18)

They always say, "The truth will out."

3.4. Sources of Semantic Enhancement

Semantic constraints on the referent of a pronoun can be produced by a variety of syntactic constituents. Sources we have found so far are:

1. Lexical properties of the pronoun indicate gender, number and animacy.
2. Predicate complements of copular constructions add predication as we saw in the 7C sentences above, as do near-copulars such as *seems/appears to be*, and *call/name/define*. For example (from BUR 1996):
“The governor called **that** a worst-case scenario.”
3. Verb semantics restrict the semantic properties of their objects. Some restrictions are very specific, for example the object of *brandish* must be a weapon, while others are less specific, for example the object of *send* can be any type of object that is movable.
4. The set size of a plural referent can be specified with a variety of constructions, for example:
“Then they both go to Elmira.”
“Then all three of them go to Elmira.”
5. Prepositional post-modification can be used in certain forms of dialog, or to achieve a dramatic effect. For example:
“Then He of the red cape arrived and saved us.”
“Those from Iowa lined up on the left and those from Illinois gathered to the right.”
6. Relative clauses can be appended to pronouns to form a uniquely-identifiable referent. For example:
“He who lives by the sword”

4. Enhancing the Givenness Hierarchy

We hope we have presented ample proof in Section 3 to demonstrate that pronouns may access non-focused entities when enough semantic information is provided. This works only when there are no competitor entities with higher salience that match all the semantic constraints. This sounds suspiciously similar to how noun phrases work. Are we claiming that definite noun phrases and pronominal noun phrases function just the same? No. The conversational implicatures of descriptive noun phrases do contrast with those of pronominal REs, however we argue that this difference depends crucially on the semantic information provided, not simply on the form. This section explores what our findings imply for the Givenness Hierarchy, and how the speaker’s choice of form does impact its assumed salience.

4.1. Revisiting Definite Noun Phrases

Gundel et al. point out that descriptive noun phrases sometimes refer to entities with high salience. In fact, their experimental results for English found definite noun phrases covering many cognitive statuses:

Status	Focus	Activated	Familiar	UI
--------	-------	-----------	----------	----

The N	30	95	47	108
-------	----	----	----	-----

Table 2: Definite Noun Phrases in English

This makes sense if we consider that a full noun phrase can re-mention a focused entity while providing new information about it. For example, in Example 18 the focused entity is re-mentioned with a definite noun phrase without violating Grice’s maxim of quantity:

Example (19)

My son_i came into the room, then the little scamp_i...

Example 19 works because the noun phrase provides new information about the entity; therefore, it is not more informative than necessary. Contrast Example 19 with Example 20, in which re-mentioning the focused entity with a name does violate the maxim:

Example (20)

John_i came into the room, then John_i...

Example 20 violates Grice’s maxim because the second mention of John provides no new information; therefore, a full noun phrase is more informative than necessary (because the information it contains is redundant). A pronoun would have been sufficient here, but since a pronoun was not used, the use of a noun phrase gives the implication that some other John is being referred to. This is the main sense in which noun phrases and pronouns differ. When there are two competing discourse referents of roughly the same salience and with roughly the same semantic constraints, the preferred interpretation of a definite noun phrase that applies to both entities will be the unfocused entity. For example:

Example (21)

- A) I met the neatest woman today. Millie_i and I were having coffee at Starbucks and she_i started telling me about her career as an astronaut.
- B) I met the neatest woman_i today. Millie and I were having coffee at Starbucks and the woman_i started telling me about her career as an astronaut.

In Example 21, Millie is the most focused entity, because she was more recently mentioned, when we reach the second clause. Using the pronoun in the A version resolves to the most salient female entity, in this case Millie. In the B version, the pragmatic implication of using the phrase ‘the woman’ signals that it should be bound not to the most focused woman, but to some other woman.

Noun phrases can easily refer to the focused entity when a very vague noun such as ‘people’ and ‘thing’ is used. These nouns function essentially like pronouns. The addressee will bind these references to the most salient entity that meets all the semantic restrictions, just like a pronoun. For example, if speaker A is describing changing his oil and says “But then the thingy broke as I was removing it”, the addressee will select the most salient entity as the referent of ‘the thingy’ - perhaps the oil filter or whatever he judges to be the most likely entity that one removes during the course of an oil change and that could have broken during its extraction.

Noun phrases are often used to access activated entities, particularly inferred entities or mentioned entities that are no longer in focus. The most salient entity that meets the description will be chosen as the referent. For example, if the oil change example above is changed to: “Then the wrench broke”, the addressee will interpret this to mean “search for a referent that is the most salient thing that is used during an oil change and that can be described as a wrench.” The addressee should bind the reference to the oil filter wrench, assuming he knows that oil filter wrenches are used during oil changes. Salience is used to bind the referent to the oil change wrench rather than some other wrench, preferring activated to uniquely identifiable referents. But the noun phrase must specify the referent in a way that is not interchangeable with a pronoun or else it will violate the maxim of quantity and suggest that a different, less focused entity is actually being discussed. For example, if the final sentence of Example 13 had been “And that particular bangle took **the jeweller**, something like 120 hours of work.” The reference seems to imply that some other jeweller is being referred to, since we already know that the intended referent is a jeweller and he is salient enough to be referred to with a pronoun.

5. Implications and Conclusions

In this paper we reviewed some examples of pronominal reference that cannot be accounted for by previous theories. Although the intuitions behind the Givenness Hierarchy are correct, we believe its definition needs to be expanded slightly to include all semantic constraints on the pronoun, not just those contributed by the RE. We can see in our examples that it is possible to provide a large amount of evidence regarding the semantic properties of the correct referent of a pronoun, even though the pronoun itself contains little semantic material. In fact, we claim that the division between pronominal REs and descriptive noun phrase REs is somewhat artificial, since the RE itself is not the only constituent that contributes semantic constraints for its referent.

Although well-substantiated as a default assumption, the commonly held claim that the form of a referring expression determines the accessibility space of its possible referents has been found to be not entirely true. For semantically enhanced pronouns, the semantic information as a whole is taken into account to determine the possible position on the Givenness Hierarchy that the referent may occupy. As we saw in Example (10), semantic enhancement is necessary for an unstressed pronoun to refer outside of the focus. When semantic enhancement is not provided, the pronoun reverts to its default interpretation, relying on salience alone. Thus there is a tradeoff between predication and salience. The more predication, the less the addressee will rely on salience to interpret the RE, and vice versa. Before committing to the interpretation of a referring expression, an addressee accumulates all the semantic constraints on the referent found in the sentence, regardless of the form of the RE. If multiple discourse entities match the semantic constraints, then Grice’s maxim is applied and the salience implications of different referring expression forms are taken into consideration.

6. Acknowledgements

This material is based on work supported by ONR research grant no. N00014-95-1-1088 and DARPA research grant no. F30602-98-2-0133. Any opinions, findings, conclusions and recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the granting institutions.

7. References

- Nicholas Asher, editor. 1993. *Reference to Abstract Objects in Discourse*. Kluwer Academic.
- Breck Baldwin. 1997. Cogniac: High precision coreference with limited knowledge and linguistic resources. In *Proceedings of the ACL-97 workshop: Operational Factors in Practical, Robust Anaphora Resolution for Unrestricted Texts*, pages 38–45.
- Kaja Borthen, Thorstein Fretheim, and Jeanette K. Gundel. 1997. What brings a higher-order entity into focus of attention? Sentential pronouns in English and Norwegian. In Ruslan Mitkov and Branimir Boguraev, editors, *Operational Factors in Practical, Robust Anaphora Resolution for Unrestricted Texts*, pages 88–93.
- Susan Brennan, Marilyn Friedman, and Carl Pollard. 1987. A centering approach to pronouns. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics (ACL’87)*, pages 155–162.
- BUR: Boston university radio speech corpus. 1996. Linguistic Data Consortium. See details on <http://morph ldc.upenn.edu/Catalog/LDC96S36.html>.
- Jaime G. Carbonell and R.D. Brown. 1988. Anaphora resolution: a multi-strategy approach. In *Proceedings of the 12th International Conference on Computational Linguistics (COLING’88)*, pages 96–101.
- Herbert H. Clark, 1977. *Thinking: Readings in Cognitive Science*, chapter Bridging. Cambridge University Press.
- Peter C. Gordon, Barbara J. Grosz, and Laura A. Gilliom. 1993. Pronouns, names, and the centering of attention in discourse. *Cognitive Science*, 17(3), 311–348.
- H.P. Grice. 1975. Logic and conversation. In P.Cole and J.Morgan, editors, *Syntax and Semantics*, pages 41–58, New York. Academic Press.
- Barbara J. Grosz, Karen Sparck Jones, and Bonnie Lynn Webber, editors. 1986. *Readings in Natural Language Processing*. Morgan Kaufmann Publishers.
- Barbara J. Grosz, Aravind K. Joshi, and Scott Weinstein. 1995. Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21(2):203–226.

- Barbara J. Grosz. 1981. Focusing and description in natural language dialogues. In Aravind K. Joshi, Bonnie Lynn Webber, and Ivan Sag, editors, *Elements of Discourse Understanding*, pages 84--105. Cambridge University Press, Cambridge.
- Jeanette K. Gundel, Nancy Hedberg, and Ron Zacharski. 1993. Cognitive status and the form of referring expressions in discourse. *Language*, 69(2):274--307.
- Jeanette K. Gundel. 1998. Centering Theory and the Givenness Hierarchy: Towards a Synthesis. In Marilyn A. Walker and Aravind K. Joshi and Ellen F. Prince, eds. *Centering Theory in Discourse*, 183—198.
- Janet Hitzeman, Chris Mellish, and J. Oberlander. 1997. Dynamic generation of museum web pages: The intelligent labeling explorer. *Archives and Museum Informatics*, 11:107—115.
- Janet Hitzeman and Massimo Poesio. 1998. Long distance pronominalisation and global focus. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics (ACL'98)*.
- Jerry Hobbs. 1986. Resolving pronoun reference. In *Readings in Natural Language Processing*. Morgan Kaufmann.
- Megumi Kameyama. 1986. A property-sharing constraint in centering. In *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics (ACL'86)*.
- Megumi Kameyama. 1998. Intrasentential centering: A case study. In M. Walker, A. Joshi, and E. Prince, editors, *Centering Theory in Discourse*, pages 89--112. Clarendon, Oxford.
- Lauri Karttunen. 1976. Discourse referents. In J. McKawley, editor, *Syntax and Semantics*, vol. 7, pages 361--385. Academic Press.
- Mitchell P. Marcus, Beatrice Santorini, and MaryAnn Marcinkiewicz. 1993. Building a large annotated corpus of english: The Penn Treebank. *Computational Linguistics*, 19(2):313--330.
- Ruslan Mitkov. 1998. Robust pronoun resolution with limited knowledge. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics (ACL'98)*, pages 869--875.
- Massimo Poesio, Sabine Schulte im Walde and Chris Brew. Lexical Clustering and Definite Description Interpretation. In *Papers from the 1998 AAAI Spring Symposium on Applying Machine Learning to Discourse Processing*. 1998. 82—89.
- Ellen Prince. 1991. The ZPG letter: Subjects, definiteness, and information-status. In S. Thompson and W. Mann, editors, *Discourse description: diverse analyses of a fund raising text*, pages 295--325. John Benjamins B.V.
- Rebecca Schiffman. 1985. *Discourse constraints on 'it' and 'that': A study of language use in career-counseling interviews*. Ph.D. thesis, University of Chicago.
- Ethel Schuster. 1988. Anaphoric reference to events and actions: A representation and its advantages. In *Proceedings of the 12th International Conference on Computational Linguistics (COLING'88)*, pages 602--607.
- Candace L. Sidner. 1983. Focusing in the comprehension of definite anaphora. In Michael Brady and Robert C. Berwick, editors, *Computational Models of Discourse*, pages 363--394. MIT Press, Cambridge.
- Matthew Stone. 1994. Discourse deixis, discourse structure, and the semantics of subordination. Unpublished Manuscript.
- Michael Strube. 1998. Never look back: An alternative to centering. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics (ACL'98)*, pages 1251--1257.
- Bonnie Lynn Webber. 1979. *A Formal Approach to Discourse Anaphora*. Garland Publishing.
- Bonnie Lynn Webber. 1983. So what can we talk about now? In Michael Brady and Robert C. Berwick, editors, *Computational Models of Discourse*, pages 331--371. MIT Press, Cambridge.
- Bonnie Lynn Webber. 1988. Discourse deixis: Reference to discourse segments. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics (ACL'88)*, pages 113--122.