

# RAFT/RAPR and Centering: A Comparison and Discussion of Problems Related to Processing Complex Sentences

Linda Z. Suri\*  
Educational Testing Service

Kathleen F. McCoy†  
University of Delaware

## 1. Introduction

Several researchers have noted the local coherence exhibited by discourse (Sidner 1979; Grosz, Joshi, and Weinstein 1983; Carter 1987; etc.). A primary component of this local coherence is the way the *local focus* of the discourse shifts from one sentence to the next and the way this shifting is marked by linguistic choices made by the writer/speaker. By local focus, we refer to that concept a sentence is most centrally about within the discourse context in which it occurs. This is sometimes called the *topic* or *center*.

A local focusing framework typically consists of focus-tracking algorithms and algorithms for suggesting referents for pronouns. Such a framework can be used in conjunction with an inferencing mechanism to resolve pronouns (and other anaphora) in a Natural Language Understanding system. The focusing framework suggests a referent for a pronoun, and the inferencing mechanism then confirms or rejects the suggested referent on the basis of *semantic factors*, i.e., semantics, world knowledge, etc. The focusing framework is useful because it only requires an inferencing mechanism to confirm a co-specification rather than requiring an inferencing to find the referent independently.

To date, there have been two major frameworks for tracking the local focus from one sentence to the next and for using focus during pronoun resolution. The first framework, *Focusing*, was introduced by Sidner (1979). In this squib, we will describe our framework, Revised Algorithms for Focus Tracking and Revised Algorithms for Pronoun Resolution (RAFT/RAPR), which is based on Sidner's work. RAFT/RAPR can be characterized as maintaining two foci for a sentence: the subject focus and the current focus, which very often have distinct contents. RAFT/RAPR maintains a set of data structures, and uses rules (which take grammatical roles into account) for pronoun resolution and computing the foci. Taken together, these rules describe how focus can (and is most likely to) shift from one sentence to the next. Note that focus tracking and pronoun resolution are mutually dependent processes: focus tracking is necessary for pronoun resolution, and pronoun resolution, in turn, affects focus tracking.

---

\* Educational Testing Service, Mail Stop 10-R, Rosedale Road, Princeton, NJ 08541. E-mail: lsuri@rosedale.org.

† Dept. of Computer and Information Sciences, 103 Smith Hall, University of Delaware, Newark, DE 19716. E-mail: mccoy@udel.edu.

Subsequent to Sidner's work, Grosz, Joshi, and Weinstein (1983) introduced *centering* to account for the same phenomena addressed by Sidner's algorithm.<sup>1</sup> Centering attempted to simplify processing by keeping fewer data structures than Sidner's framework did. In particular, the centering literature claims that, rather than two foci, only one focus is needed, termed the *backward-looking center* (*C<sub>b</sub>*). Pronoun resolution within the centering framework is largely based on an ordering of preferred focus (centering) moves.

Other research on discourse (e.g., Grosz 1981; Grosz and Sidner 1986; Reichman 1978) has studied another phenomenon, the *global focus* of discourse. The term global focus generally refers to the entity or set of entities that are relevant to or salient in the overall discourse; the identification of global focus typically interacts with the identification of *discourse segments*. Global focus and discourse segmentation are distinct from the phenomenon of local focusing that is addressed in this paper. However, we should point out that the centering literature has noted that centering "... is intended to operate within a [discourse] segment" (Walker 1989, p. 253). In our work on RAFT/RAPR we do not restrict the domain of the algorithms to within a discourse segment.

Given that multiple frameworks for focus tracking and pronoun resolution have emerged, we would like to do a comparison to see how the frameworks are the same and how they differ. Previous assessments and comparisons of local focusing frameworks have relied on comparing how frameworks process a small number of constructed discourses, but this kind of comparison is inadequate. Instead the question that must be answered is which framework performs best on naturally occurring text. However, such a comparison is not possible at this point because no framework has fully specified how to handle complex sentences (see Suri [1993] for the details of this argument).

In light of this, we propose a comparison of RAFT/RAPR and centering along two lines. First, it is instructive to take a careful look at how the frameworks handle certain kinds of constructed discourses involving simple sentences. This comparison proves useful for understanding why the frameworks suggest the referents that they do. It is interesting to note that, while the methodologies used in RAFT/RAPR and centering are quite different from one another, the frameworks very often have the same preferences for pronoun resolution for text that is not discourse-initial (nor discourse-segment-initial) and that involves only simple sentences. Despite this similarity, we point out places where the two frameworks differ. A major difference between centering and RAFT/RAPR is that while RAFT/RAPR stacks old focus information, centering keeps information about the previous sentence only. We show why this is problematic for centering. We point out other differences that arise because centering keeps one focus and does not take the grammatical roles of pronouns and potential antecedents into account during pronoun resolution. This difference is evident in the examples discussed in this paper involving discourse-initial text, and even in an example discussed in this paper that (we believe) does not involve a discourse segment boundary. Note that because the centering literature claims that centering should operate only within a discourse segment, and because this claim is used to explain some otherwise problematic cases of pronoun use, not being able to adequately handle discourse seg-

---

<sup>1</sup> Notice that we use the term focusing to cover all local focusing frameworks, Sidner's focusing framework (Sidner 1979), Carter's extensions to Sidner's framework (Carter 1987), the centering framework (Grosz, Joshi, and Weinstein 1983 and others), our framework (RAFT/RAPR), PUNDIT (Dahl [1986] and others), etc. We use uppercase ("Focusing," or "Local Focusing"), or "Sidner's Focusing Algorithm/Framework" to refer to Sidner's work. We use RAFT/RAPR to refer to our work.

ment initial text is much more of a problem for the centering frameworks than may at first be apparent.

While the observations we make in this first comparison are intriguing, it would be inappropriate to assess and compare the frameworks only on the basis of a handful of constructed texts. However, in order to do a corpus analysis to compare focusing frameworks one must be able to handle many kinds of complex sentences. Thus, we developed a methodology for determining how people process a particular kind of complex sentence.<sup>2</sup> Our second line of comparison of frameworks involves studying how well each framework can be extended to account for such findings. Suri (1993) presented preliminary results for processing sentences of the form “SX because SY,” where SX and SY each consist of a single clause. In this squib, we report some of those findings, and discuss extending RAFT/RAPR and centering in light of these findings. In closing, we summarize the abstract similarities and differences between centering and RAFT/RAPR.

## 2. Our Focusing and Pronoun Resolution Algorithms

Below, we discuss the behavior of our algorithms for simple (i.e., single-clause) sentences. (See Suri [1993] for a fuller discussion of our focusing and pronoun resolution algorithms, and a discussion of how our algorithms differ from Sidner’s.)

### 2.1 Data Structures

Our algorithms maintain more focusing data structures than centering does, but each data structure is motivated by discourse processing needs. Below are the data structures that RAFT/RAPR uses:

- Current Focus (CF): the item computed to be the *local focus* of the sentence.
- Potential Focus List (PFL): all NPs other than the CF and SF, ordered according to the following: direct object, indirect object, all other NPs in surface order within the clause.
- Subject Focus (SF): the surface subject of the clause, except in certain cases as mentioned later. (The need for an SF as well as a local focus is discussed in Section 4.3.)
- The Potential Subject Focus List (PSFL): all NPs other than the SF and CF, ordered as follows: direct object, indirect object, all other NPs in surface order within the clause.<sup>3</sup>
- CF stack, SF stack, PFL stack, PSFL stack. We stack the foci and foci lists after each sentence. (See Section 4.1.)

### 2.2 Resolving Pronouns (in Simple Sentences)

RAFT/RAPR resolves pronouns based on the grammatical role of the pronoun and the focusing data structures from the previous sentence. For a *nonsubject third person*

<sup>2</sup> Suri (1993) discusses the issues that one faces in trying to determine how to process complex sentences, and argues why this methodology is better than alternative methodologies.

<sup>3</sup> Determining whether we truly need both a PFL and PSFL would require a corpus analysis, which is beyond the capabilities of the current technology, as discussed in Suri (1993). In fact, the motivation for maintaining both a PFL and a PSFL is based on processing complex sentences (see Suri 1993).

*singular pronoun*, our algorithm first proposes the CF (of the last sentence) as the co-specifier, then the SF, then members of the PFL, and then, under preferences yet to be determined, the members of the CF stack, SF stack, PFL stack, and PSFL stack.

We thus prefer the pronoun to co-refer with the last focus, then the subject focus of the previous sentence, then some other NP introduced in the previous sentence, and then elements that have been in focus or mentioned in previous sentences. Each attempted co-specification may be rejected by a separate inferencing component on the basis of semantic factors (semantics, world knowledge, etc.) or based on syntactic constraints.

For a *subject third person singular pronoun*, we first try the SF, then the CF, then members of the PSFL, then the stacked elements.

### 2.3 SF and CF Computation (in Simple Sentences)

For a *there*-insertion sentence, the SF is the deep subject of the sentence, but for most simple sentence types the SF is the surface subject of the simple sentence.<sup>4</sup>

Our algorithms compute the CF of the current sentence based on the following *interacting* criteria:

1. Co-specification: Prefer elements that co-specify an element in a focusing data structure over elements just introduced. If an element being talked about now has been talked about before, it is more likely to be the topic (and thus to continue to be talked about) than something that has just been introduced.
2. The type of realization of each element: Prefer NPs realized as pronouns over those realized with full NPs. A pronoun is more likely to be talked about in subsequent text than a full NP (Brown 1983).  
To appreciate this preference, consider the following. A pronoun carries less semantic information than a full NP. If a writer chose to communicate an element using a pronoun, he or she must have believed that the element was highly focused enough that the reader would not have difficulty interpreting the pronoun without the extra semantic information that would be communicated with a full NP.
3. Which focusing data structure, if any, is co-specified by each NP.  
In general, we believe a writer/speaker is more likely to keep discussing the focus than to move the local focus to some other discourse entity. Thus, we prefer the CF to remain constant from one sentence to the next. We refer to this preference as the *focus retention preference*. We also believe a writer is more likely to switch focus to an element that was just mentioned in the previous sentence than to one that was discussed earlier. Thus, in sum, we prefer for the CF to co-specify the last CF rather than something on the last PFL or the last SF, and we prefer for the CF to co-specify something on the last PFL or the last SF rather than a stacked element.
4. Syntax: We prefer for the CF to be a nonsubject rather than a subject, although the CF can be the subject.

---

<sup>4</sup> We need to explore how to compute the SF for other sentences that involve pleonastic subjects.

Recall that we compute the SF based on subject NPs. We prefer for the CF and SF to be distinct, and for the SF to be a subject and the CF to be a nonsubject since we believe that it is often the case that there are two elements that are being focused on, and we believe that a subject pronoun is more likely to refer back to a subject and a nonsubject pronoun is more likely to refer back to a nonsubject. Thus, we want to record information in the focusing data structures in a manner that will allow us to take the grammatical roles of potential antecedents into account.

5. Syntactic forms and clue words: The use of a particular linguistic form (e.g., a *there*-insertion structure) or clue words (e.g., “but”) also influences the choice of the local focus.

Again, these (abstract) criteria *interact* in determining the choice of the CF. For example, among nonsubject NPs co-specifying elements of the PFL, we prefer a pronoun to a full NP; on the other hand, if the only pronoun is the subject, and it co-specifies a member of the PFL, and a nonsubject nonpronoun co-specifies the CF, we will not shift the CF to a member of the last PFL. That is, the CF will remain the same despite the pronominalization. Also, we prefer a nonsubject nonpronominal NP co-specifying a member of the PFL over a subject pronoun co-specifying the CF.

#### 2.4 An Example of Pronoun Resolution with RAFT/RAPR

The following example (modified from Brennan, Friedman, and Pollard [1987]) illustrates RAFT/RAPR.

##### Example 1

(S1) Susan drives a Ferrari. SF= [Susan], CF=[Ferrari] (based on surface syntactic preferences)

(S2) She drives too fast. SF= [Susan], CF=[Susan]

(S3) Lyn races her on weekends. SF= [Lyn], CF=[Susan], PFL=([weekends])

(S4) She often beats her.

For the first sentence, we compute the SF to be the surface subject, [Susan], and the CF to be the direct object, [a Ferrari].

For S2, we first try to resolve the subject, “She,” using SF(S1); since this is not rejected by inferencing with semantic factors and syntactic information, the pronoun is resolved to refer to [Susan]. SF(S2)= [Susan], the surface subject of the sentence. Since the only co-specifier to something previously mentioned is [Susan], CF(S2)=[Susan] as well.

For S3, we try to resolve the nonsubject pronoun “her” using CF(S2). Since this is not rejected by inferencing with semantic factors and syntactic information, the pronoun is resolved to refer to CF(S2), [Susan]. We compute SF(S3) to be the surface subject, [Lyn]. Since [Susan] (“her”) is the only thing that specifies anything previously mentioned, CF(S3)=[Susan].

In (S4), we first try to resolve “She” using the SF and “her” using the CF. This interpretation is not rejected on the basis of inferencing with semantic factors or syntactic information. Thus, we get the reading “Lyn often beats Susan,” which is the same reading that native speakers of English preferred in an informal poll.

**Table 1**  
Centering transitions.

	Cb(Un)=Cb(Un-1)	Cb(Un)≠Cb(Un-1)
Cb(Un)=Cp(Un)	Continue	Smooth-Shift
Cb(Un)≠Cp(Un)	Retain	(Rough-)Shift

**No-Cb:** No element in  $U_n$  realizes an element of  $Cf(U_{n-1})$ .

### 3. Centering

In this section, we give a brief introduction to the centering framework. This framework was introduced by Grosz, Joshi, and Weinstein (1983), it has been discussed and/or expanded on in several other works, including Brennan, Friedman, and Pollard (1987), Walker (1989, 1993), Kameyama (1986, 1993), Walker, Iida, and Cote (1992), Brennan (1993), Kameyama, Passonneau, and Poesio (1993), Linson (1993), and Hoffman and Turan (1993).

#### 3.1 Computing Centers

For each utterance, centering computes the following (Grosz, Joshi, and Weinstein 1983; Brennan, Friedman, and Pollard 1987; Walker, Iida, and Cote 1992): the *backward-looking center* (*Cb*), which is intended to capture that item which ties the current sentence in with the previous sentence in the discourse, and a list of *forward-looking centers* (*Cf*), or elements that can potentially be the *Cb* of the next sentence. For English, the *Cf* is ordered or ranked by grammatical relations to the main verb (Walker 1989); the order is “first the subject, object, and object2, followed by other subcategorized functions, and finally adjuncts.” (Brennan, Friedman, and Pollard 1987, p. 156)

The first element of the *Cf* list is the *preferred center* or (*Cp*). As the name implies, the *Cp* is the element that is considered most likely to become the *Cb* of the next utterance.

After the first sentence, the *Cb* of utterance  $U_n$ ,  $Cb(U_n)$ , is the highest ranked element of the *Cf* of the previous utterance,  $Cf(U_{n-1})$ , that is realized in the current utterance. We will refer to this condition as *the Cb constraint*.

#### 3.2 Resolving Pronouns

The centering literature has identified four kinds of transitions between sentences (Walker, Iida, and Cote 1992), where  $U_n$  is the current utterance, and  $U_{n-1}$  is the previous utterance. The transitions are described in Table 1.<sup>5</sup>

Centering has used the following rules.

##### Rule 1

“If some element of  $Cf(U_{n-1})$  is realized as a pronoun in  $U_n$ , then so is  $Cb(U_n)$ .” (Brennan, Friedman, and Pollard 1987, p. 156).

##### Rule 2

Continuing is preferred over retaining, retaining is preferred over smooth-shifting,

<sup>5</sup> In describing these transitions, we assume the sentence has a subject.

and smooth-shifting is preferred over shifting. (Brennan, Friedman, and Pollard 1987; Walker, Iida, and Cote 1992).

In resolving pronouns, the Brennan, Friedman, and Pollard (1987) (and Walker, Iida, and Cote [1992]) algorithm uses *Rule 1*, *Rule 2*, and the *Cb constraint*. Basically, the algorithm generates all possible co-specifications for pronouns in the current sentence by generating all possible Cb and Cf pairs for the current sentence (corresponding to various co-specifications), and then it *filters* and *ranks* these possible co-specifications. Filtering eliminates co-specifications that violate the *Cb constraint* or *Rule 1*, and co-specifications involving conraindexing problems. The ranking of co-specifications is based on the *Rule 2* transition preferences. In this way, the algorithm proposes co-specifications for pronouns that comply with the constraints and rules of centering, in an order that corresponds to centering's preferences for centering moves. Notice that the Cb of the current sentence is computed as a side effect of doing the pronoun resolution.

### 3.3 Example 1 with Centering

To illustrate centering, we look again at Example 1.

#### Example 1

- 1) Susan drives a Ferrari. Cf=(*[Susan]*, *[Ferrari]*)
- 2) She drives too fast. Cb= *[Susan]*, Cf=(*[Susan]*)
- 3) Lyn races her on weekends. Cb= *[Susan]*, Cf=(*[Lyn]*, *[Susan]*)
- 4) She often beats her.

According to the Cb constraint, Cb(4)=*[Lyn]*, since Lyn is the highest ordered element of Cf(3). Thus, if "She" were *[Lyn]* and "her" were *[Susan]*, then Cb(4)=*[Lyn]*, Cp(4)=*[Lyn]*, and this would be a smooth shift. But, if "She" were *[Susan]* and "her" were *[Lyn]*, then Cb(4)=*[Lyn]*, Cp(4)=*[Susan]*, and this would be a shift. Thus, centering prefers the reading "*[Lyn]* often beats *[Susan]*." Thus, centering and RAFT/RAPR agree on how to resolve these pronouns.

## 4. Discussion of Some Differences

Although centering and RAFT/RAPR seem to resolve pronouns similarly in many cases, there are some situations where the two algorithms differ.

### 4.1 Stacking Focus Information

A major difference between the two frameworks is that by maintaining stacks of CFs and SFs, RAFT/RAPR allows a writer to pronominally refer to a CF or SF of a sentence before the previous one. Centering, which has no counterpart to our stacks, does not.

The stacking of focusing information for RAFT/RAPR (or focusing algorithms in general) is motivated by examples of text in which the referent of a pronoun is more than one sentence back, and for which there is no indication that there is an intervening discourse segment. Example 2 is such an example.

#### Example 2

With a strain, *he* could see a glimpse of the river to the northwest. The walls were Sheetrock and bare. She had picked out some artwork. *He* determined that the Ego Wall would face the desk, behind the wing chairs. (Grisham 1991, p. 139).

Note that one might argue that the sentence “She had picked out some artwork” constitutes a discourse segment unto itself marked by the use of the past perfect (note the simple past is used in other sentences).<sup>6</sup> However, even if this were the case (and that sentence was ignored in the pronoun resolution), the centering framework would still not appropriately resolve the “He” in the final sentence because there would still be an intervening sentence (“The walls were Sheetrock and bare.”) that does not contain the antecedent of “He.” We can see no reason why this (other) intervening sentence should be seen as a discourse segment.

#### 4.2 Discourse-Initial Sentences

Recall that centering bases its pronoun decisions on the type of “centering” transitions that would be manifested by various interpretations. These transitions require that the Cb of the previous sentence be defined. As a result, centering can make no predictions about pronouns until the third sentence in a discourse (segment). Consider the following:

##### Example 3

- 1a) Lyn races Susan on weekends.
- 2a) She races Jack during the week.

Because no Cb is established in 1a (since 1a is discourse-initial), centering makes no prediction as to whether the pronoun in 2a should refer to [Lyn] or [Susan]. Our pronoun resolution algorithms predict that “she” is [Lyn], the SF of 1a.

Of course, one way around this problem is for a likely Cb for the first sentence to be calculated and used in pronoun resolution. For instance, because subjecthood is very important in centering, one might take the subject (or Cp) of a discourse-initial (or discourse-segment-initial) sentence as its Cb.<sup>7</sup> This would allow centering to resolve the referent of “She” for this text correctly. However, such a decision would cause centering to resolve “her” in 2b below incorrectly:

##### Example 4

- 1b) Lyn races Susan on weekends.
- 2b) Jack races her during the week.

An informal poll suggests that native speakers prefer to interpret “her” as referring to [Susan]. The problem here is that for centering, if there is a single pronoun in a sentence, it becomes the Cb, and centering prefers to resolve the Cb so that it co-refers with the Cb of the previous sentence regardless of the grammatical role of the Cb in the previous sentence. Thus, if centering is to stay within its paradigm which was set up for processing discourse segment internal pronouns, it will resolve a pronoun in one of these two examples inappropriately. (Notice no semantic or world knowledge inference could flag the mistake.)

---

<sup>6</sup> We discuss aspectual classification and the possible import for focusing and pronoun resolution in Suri (1993).

<sup>7</sup> In fact, this is what Brennan, Friedman, and Pollard (1987) appear to do for their examples.



Our algorithms consider the grammatical role played by a pronoun and its potential antecedents. They correctly resolve the pronouns in these examples.<sup>8</sup>

### 4.3 Pronominalization and Grammatical Roles

Brennan, Friedman, and Pollard (1987) discuss an ambiguous discourse. (A modified version is shown in Example 5 below.) Centering resolves the ambiguous pronoun ("She" in 4) one way, and RAFT/RAPR resolves it another way. This is not particularly interesting since the discourse is ambiguous.

#### Example 5

- 1) Susan drives a Ferrari.
- 2) She drives too fast.
- 3) Lyn races her on weekends.
- 4) She wins a lot of trophies.

(RAFT/RAPR: She=[Lyn], centering: She=[Susan])

What is more interesting is that if "her" in 3 is replaced by a full NP, then readers do have a strong preference for the pronoun to be resolved with [Lyn].

#### Example 6

- 1) Susan drives a Ferrari.  
SF=[Susan], CF=[Ferrari]; Cf=(*[Susan]*, *[Ferrari]*)
- 2) She drives too fast.  
SF=[Susan], CF=[Susan]; Cb=*[Susan]*, Cf=(*[Susan]*)
- 3) Lyn races Susan on weekends.  
SF=[Lyn], CF=[Susan]; Cb=*[Susan]*, Cf=(*[Lyn]*, *[Susan]*)
- 4) She wins a lot of trophies.  
(RAFT/RAPR: "She"=*[Lyn]*; centering: "She"=*[Susan]*)

Notice that RAFT/RAPR computes the "She" in 4 to be [Lyn] since it prefers to resolve subject pronouns to be the SF of the previous sentence. This interpretation agrees with native speaker preferences for this text. On the other hand, centering computes the Cb in 3 to be [Susan] and as a result it will prefer that "She" is [Susan] in 4 (by preferring a continue transition). Thus, centering's prediction does not agree with native speakers' preferences for this text.

The problem for centering may be viewed as resulting from the fact that if a sentence, Un, has only one pronoun, then centering (Brennan, Friedman, and Pollard 1987) makes that element the Cb(Un) (assuming it co-specifies something in Cf(Un-1)), and centering prefers to resolve pronouns so that Cb(Un)=Cb(Un-1) regardless of the grammatical roles that Cb(Un-1) played in Un-1.

Because our algorithms track two foci and consider the grammatical roles of a pronoun and its possible antecedents during pronoun resolution (by preferring the

<sup>8</sup> Working within centering, Kameyama (1986) proposes preferences based on grammatical roles in resolving pronouns in certain situations. However, she does not address the discourse-initial problems. Also, Brennan, Friedman, and Pollard (1987) claim Kameyama's proposed rule falls out of their transition preferences, but whether this is true in all contexts is not immediately evident.

SF to the CF when resolving a subject, and the CF over the SF when resolving a nonsubject), we prefer for “She” to be [Lyn], the SF of 3.<sup>9</sup>

One way around this example is for centering to claim that 3 actually starts a new discourse segment.<sup>10</sup> However, recall that centering is intended to operate within a discourse segment and discourse-segment-initial processing is problematic for centering (see Section 4.2). Thus, more work would still be required to enable centering to handle pronouns in Example 6 and the identical discourse with the fourth sentence replaced by “Jack races her during the week.”

#### 4.4 What Is the Import of the Above Analysis?

Two observations should be made at this point. First, it is very difficult to identify the focus of a discourse-initial sentence, and neither framework is likely to process all discourse-initial text correctly. This in itself is not terribly problematic, since discourse-initial text is relatively rare (since there is only one discourse-initial sentence for each naturally occurring text).<sup>11</sup> The problem for centering is compounded, however, since centering is intended to work *within a discourse segment*. This means that all *discourse-segment-initial* text will be problematic. The proportion of discourse-segment-initial text is much larger. This fact is much more evident if one notes that presumably all cases where a stack is needed by RAFT/RAPR (see Section 4.1) are cases where a centering theorist would need to argue for an intervening discourse segment.

Second, the reader should observe that we have pointed out a potential problem for centering by citing two examples (Examples 3 and 4) that could not both be handled by centering. This is clearly insufficient. In fact, while RAFT/RAPR does handle both of these examples appropriately, there is no doubt that one could come up with a discourse-initial example where RAFT/RAPR would have difficulty.<sup>12</sup> The question that must be addressed is which framework would correctly process the most *naturally occurring* discourse-initial texts. Likewise, in Section 4.3 we showed that RAFT/RAPR could appropriately handle a non-discourse-initial text, which centering could not. But, again, the question we would want to address is which framework would correctly process the most *naturally occurring* non-discourse-initial texts. The intent of the above analysis was to point out differences in the abstract/underlying preferences of the two frameworks, and how those differences manifest themselves during pronoun resolution. In particular, RAFT/RAPR’s abstract/underlying preference for considering both the grammatical role of a pronoun and the grammatical roles of its potential antecedents is not shared by centering.

### 5. Extensions for Handling Complex Sentences

While we have shown several problematic cases for centering for discourses involving constructed discourses consisting of simple sentences, clearly such a comparison is

9 Note that the tracking of two foci is important in order to correctly handle this text *and* Example 5 with sentence 4 changed to “Jack races her during the week.”

10 In fact, that “Susan” is not pronominalized in S3 does seem to signal a shift in subject focus from [Susan] to [Lyn]. But whether this would constitute a new discourse segment is unclear. Furthermore, to rely on the claim that this is a new discourse segment, centering would need to better specify how to identify the start of a new discourse segment, and to better specify how to handle discourse-initial and discourse-segment-initial text.

11 Furthermore, it is a bit odd to speak of the focus or foci (e.g., CF and SF) of a discourse-initial sentence, since it is usually difficult for a reader to determine what the topic of a discourse-initial sentence is from the discourse-initial sentence alone.

12 In particular, we think that there are likely to be texts that do not involve the kind of “parallel verb structures” of “NP<sub>1</sub> races NP<sub>2</sub> ...” exhibited in the examples discussed above for which centering would predict the correct referent, but RAFT/RAPR would not.

inadequate. A better comparison of the the two frameworks might be based on a large corpora analysis to determine which framework gets the correct pronoun resolution more often and by suggesting the fewest number of referents on average. There are a number of problems one faces when trying to do a corpora analysis. One problem is that pronoun resolution is affected not only by focusing information, but also by semantic factors. In order to identify which framework best captures the pronoun resolution preferences speakers have *based solely on the type of information intended to be captured by focusing frameworks*, we would need to control for how often semantic factors are affecting pronoun resolution.

In this squib, we are concerned with another, perhaps more interesting, problem that precludes undertaking a corpora analysis: although complex sentences are prevalent in written English, neither framework has explicitly specified how to handle complex sentences.<sup>13</sup> In fact, some previous research handled complex sentences in an inconsistent manner (Sidner 1979; Linson 1993). Other research made claims that were not accurate because it did not acknowledge the need to specify how to process complex sentences. In particular, centering literature (Grosz, Joshi, and Weinstein 1983) questioned Sidner's use of an Actor Focus using text that involved complex sentences. Suri (1992) argued that the reason Sidner's algorithms could not handle that text was not because they used two foci, *but because a method for processing of complex sentences was not specified*.

The central question is whether one should process a complex sentence as multiple sentences or as a single sentence. For instance, should the clauses of a complex sentence be processed in linear order, updating the focusing (e.g., RAFT/RAPR or centering) data structures as if one were processing a linear sequence of simple sentences? Or, is some other order more appropriate? More specifically, 1) How do you resolve pronouns within a complex sentence? Do you prefer to resolve a pronoun in a noninitial clause with elements from the previous clause (or the matrix clause) or with elements from the previous sentence? 2) How do you resolve pronouns in a subsequent sentence? In other words, how should one update focusing (e.g., RAFT/RAPR or centering) data structures after a complex sentence? Answers must be provided based on the kinds of syntactic structures (e.g., sentential verb complements, conjunctions) used in the sentence. Answers might also depend on other things, such as whether a verb is subject-control or object-control (e.g., processing "NP1 promised NP2 SX" vs. "NP1 persuaded NP2 SX").

Suri and McCoy (1993) proposed a two-part methodology for determining how to process a given kind of complex sentence. (Suri [1993] explains why this methodology is superior to alternative methodologies.) Note that a local focusing framework (e.g., RAFT/RAPR or centering) is intended to capture preferences for pronoun resolution *independent* of world knowledge, semantics, and rhetorical relations. Thus, the aim of the first part of the methodology is to identify how people prefer to resolve pronouns based on previous pronominalization, focus history, and syntax. Because it is difficult to create texts where pronoun resolution is not affected by world knowledge, semantics, and rhetorical relations, we develop texts that are intentionally *semantically slanted* for pronoun interpretation based on world knowledge, semantics, and rhetorical relations. Several semantically slanted texts are created in order to test the effect of the syntactic form of interest (e.g., "SX because SY") on pronoun resolution within the complex sentence, and in a sentence following that kind of complex sentence. We then

<sup>13</sup> Suri (1992) identified the need to address this problem, and made some preliminary proposals for particular sentence types. Walker (1989) indicated that centering needs to handle multiple subjects, but did not specify how to do that.

gather judgments about the acceptability of these texts from native English speakers. These judgments suggest that the readers' preferences are based on syntax, pronominalization, and focusing history. For example, judgments suggested that readers prefer the subject of the SX clause over the subject of the previous sentence when resolving a subject in the SY clause of an "SX because SY" sentence. That is, *when this hypothesized preference agreed with the semantic-slanting of the text, the text was judged acceptable. But when this hypothesized preference was at odds with the slanting of the text, the text was judged awkward or ambiguous.*<sup>14</sup> We then verify such findings by analyzing the use of anaphors (within and following sentences of the form of interest) in naturally occurring text.

### 5.1 "SX because SY" Sentences

Based on our analysis of grammaticality judgments from native speakers of semantically slanted discourses involving a simple sentence, an "SX because SY" sentence ( $S_n$ ), and a simple sentence, we concluded (Suri and McCoy 1993):

1. Readers prefer to resolve Subject(SX) with Subject( $S_{n-1}$ ).
2. Readers prefer to resolve Subject(SY) with Subject(SX) (over Subject( $S_{n-1}$ )).
3. Readers prefer to resolve Subject( $S_{n+1}$ ) with Subject(SX).

Note that these preferences<sup>15</sup> refute an assumption sometimes made by researchers regarding complex sentences: that the clauses of complex sentences can be processed in a linear order. While our findings indicate that *the pronoun resolution* within S2 does happen linearly, the appropriate contents of *the focusing data structures* after processing S2 should be much more heavily influenced by the SX clause (and not by the SY clause, as the previous assumption would require).

Given these findings, we extended RAFT/RAPR to process sentences of the form "SX because SY" as follows:

1. For resolving a Subject(SX) pronoun, first propose SF( $S_{n-1}$ ) as the referent.
2. For resolving a Subject(SY) pronoun, first propose Subject(SX).
3. Compute the SF of a sentence of the form "SX because SY" to be Subject(SX).

These extensions allow RAFT/RAPR to propose antecedents for discourses of the same form as the discourses judged to be acceptable more efficiently than for discourses of the same form as those judged to be awkward. In this way, the RAFT/RAPR processing would "match" native speakers' judgments.

### 5.2 Centering and "SX because SY" Sentences

One must question whether or not the results of our analysis are directly applicable to the centering framework. However, we find several problematic cases for centering in processing this kind of complex sentence. First, consider the approach of processing "SX because SY" sentences one clause at a time, linearly, where the centering information from SY will be used in processing S3. (Note: this kind of processing for extending

<sup>14</sup> See Suri and McCoy (1993) or Suri (1993) for a discussion of the methodology.

<sup>15</sup> Results in Caramazza, Grober, Garvey, and Yates (1977) suggest that perhaps we should consider whether the verb in the SX clause is an NP1-bias or NP2-bias verb in formulating these preferences.

RAFT/RAPR was rejected by our analysis.) Under this approach, for Example 7, we find:

**Example 7**

- S1) Dodge was robbed by an ex-convict    Cb=Dodge?; Cf=Dodge, ex-con, night.  
the other night.
- S2X) The ex-convict tied him up because    Cb=Dodge; Cf=ex-con, Dodge; retain?
- S2Y) he wasn't cooperating.                Cb=Dodge; Cf=Dodge; continue.
- 3b) # Then he started screaming for help.    Cb=Dodge; Cf=Dodge; continue.

This text looks highly coherent according to centering, but it was judged to be awkward or ambiguous by native speakers of English in a controlled survey (judgments gathered were: 3 acceptable; 23 awkward; 5 ambiguous). One might conclude that the apparent problem for centering stems from violating the conclusion that the text should not be processed linearly. Consider altering centering so that the Cb and Cp of S2 is chosen from the SX clause in a manner consistent with the solution for RAFT/RAPR that we outlined above. Let's look at Example 7 using this modification:

**Example 7**

- S1) Dodge was robbed by an ex-convict    Cb=Dodge?; Cf=Dodge, ex-con, night.  
the other night.
- S2) The ex-convict tied him up because    Cb=Dodge; Cf=ex-con, Dodge; retain?  
he wasn't cooperating.
- S3b) # Then he started screaming for    Cb=Dodge; Cf=Dodge; continue.  
help.

Centering would still predict that this text should be highly coherent, although the text was judged awkward or ambiguous by native speakers.<sup>16</sup> Furthermore, let's look at another example:

**Example 8**

- S1) Dodge was robbed by an ex-convict    Cb=Dodge?; Cf=Dodge, ex-con, night.  
the other night.
- S2) The ex-convict tied him up because    Cb=Dodge; Cf=ex-con, Dodge; retain?  
he wasn't cooperating.
- S3a) Then he took all the money and ran.    Cb=ex-con; Cf=ex-con; smooth-shift

This text would probably also be judged acceptable/coherent by centering using this strategy, and it is acceptable to native speakers (judgments gathered were: 29 acceptable; 3 awkward). What is interesting to note is that this example would probably be considered less acceptable than Example 7, since this example ends in a smooth-shift while the other ends in a continue.

**5.3 A Possible Centering Explanation**

Work by Walker (1993) suggests that sentences beginning with "Now" (which she assumes mark new discourse segments) are more likely to be followed by a smooth-shift or retain transition rather than by a continue. While Walker bases her findings

---

<sup>16</sup> Note that if one were to try to use centering for natural language generation, centering would (probably) predict that one could pronominalize the subject of S3 in Example 7 (since S3 is a continue), while in fact a full NP is required to make the text less awkward/ambiguous.

on a (spoken English) corpora analysis and does not indicate how she processed any complex sentence preceding a sentence starting with “Now,” her results suggest one should study whether sentences starting with “Then” are more likely to be followed by a retain or smooth-shift (than a continue).<sup>17</sup> If this is found to be the case, then an appropriate corresponding revision of centering’s *Rule 2* would allow centering, like RAFT/RAPR, to propose antecedents for discourses of the same form as Example 8, i.e., discourses judged to be acceptable, more efficiently than for discourses of the same form as Example 7, i.e., those judged to be awkward or ambiguous. Therefore, further study of complex sentences and clue words such as “Then” is needed.

#### 5.4 Conclusions Regarding “SX because SY” Extensions

In sum, we extended RAFT/RAPR to process sentences of the form “SX because SY,” in a manner that reflected the judgments given by native speakers, by simply specifying how to resolve pronouns in this kind of sentence and how to update the focusing data structures after processing this kind of sentence. However, there appears to be no similarly straightforward way to extend the existing centering framework to reflect these judgments.

### 6. Comparing the Frameworks for Processing Simple Sentences

There are a number of distinctions between RAFT/RAPR and centering that should be examined. This discussion assumes the algorithms are being applied to single-clause sentences. Many of the comments generalize to processing complex sentences.

- RAFT/RAPR maintains two foci: a subject focus and a current focus. Centering maintains one focus, the Cb.<sup>18</sup>
- RAFT/RAPR resolves nonsubject pronouns in a different manner than subject pronouns, while centering does not. This results in RAFT/RAPR and centering having different preferences for Example 6 (where RAFT/RAPR matches native speaker intuitions), for some discourse-initial texts (see Section 4.2), and for some ambiguous discourses (e.g., Example 5).
- RAFT/RAPR resolves pronouns by searching data structures in an order based on several factors, including the importance of grammatical roles and preferences for focusing movement. Centering resolves pronouns by generating all possible co-specifications and then filtering and ranking them based on a number of constraints and rules.
- By using *Rule 1* to eliminate possible co-specifiers, centering mixes the process of pronoun resolution with focus computation. The RAFT/RAPR approach resolves pronouns and then updates the focus. Furthermore, although RAFT/RAPR recognizes pronominalization as a signal of focus when computing the CF, it does not apply a rule like centering’s *Rule 1*

<sup>17</sup> Related to this matter, when we gave readers the same discourses discussed earlier with the “Then” deleted from the S3 sentences, the judgments were still supportive of our previous findings, but they were more divided than for the discourses including the “Then.” This supports the possibility that “Then” may be playing a role in the pronoun resolution for the S3 sentences. See Suri (1993) for further discussion.

<sup>18</sup> However, we note the Cp is very similar to the SF in that both are computed (for English) to be the surface subject of the sentence (for simple sentences). In this way, one could argue that centering does track two foci, to the same extent that RAFT/RAPR does.

without considering the grammatical roles of anaphors and focusing history.

- Because the *abstract preferences underlying* the frameworks share much in common, the approaches very often make the same predictions. RAFT/RAPR's focusing and pronoun resolution algorithms reflect the following underlying abstract preferences:
  1. We prefer for the CF to be something that is co-referential with an element mentioned earlier in the discourse.
  2. We prefer for the CF to be a pronoun rather than a full definite description.
  3. We use preferences for resolving pronouns and computing the CF that involve the grammatical role of the pronoun.
  4. We prefer for the CF to be the same as the last CF and the SF to be the same as the last SF.

Centering shares the first two of these in common with us. But, centering prefers the local focus to stay the same *and* for the subject to be the local focus (which is stronger than preference 4). As noted earlier, we do not apply preference 2 without regard for the grammatical roles of the NPs and which focusing data structures they co-specify. Furthermore, centering invokes the second preference during pronoun resolution, while we invoke it only after resolving the pronouns. In addition, centering requires that the focus (Cb) realizes an element in the immediately preceding sentence, as opposed to merely preferring for the focus to be something in the immediately preceding sentence but allowing it instead to be co-referential with an element further back in the preceding text (as RAFT/RAPR does).

In sum, the order in which these abstract preferences are applied differs between the two frameworks. Furthermore, as a result of (some) differences between the frameworks with respect to these preferences, resolving a pronoun to have the same grammatical role as its antecedent is less important in centering than it is in RAFT/RAPR (Section 4.3).

- RAFT/RAPR presents possible referents for pronouns one possibility at a time, and if pragmatic, semantic, and general knowledge inferencing rejects a referent, RAFT/RAPR proposes an alternative. Centering, on the other hand, (in addition to requiring this same kind of inferencing) sometimes suggests multiple possibilities for the co-specifications of pronouns in a sentence (Walker 1989). Is pragmatic inferencing applied at this point to pick among the possible antecedents? We assume this to be the case, otherwise it is not clear how centering will choose among the multiple potential referents. Yet, assuming this is the case, centering seems to involve more complex inferencing than our approach involves; inferencing for centering must pick which antecedent is better. RAFT/RAPR only asks inferencing to confirm a co-specification.<sup>19</sup>

---

<sup>19</sup> We do not check for ambiguity of pronouns in the same way that Sidner did. However, even if we were to incorporate similar checks for ambiguity, such checks would involve using inferencing to confirm two possible co-specifications for a pronoun, while centering might require the use of inferencing to confirm more than two. Thus, the inferencing required by RAFT/RAPR would still be more limited than that required by centering.

## 7. Conclusions

RAFT/RAPR and centering often make the same predictions for resolving pronouns. One reason for this is that the (abstract) preferences underlying the two frameworks share much in common. A major difference between the two approaches is that RAFT/RAPR keeps two foci, while centering keeps only one. In addition, RAFT/RAPR has more of a commitment to resolving pronouns according to grammatical role.

In Section 5, we argued that one is able to extend RAFT/RAPR to process sentences of the form "SX because SY," by specifying how to resolve pronouns and how to update RAFT/RAPR data structures, to match speaker preferences for discourses involving sentences of that form. On the other hand, centering could not be extended in a similar manner. Recent work (Section 5.3) suggests that centering might be able to match judgments by modifying *Rule 2*; however, the appropriateness of such a modification requires further study.

We presented examples of discourse-initial texts (Section 4.2) and a non-discourse-initial text (Section 4.3) involving only simple sentences that were problematic for centering. A second problem for centering results from not stacking focus information.

Although the claims in this paper are based on a limited number of discourses, we noted that it is difficult to perform a corpora analysis to determine which framework performs better on average because of the prevalence of complex sentences and the lack of work on complex sentences in either framework. Furthermore, determining how to process particular kinds of complex sentences is a crucial step toward enabling such a corpora analysis. We discussed the results of an analysis of one kind of complex sentence ("SX because SY"). On the basis of this analysis, we extended RAFT/RAPR to handle this kind of complex sentence. In addition, we showed how the judgments collected from native speakers of English as part of this analysis were difficult to explain within the current centering framework. Furthermore, the findings discussed in this paper indicate that one cannot evaluate the RAFT/RAPR and centering frameworks by simply applying the algorithms to finite clauses of naturally occurring text in linear order, without taking into account the occurrence of complex sentences and clue words.

### Acknowledgments

We thank Sandee Carberry, John A. Hughes, and Jeff Lidz for their comments and advice. We are also very grateful to the many people, too numerous to name here, who provided us with judgments. This research was supported in part by NSF Grant #IRI-9010112.

### References

- Brennan, Susan E. (1993). "Reference and local structure in interactive discourse." In *Proceedings, Workshop on Centering Theory in Naturally Occurring Discourse*, Institute for Research in Cognitive Science, University of Pennsylvania.
- Brennan, Susan E.; Friedman, Marilyn W.; and Pollard, Carl J. (1987). "A centering approach to pronouns." In *Proceedings, 25th Annual Meeting of the Association for Computational Linguistics*, 155–162.
- Brown, Cheryl (1983). "Topic continuity in written English narrative." In *Topic Continuity in Discourse*, edited by T. Givon, 317–341. J. Benjamins.
- Caramazza, Alfonso; Grober, Ellen; Garvey, Catherine; and Yates, Jack (1977). "Comprehension of anaphoric pronouns." *Journal of Verbal Learning and Verbal Behavior*, 16, 601–609.
- Carter, David (1987). *Interpreting Anaphors in Natural Language Texts*. John Wiley and Sons.
- Dahl, Deborah A. (1986). "Focusing and reference resolution in PUNDIT." In *Proceedings, 1986 National Conference on Artificial Intelligence*, Philadelphia, PA, 1083–1088.
- Grisham, John (1991). *The Firm*. Dell.
- Grosz, Barbara J. (1981). "Focusing and description in natural language dialogues." In *Elements of Discourse Understanding*, edited by A. Joshi,



- B. Webber, and I. Sag, 85–105. Cambridge University Press.
- Grosz, Barbara J., and Sidner, Candace L. (1986). "Attention, intentions, and the structure of discourse." *Computational Linguistics*, 12(3), 175–204.
- Grosz, Barbara J.; Joshi, Aravind K.; and Weinstein, Scott (1983). "Providing a unified account of definite noun phrases in discourse." In *Proceedings, 21st Annual Meeting of the Association for Computational Linguistics*, Cambridge, MA, 44–50.
- Hoffman, Beryl, and Turan, Ümit Deniz (1993). "Word order and centering in Turkish." In *Proceedings, Workshop on Centering Theory in Naturally Occurring Discourse*, Institute for Research in Cognitive Science, University of Pennsylvania.
- Kameyama, Megumi (1986). "A property-sharing constraint in centering." In *Proceedings, 24th Annual Meeting of the Association for Computational Linguistics*, Cambridge, MA, 200–206.
- Kameyama, Megumi (1993). "Intrasentential centering." In *Proceedings, Workshop on Centering Theory in Naturally Occurring Discourse*, Institute for Research in Cognitive Science, University of Pennsylvania.
- Kameyama, Megumi; Passonneau, Rebecca; and Poesio, Massimo (1993). "Temporal centering." In *Proceedings, 31st Annual Meeting of the Association for Computational Linguistics*, 70–77.
- Linson, Brian (1993). "A distributional analysis of discourse transitions." In *Proceedings, Workshop on Centering Theory in Naturally Occurring Discourse*, Institute for Research in Cognitive Science, University of Pennsylvania.
- Reichman, Rachel (1978). "Conversational coherency." *Cognitive Science*, 2, 283–327.
- Sidner, Candace L. (1979). *Towards a computational theory of definite anaphora comprehension in english discourse*. Doctoral dissertation, MIT.
- Suri, Linda Z. (1992). "Correcting illegal NP omissions using local focus." In *Proceedings, 30th Annual Meeting of the Association for Computational Linguistics (Student Session)*, 273–275.
- Suri, Linda Z. (1993). *Extending focusing frameworks to process complex sentences and to correct the written English of proficient signers of American Sign Language*. Doctoral dissertation, University of Delaware.
- Suri, Linda Z., and McCoy, Kathleen F. (1993). "Focusing and Pronoun Resolution in Particular Kinds of Complex Sentences." In *Proceedings, Workshop on Centering Theory in Naturally Occurring Discourse*, Institute for Research in Cognitive Science, University of Pennsylvania.
- Walker, Marilyn A. (1989). "Evaluating discourse processing algorithms." In *Proceedings, 27th Annual Meeting of the Association for Computational Linguistics*, 251–261.
- Walker, Marilyn A. (1993). "Initial contexts and shifting centers." In *Proceedings, Workshop on Centering Theory in Naturally Occurring Discourse*, Institute for Research in Cognitive Science, University of Pennsylvania.
- Walker, M.; Iida, M.; and Cote, S. (1992). "Japanese discourse and the process of centering." Technical Report 92-14, Institute for Research in Cognitive Science, The University of Pennsylvania.