

PRONOUN RESOLUTION BY THEMATIC ROLES

Bachelor project

Erick Wilts, s1635085, erickwilts@gmail.com

Abstract: This research examines focusing preferences for referents of pronouns following two types of verbs. The first type, transfer verbs (e.g., *pass*, *bring*), is a special case of the second type, experiential verbs (e.g., *scare*, *trust*). The frequency data from these corpus studies is compared with previous experimental research. Both studies confirm that with thematic roles (Goal and Source; Stimulus and Experiencer), continuations with pronouns are predictable. In experiential verbs, this is dependent on verb type (Stimulus-Experiencer vs. Experiencer-Stimulus), so there is an interaction between thematic roles and grammatical functions, which is consistent with Arnold (2001). For transfer verbs, there was an overall Goal bias. Guarantee of transfer (a distinction made by Hartshorne & Snedeker) did not influence the predictability of the continuation based on thematic roles. The object bias percentages per verb found by Hartshorne & Snedeker (*in press*) show a moderate correlation with the percentages from this study. The fact that this correlation is not higher could be caused by the difference in structure of the studies: a corpus study produces noisier data than a controlled experiment.

Keywords: Pronoun resolution; transfer verbs; experiential verbs; implicit causality; thematic roles

1. Introduction

If we want computers to understand or analyze natural language, pronoun resolution is a problem. Pronoun resolution is about finding out what a pronoun refers to, usually an entity introduced by a noun (phrase) or a name. But pronouns contain only a minimal amount of identifying information – usually only number and gender. The pronoun and its antecedent are co-referential (referring to the same object) with an object in the real world.

In sentence (1), the pronoun is *she*. The words *she* and *Mary* refer to the same object: the person Mary. We say that the anaphor *she* is co-referential with the antecedent *Mary*.

(1) Mary is happy because she finished a paper.

Problems in pronoun resolution occur mostly in interpersonal verbs (e.g., *amaze*, *throw*, *call*, *bring*), that describe events between multiple persons, because a pronoun could

refer to multiple entities. Introducing an ambiguous pronoun in a sentence is a useful way to examine the way humans resolve pronouns.

For example, in (2), the pronoun *he* clearly refers to Bill, because the ball is in possession of Bill, who is the only one who is able to throw it over the net. This contextual information is essential to understand the pronoun, but a computer does not have this world knowledge. Because John is masculine, *he* – a masculine pronoun – could just as well refer to John, which is the case in (3).

(2) John threw the ball to Bill. He lobbed it over the net.

(3) John threw the ball to Bill. He was amazed by his beautiful catch.

Three steps are necessary for a computer program to understand the pronoun *he*: first, the sentence should be parsed. This requires knowledge of grammar, words, names, irregularities in language, and still, with this knowledge, this is a hard task, because sentences are often ambiguous. Second, the consequence of the action described by the verb should be discovered. In case (2) and (3), the consequence of the first sentence is that John does not have the ball, and Bill probably has. I use the word *probably*, because throwing a ball does not guarantee a successful transfer (this is a distinction I will use later). Finally, all possible referents of the pronoun in the second sentence (in this case, *John* and *Bill*), should be tested with information about the consequences of the event in the first sentence. Now, all but one of the referents should be discarded as referents of the pronoun *he*.

Although this procedure would produce a robust solution to pronoun resolution, it would require lots of effort to put this into a program. This is difficult given the dynamic nature of human language. Therefore, the question arises whether there are other solutions. The rest of the introduction describes solutions provided by previous research, followed by a brief description of this research.

Au (1986) was one of the first to notice that humans use information implicit in verbs to find causes and consequences of events. She made subjects finish sentences starting with an interpersonal event and followed by the word *because* or *so*, to find causes and consequences of the event, respectively. An example of the sentences she provided for her subjects is given in (4).

(4) John feared Mary because ...

Au investigated two categories of verbs. Verbs from the first category are experiential (e.g., *thank*, *confess*, *punish*, *admire*) and describe how one person (the Experiencer) experiences stimuli of the other (the Stimulus). *Fear* is an example of an Experiencer-Stimulus verb, because the Experiencer (the subject of the sentence) occurs before the

Stimulus (the object of the sentence). In Stimulus-Experiencer verbs, the Stimulus is the object and the Experiencer is the subject. For example, the Stimulus-Experiencer equivalent of *fear* is *frighten*. Au took both Experiencer-Stimulus and Stimulus-Experiencer verbs into account. For these verbs, the Stimulus was judged as the cause of an event in 94% of the cases. If (4) would be one of the 94%, *because* would be followed by *she*; Mary would be the cause of John fearing her.

The second category of verbs consists of action verbs (e.g., *help*, *interrupt*, *congratulate*, *thank*), that describe actions of one person (the Agent) on another (the Patient). This category has also two subcategories: Action-Agent verbs (such as *hit*) and Action-Patient verbs (such as *correct*). The former describe intended or self-initiated actions, whereas the latter describe actions that respond to the situation.

In 84% of Action-Patient verbs, the Patient was judged as the cause of the sentence, but in Action-Agent verbs, no significant bias towards the Agent or the Patient as the cause was found.

The conclusion is that many verbs contain implicit information about the cause and the consequence of the event they describe.¹

This effect is later called ‘implicit causality’ by others (e.g., Arnold, 2001). Implicit causality, Arnold states, influences the interpretation of a pronoun that is followed by the use of a verb from a specific group. For both these verbs and for verbs followed by the word *because*, which marks an effect-cause discourse relation, we can predict the antecedent for a pronoun that follows the use of a verb. These constraints to the predictability of pronoun references show that implicit causality is not a sufficiently general solution for the problem of pronoun resolution.

A more general explanation might be that the subject (Arnold, *unpublished manuscript*) or first-mentioned noun phrase (McDonald & MacWhinney, 1995) is more accessible (i.e., more likely to be referred to).

Another explanation is that thematic roles (like Experiencer, Stimulus, Agent and Patient) influence the accessibility of referents, just as they influence causes and consequences (Au, 1986).

Arnold (2001) supports the latter explanation and focuses on transfer verbs (e.g., *sell*, *toss*, *kick*, *give*). Transfer verbs describe the transfer of an object from one person (the Source) to another (the Goal). In her experiment, participants were asked to read stories aloud and add a natural continuation to the story at the end. She gave two sentences to set the context and a third to introduce a verb with Goal and Source arguments, for example:

¹ Later in this paper, I will state that the connector words *because* and *so that* Au used influence this effect, but she used a test case and showed that these connectors did not significantly influence her results.

- (5) There was so much food for Thanksgiving, we didn't even eat half of it. Everyone got to take some food home. Lisa gave the leftover pie to Brendan.

Subjects were then asked to add a fourth sentence that continued the story. The only restriction in this method is that the continuation has to be a new sentence. This allows investigating what the new focus (first mentioned character) is, and whether it is referred to with a pronoun or a name. Arnold found an overall bias to Goal referents, which is stronger if the Goal is the subject of the sentence. This might, however, be a task-specific finding, because in this research, in contrast to previous research, the object is more referred to than the subject (McDonald & MacWhinney, 1995). Arnold also found that Goals were referred to by pronouns more often than Sources were, but that effect was most prominent when the Goal was the object.

She also looked at what type of continuation the sentence was: (1) specifying the cause, (2) discussing the endpoint (i.e. discussing the following conditions of a transfer event), or (3) another continuation. The causal continuation type yielded the most Goal continuations.

Because the conclusion of Arnold's first study (that the object was more referred to than the subject) deviated with conclusions from previous research, she conducted a corpus study. She sampled 174 sentences with experiential verbs from her first study from the Aligned-Hansard corpus, a collection of transcriptions of spoken text from the Canadian Parliament. 84 of those sentences contained Source-Goal verbs. To compare her study to my study, the numbers of samples are listed below:

- *give*: 22 samples
- *send*: 20 samples
- *teach*: 1 sample
- *offer*: 21 samples
- *pay*: 20 samples

The rest of the samples (90 sentences) contained Goal-Source verbs (*get*, *accept*, *receive*, *buy*, *take*, and *learn*). From this corpus study, Arnold made three conclusions.

First, subject referents were more often referred to than object referents. Therefore, she concluded that the results from her first study, that the object was more referred to than the subject, were indeed a task-specific findings.

Second, Goals were more often referred to than Sources. This was, again, mainly true if the Goal was the object of the sentence. This was consistent with her experimental findings.

Finally, there is an interaction between grammatical functions (subject, object) and thematic roles (Goal, Source): objects referred to first were far more often Goals than Sources, whereas when the subject was the new focus, the percentages of Goal references and Source references were almost equal.

Rohde, Kehler and Elman (2006) offer more experimental results that show the relationship between implicit causality and continuations with pronouns. These results seem to be more cognitively plausible, because they deal with context understanding.

They divided transfer verbs in 3 classes, distinguishing between co-locatedness (whether Goal and Source are in the same location given the sentence) and guarantee of transfer (whether the object transferred is guaranteed to arrive):

Table 1.1: Classes of transfer verbs (Rohde et al, 2006) with examples, categorized by co-locatedness and guarantee of transfer.

	Co-located	Not co-located
Guaranteed transfer	Class 1: <i>hand, give, deliver</i>	
No guaranteed transfer	Class 2: <i>throw, lob, fling</i>	Class 3: <i>mail, fax, ship</i>

It is not possible to think of a verb that describes a successful transfer between a Source and a Goal that are not co-located, so that class does not exist.

Rohde et al. (2006) also made a distinction in verb aspect:

(6) John handed a book to Bob. He...

(7) John was handing a book to Bob. He...

Sentence (6) has a perfective verb and describes a completed event. Sentence (7) has an imperfective verb and describes an ongoing process, so it cannot be continued by explaining the cause of the process.

They found that, in written sentences with ambiguous pronouns like (6) and (7), the imperfective yielded more Source than Goal resolutions for all three classes; this is also true for perfectives from class 3, but not significantly for classes 1 and 2. This insight leads them to the conclusion that pronoun resolution is dependent on event structure, i.e. how the sentences are related to each other.

To show this is also true for perfectives, they further investigated their participant's continuations. Following Arnold (2001), who categorized the coherence relations into 3 classes ((1) specifying the cause, (2) discussing endpoint and (3) other), Rohde et al. categorize them into 5 classes (here sorted by category size, greatest first).

- Occasion: the second sentence describes an event that is possible with the consequence of the event described in the first sentence.
- Elaboration: the second sentence describes what other event happened than described in the first sentence.
- Explanation: the second sentence explains the event in the first sentence.
- Result: the second sentence describes the result of the first sentence.

- Parallel: the second sentence describes what else happened to the entities mentioned in the first sentence.

In Occasion continuations, Goals were most often referred to first. In Elaborations and Explanations, Sources were more salient. For the other two categories (containing 24 and 8 sentences, respectively), no significant results were found.

Given these results, they concluded that pronoun interpretation is dependent on event structure and that bias caused by thematic roles is a side effect.

As said before, implicit causality is information inherent in the meaning of particular verbs. Rohde et al. (2006) seem to have clearly shown that its effects are mitigated by depending on the coherence relations they occur with, but still the term is used in many research fields (Hartshorne & Snedeker, *in press*). The reason for this, they state, is that it links the linguistic process of pronoun resolution to knowledge about the world. Two existing theories about the mechanisms of implicit causality are what Hartshorne & Snedeker call the *world knowledge* account (Au, 1986) and the *arbitrary semantic tag* account (Garvey & Caramazza, 1974).

The world knowledge theory states that knowledge about the causal structure of the world impacts pronoun resolution according to the same causal structure in verbs that is used to describe the world. According to this theory, implicit causality bias is in itself not a linguistic phenomenon, but a result of world knowledge projected on individual verbs.

In contrast, the arbitrary semantic tag theory states it is not possible to classify verbs according to the direction of the implicit causality bias. Rather, each verb has a bias, but there are no rules according to which this bias is distributed.

Hartshorne & Snedeker prefer a theory that is somewhere in between these two: the *semantic structure* theory. The world knowledge theory provides too few verb classes (and therefore classes with too much variation) to extract significant theories for these classes, whereas the arbitrary semantic tag theory provides too many (namely, each verb in its own class). The semantic structure theory divides 3769 verbs into 274 classes, using VerbNet (Kipper et al., 2009), “the most comprehensive semantic classification scheme”, according to Hartshorne & Snedeker, but also very similar to other classification proposals. Verb classes vary in:

- Types of grammatical arguments they can take (subject, direct object, indirect object);
- Whether the order of grammatical arguments can be changed;
- The way in which the verb influences the object (e.g., *turn x to y* changes x into y, in contrast to *convert x to y*, which makes x an instance of y)

Hartshorne & Snedeker investigated the 720 most frequently used interpersonal verbs in English. They found four semantic classes of experiential verbs (in total 63 verbs) with a significant bias toward the object or the subject:

- Class 45.4 (verbs like *accelerate, reverse, shrink*) and 31.1 (verbs like *affect, bore, trouble*) from VerbNet show a bias toward the subject;
- Class 31.2 (verbs like *admire, bear, love*) and 33 (verbs like *congratulate, curse, greet*) show a bias toward the object.

They concluded that, with moderate-sized verb classes, indeed predictions can be made about the new focus.

This paper combines the aforementioned research by doing two investigations, both corpus studies. The first study is about transfer verbs: I will use class 1 and class 2 verbs from Rohde et al. (2006) (see table 1.1) and test whether they have Goal or Source continuations. This is an expansion of the corpus study by Arnold (2001), investigating 602 sentences of Source-Goal verbs (where Arnold had 84 sentences). Moreover, it uses sentences with both dative alterations (dative = indirect object), e.g., with a prepositional phrase (PP) or with an indirect object. In English, sentences with an indirect object can be formed in two ways. In (6) and (7), Bill is the indirect object of the sentence. Sentence (7) is the altered version of sentence (6).

(8) John threw the ball to Bill.

(9) John threw Bill the ball.

In a way, this study fine-tunes Arnold's (2001) research: I only investigate Source-Goal verbs (e.g. *give*, but not *receive*). I also limited my samples to sentences that describe events of physical transfer (so, *give* is allowed whereas *teach* is not).

This study also builds on the research of Rohde et al. (2006) by using their verb classes. It attempts to reproduce their findings in a corpus study.

The second study is about experiential verbs: I will investigate what circumstances (Stimulus-Experiencer vs. Experiencer-Stimulus verbs; connecting conjunctive) influence a continuation with a Stimulus or an Experiencer. This study builds on Au (1986), who argued that verb types influence the cause of a sentence; this study attempts to extrapolate her findings about causes and consequences to pronoun resolution.

Also, this study attempts to reproduce Hartshorne & Snedeker's (2006) conclusions by comparing object bias percentages. Hereby, an answer to the following research question is to be found:

Do we find frequency data that correlate with the experimental results showing focusing preferences for transfer verbs and experiential verbs?

2. Methodology

I used the Corpus of Contemporary American English (COCA; Davies, 2008) to collect data. It is an online, freely available, corpus that allows searching for words and inflexions in fiction, magazines, newspapers, academic texts and transcriptions of spoken text from between 1990 and 2012. COCA is a balanced corpus, i.e. the amount of words per genre and per year is approximately equal: about 4 million words per year per genre.

I used COCA because I needed many examples of what are often rather infrequent verbs, in a specific syntactic construction. COCA contains 450 million words, and as such is the largest freely available annotated corpus of English. COCA is not fully syntactically parsed, which made searching sometimes difficult. However, existing parsed corpora are much smaller. In COCA, I found 602 sentences with transfer verbs and 255 sentences with experiential verbs.

2.1. Study 1: transfer verbs

For my first study, I used search strings like the ones in Figure 2.1. The parts of the upper search strings look for the following sentence fragments, consecutively:

- [throw].[v*] – [throw] represents any form of the word *throw* (e.g. *throwing*, *throws*), . [v*] means that it has to be a form of the verb *throw* (and not a form of the noun *throw*)
- * – Any word (in most cases, this was *a*, *the* or a personal pronoun)
- [nn*] – Any noun
- to – The word *to*

This search string finds results like “throwing a stone to” and “throws a party to”, with their context. The latter example is discarded, because the word *to* does not denote a transfer, but means ‘in order to’. I saved a small piece of context to show what the Source and Goal are. I also saved the type of text the verbs come from, whether it was continued with *who* (in other cases, I called it a ‘regular’ continuation), and the new focus of the succeeding discourse (Source or Goal).

The second string is the dative alteration² of the first: it searches for a form of the verb *throw* that is followed by a personal pronoun. Then, *a|the* searches for the word *a* or the word *the*. The last part searches for any noun.

To find more sentences, I searched for longer noun phrases placing one or more extra asterisks (*) in the search string.

² See section 1.5 for an explanation of dative alteration.


```
[throw].[v*] * [nn*] to
[throw].[v*] [p*] a|the [nn*]
```

Figure 2.1: Two search strings for COCA used for this corpus research to find sentences with a transfer verb, a Source and a Goal. The text above explains how they work and what they find.

2.2. Study 2: experiential verbs

The string in Figure 2.2 is used to find Stimulus-Experiencer verbs and searches for sentences that contain the following consecutive words:

- [blame].[v*] – Any form of the verb *blame*
- * – Any word (mostly a reference to a person)
- and|so|because – Any one of the words {and, so, because}
- [p*] – Any personal pronoun

To find more sentences, I searched for longer noun phrases placing one or more extra asterisks (*) in the search string.

```
[blame].[v*] * and|so|because [p*]
```

Figure 2.2: A search string for COCA used to find sentences with an experiential verb, an Experiencer and a Stimulus.

3. Results

3.1. Study 1: Transfer verbs

In my study, the verb aspect (perfective and perfective) was not taken in account, contrary to Rohde et al. (2006). I did compare two of the verb classes (classes 1 and 2, see Table 1.1) they distinguished between.

These classes, as shown in Figure 3.1, do not differ from each other significantly ($\chi^2 = 2.528$; $df = 1$; $p = 0.112$; $\phi = 0.004$). Because this research does not take verb aspect into account, I cannot compare these results with those of Rohde et al. (2006).

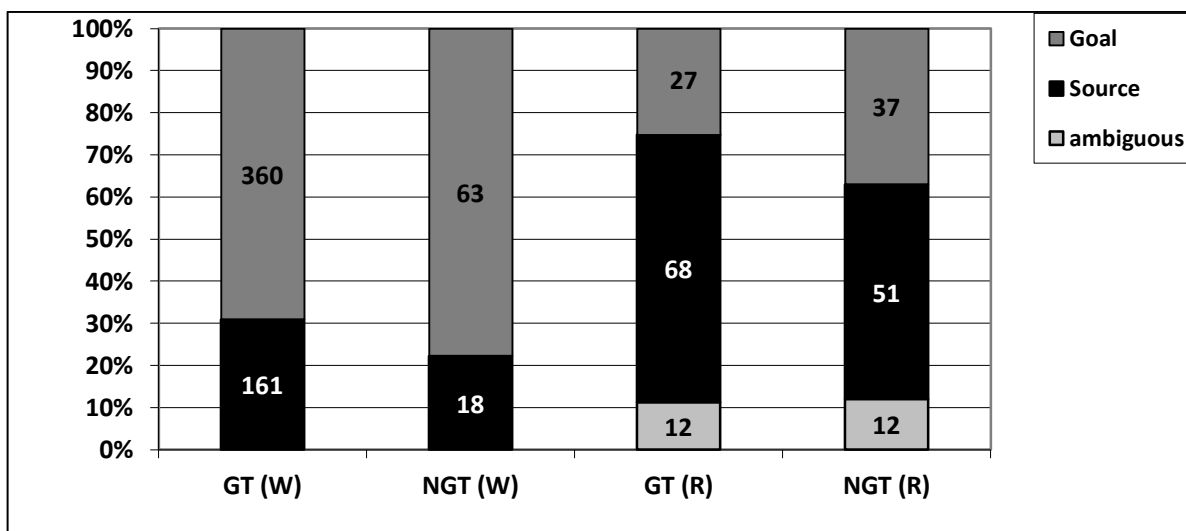


Figure 3.1: In this figure, class 1 and class 2 from Rohde et al. (2006) (see also Table 1.1) are compared, not taking verb aspect into account. (GT = Guaranteed Transfer (class 1); NGT = No Guaranteed Transfer (class 2); W = Wilts (this research); R = Rohde et al. (2006)) The bars show the number of cases and the percentages in which the new focus was the Goal, the Source or ambiguous.

Hartshorne & Snedeker (*in press*) calculated an ‘object bias’ percentage (the percentage in which the object was the focus of the discourse after the verb) for 720 verbs. I summarized the data for the experiential verbs in the same way Hartshorne & Snedeker did in their Appendix A. Because the Goal is always the object in transfer verbs used in this work, and the Goal was already determined for each sample, this was easily done. The verbs that were investigated in both studies were the following: *give, hand, offer, pass, sell, kick, roll, throw* and *toss*. The results are summarized in Figure 3.2.

I used a variant of the R^2 method to find the correlation between the data I found and the data from Hartshorne & Snedeker. R^2 uses the distance from data points to a regression line (a line that best fits the data) and the difference from the mean to calculate an R^2 value, which always lies between 0 and 1. If $R^2 = 1$, the fit is perfect. This means all found data points fit the expected regression line; however, it does not mean that there is causality between the data and the expected values.

My data were supposed to fit the line $y = x$ in figure 3.2 instead of the regression line (the line that best matches the data)³. I used the distance from the points to the line $y = x$, so I can still use the R^2 method. The result of this comparison was $R^2 = 0.557$, so the correlation is moderate.

³ Where x stands for the object bias percentages Hartshorne & Snedeker found, and y for the object bias percentages from this research. I use this line because the ideal case would be $\forall(x,y) y = ax + b$ with $a = b = 0$. The regression line (best fitting line to the data) has the approximate values $a = 78$ and $b = -0.14$.

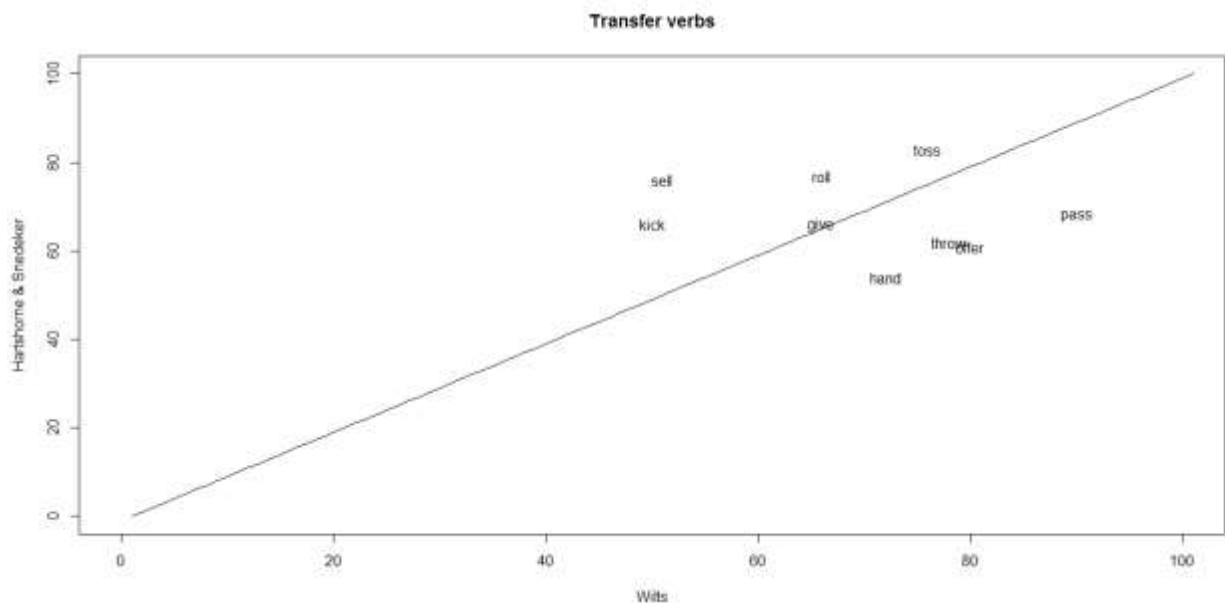


Figure 3.2: This figure shows the similarity between the findings of Hartshorne & Snedeker (*in press*) (y-axis) and my findings for transfer verbs (x-axis). If my findings would be exactly the same as those of Hartshorne & Snedeker, all words would lie on the line, because the found object bias would be the same in Hartshorne & Snedeker's research. The R^2 -correlation between the two researches is 0,557. See section 3.1 for a description about interpreting R^2 .

3.2. Study 2: Experiential verbs

The reference types for the two types of experiential verbs I investigated are shown in figure 3.3.




Figure 3.3: This figure shows how often a certain thematic role was the new focus of the discourse after an Experiencer-Stimulus (ES) verb or a Stimulus-Experiencer (SE) verb in my research (W) or Au's research (A). In my research, the two classes differ significantly ($\chi^2 = 12.224$; $df = 1$; $p < 0.001$; $\phi = 0.048$), so the thematic role of the new focus is dependent on the

verb class. In the bars summarizing Au's research, percentages instead of numbers are shown, because Au did not provide numbers.

I can also compare my results for experiential verbs with the results of Hartshorne & Snedeker (in press), just as I did with the transfer verbs. The words that occurred in both studies were: *admire, detest, dislike, fear, hate, trust, amaze, amuse, blame, bore, confess, criticize, disappoint, enrage, envy, frighten, inspire, mesmerize, phone, pity, praise and scare*. The result is shown in Figure 3.4.

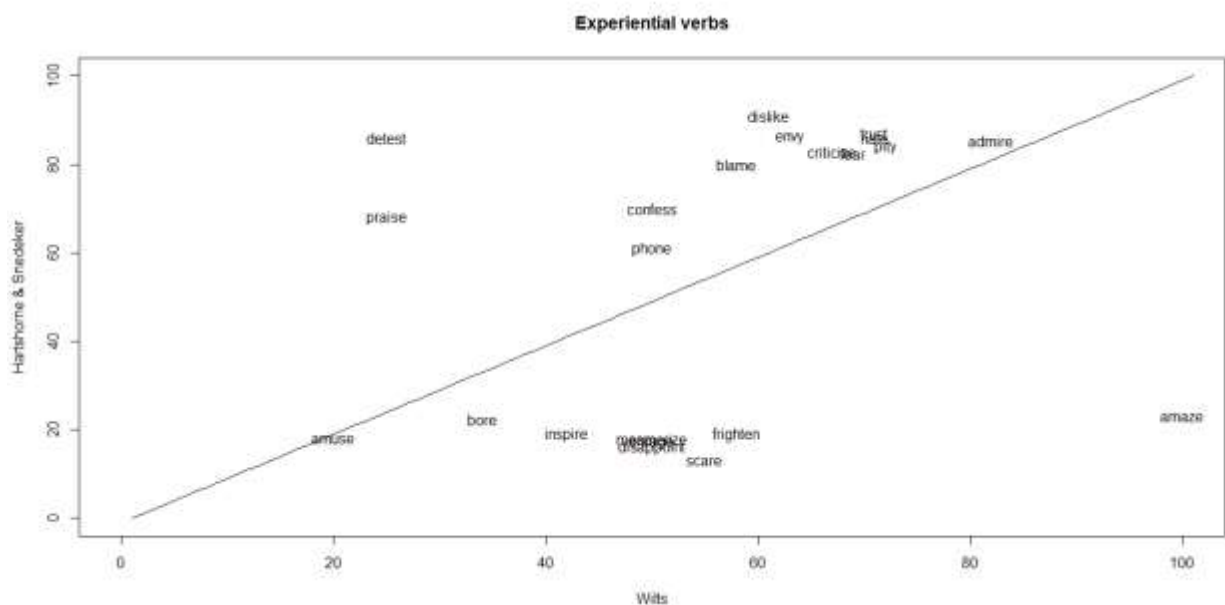


Figure 3.4: This figure shows the similarity between the findings of Hartshorne & Snedeker (*in press*) (y-axis) and my findings for experiential verbs (x-axis). If my findings would be exactly the same as those of Hartshorne & Snedeker, all words would lie on the line, because the found object bias would be the same in Hartshorne & Snedeker's research. The R^2 -correlation between the two researches is 0.291 (low correlation). After removing outliers of with 7 or less samples were collected (*detest, amaze, frighten* and *mesmerize*), R^2 was 0.479 (moderate correlation). See section 3.1 for a description about interpreting R^2 .

4. Discussion

To give an answer to the research question, it is necessary to compare the results of this research with the research of Au (1986), Arnold (2001), Rohde et al. (2006) and Hartshorne & Snedeker (*in press*). What were their conclusions and do they match the results of this corpus research?

One of the conclusions of Au (1986) was that in 94% of the cases, the Stimulus was the cause of the new sentence. I did not investigate the cause of a sentence, but the new

focus. Au made participants finish sentences with the connectors *because*, *so* and sentences without connector. The cause of a sentence, however, is mostly the same as the new focus – I defined the new focus as the first-mentioned thematic role, whereas Au examined which role the sentence was mainly about. One would expect to find a high correlation between a new focus and a cause preceded by *because*, and hence a similar percentage. Au found 94% Stimulus bias. Figure 3.3 supports Au's conclusion that the Stimulus is more often referred to in Experiencer-Stimulus verbs. Because Au did not provide numbers but only percentages, I could not perform a χ^2 test, but this graph shows a slight correlation. Au's numbers are more extreme (closer to 0% and 100%), because her data (consisting of sentences that all have the same structure) is less continuous than my corpus data. If this difference in research is counted in, Au's and my conclusions are very similar.

Arnold (2001) found that Goals were more accessible than Sources, especially when the Goal was the object. Although I did not examine grammatical roles, I did find an overall Goal bias (see Figure 3.1). However, this study only investigates Source-Goal verbs, not Goal-Source verbs like *receive*, so the Goal was the object in most of the cases. So for this claim, my results concur with Arnold (2001).

Figure 3.1 shows the comparison between class 1 (Guaranteed Transfer) and class 2 (No Guaranteed Transfer) verbs (see also Table 1.1), where Rohde et al. (2006) found a difference in guarantee of transfer combined with verb aspect; this effect is nowhere to be found in my data in which verb aspect is not taken into account. The fact that there is no difference between the classes overall could have various reasons. First, the methodology of this work is different. Rohde et al. did not find significant focusing preferences for perfectives in their first study, and my corpus study may have yielded many sentences in perfective. Second, if there were differences caused by verb aspect, they might have been averaged out.

Finally, the results of Hartshorne & Snedeker (*in press*) showed a moderate correlation with my results. In both graphs comparing my results with their results (Figures 3.2 and 3.4), one can easily see a bimodality in Hartshorne & Snedeker's results – verbs cluster around 20% and 80% object bias, in contrast to my results. This bimodal distribution could be an effect of their research method, in which participants were asked to identify the *dax* in sentences with the following structure: "Mary (VERBed/was VERBing) Sally because she is a dax". My prediction was that most participants would agree on the same *dax* for each verb, because the sentence structure is always the same. This prediction is confirmed by the clusters around 20% and 80%. In a corpus study, the data is noisier and gives a more continuous range of percentages, due to other factors playing a role, including current topic, context etc., which are factored out in a controlled study of the type Hartshorne & Snedeker did. Even though very good conclusions about the

object bias can be drawn from a bimodal distribution, this line of reasoning does explain the difference between their study and my corpus studies, and hence the low R^2 values.

Arnold (2001) did not use the word *because* as a connector between the verb-containing part of the sentence and the continuation of the participants. The word *because* denotes the reason for the described event to happen and the verb describes the event.⁴ So, if the word *because* would occur in a sentence, the subject of the verb would be more often referred to than with other connectors. My research contains continuations with the connectors *and*, *because* and *so*. Figure 4.1 shows that the focusing preferences caused by the connectors indeed differ significantly ($\chi^2 = 7.203$; $df = 2$; $p < 0.05$; $\phi = 0.028$). Although the *because* connector did not differ significantly from the rest of the data ($\chi^2 = 3.071$; $df = 1$; $p = 0.080$; $\phi = 0.012$), Arnold was right to not use connector words.




Figure 4.1. This figure shows the number of Experimenter and Stimulus continuations with different kinds of connector words. A *because* connector leads to an increasing likelihood that the Stimulus is the new focus of the sentence, because it influences the probability of a subject being the cause of the sentence.

The words *because* and *so* remind one of the coherence relations Explanation and Result from Rohde et al. (2006). For Explanation, they found a significant focusing preference toward Sources. For Result relations, the preference was toward Goals, but most of the sentences were of the form *X transfers Y to Z. Z thanks X*. This is why they made no extrapolated conclusion about this type. My data does not show a similar pattern. For the connector *so*, there is a significant focusing preference ($\chi^2 = 5.259$; $df = 1$;

⁴ This was exactly what Au (1986) wanted to investigate (*because* leading to the cause of the sentence, and *so* to the consequence of the sentence), but Arnold wanted a more general explanation, without the use of connector words (Au only omitted the connector word as a control condition).

$p < 0.05$; $\phi = 0.020$); for the connector *because*, there was no significant preference ($\chi^2 = 3.071$; $df = 1$; $p = 0.080$; $\phi = 0.012$).

With the data I collected, a wide variety of other studies could be done. To answer my research question for each aspect of the used research, the grammatical roles should be investigated instead of only the thematic roles. Moreover, the continuation type should be analyzed, just as Rohde et al. (2006) did. But even without examining these extra features, this work is still a promising step towards the conclusion that corpus data and experimental data complement each other in the study of pronoun resolution. Where structural experiments show focusing preferences based upon grammatical functions, thematic roles or discourse coherence, a corpus study can confirm or reject these findings with noisier (that is, more continuous) data. A corpus study can make conclusions about 'real world' text. This helps to satisfy the goal we stated in the introduction, that a computer understands or analyzes natural language.

Given the diverging conclusions about what the underlying mechanisms of implicit causality are and the fact that implicit causality is not a general solution, a more refined theory must be formulated about the prediction of pronoun references. Whether the focus on a certain grammatical role is dependent on verb (type) or event structure (or maybe a combination of those) is still unclear. But from the verbs that have implicit causality, some conclusions can be drawn in terms of grammatical roles. Given that some verbs have biases percentages close to 50% toward a certain thematic role, it might not be useful for predicting what a certain pronoun refers to, but this percentage (in combination with other factors) can be taken into account when estimating the likelihood of a pronoun referring to a certain referent.

For further research, I would recommend taking grammatical annotations as well as coherence relations into account, so all aspects of mentioned research can be verified. For a broader study, the number of verb types could be expanded. Class 3 verbs from Rohde et al. (describing not co-located events) and Goal-Source verbs (where the Goal is the subject and the Source is the object) are examples of such verbs.

5. References

Arnold, J.E. (*unpublished manuscript*). Marking salience: The similarity of topic and focus.

Arnold, J.E. (2001). The effects of thematic roles on pronoun use and frequency of reference continuation. *Discourse Processes*, 31(2), 137-162.

Au, T. K. (1986). A verb is worth a thousand words: the causes and consequences of interpersonal events implicit in language, *Journal of Memory and Language*, 25(1), 104-122.

Davies, M. (2008-) *The Corpus of Contemporary American English: 425 million words, 1990-present*. Available online at <http://corpus.byu.edu/coca/>.

Garvey, C. & Caramazza, A. (1974). Implicit causality in verbs. *Linguistic Inquiry*, 5, 459-464.

Hartshorne, J. K. & Snedeker, J. (*in press*). What is implicit causality: world knowledge, an arbitrary feature or an effect of semantic structure?

Kipper Schuler, K., Korhonen, A. & Brown, S. W. (2009). *VerbNet overview, extensions, mappings and apps*, Tutorial, NAACL-HLT.

McDonald, J. L. & MacWhinney, B. J. (1995). The time course of anaphor resolution: Effects of implicit verb causality and gender. *Journal of Memory and Language*, 34, 543-566.

Rohde, H., Kehler, A. & Elman, J. L. (2006). Event structure and discourse coherence bias in pronoun interpretation, *Proceedings of the 28th Annual Conference of the Cognitive Science Society*.