Names and pops and discourse structure

Nicholas Asher, Pascal Denis, Brian Reese

Names and Pops and Discourse Structure

Nicholas Asher
CNRS-IRIT, Toulouse
asher@irit.fr

Pascal Denis
University of Texas at Austin
denis@mail.utexas.edu

Brian Reese
University of Texas at Austin
bjreese@mail.utexas.edu

June 22, 2006

1 Introduction

Anyone who has studied anaphora has surely wondered why natural languages have so many types of expression which can be used anaphorically. Work in the philosophy of language on proper names and definite descriptions has not addressed this question, but linguists such as Prince (1981),Ariel (1988), and Gundel et al. (1993) have, hypothesizing a hierarchy of referential expressions correlating linguistic form to the level of salience of the entity referred to.[1] One such *Referential Hierarchy* is provided below.

Definition 1 (Referential Hierarchy)**.** $\epsilon_1 >> \epsilon_2$ is glossed as "Referring expression ϵ_1 requires a more salient antecedent than expression ϵ_2":

> Zero pronouns >> Pronouns >> Demonstrative NPs >> Definite descriptions >> Short PNs >> Full PNs >> Full PNs + appositive

A model of *salience* such as that provided by Centering Theory gives us a most salient antecedent, but to see what the observations of the Referential Hierarchy (RH) amount to, we need a theory of salience that establishes a more complete ordering on the entities in the discourse context. Various probabilistic algorithms (cf. Mitkov 1994) might be used to provide a more extensive ordering, but they fail to be dynamic in the right way. The use of certain referential expressions not only has an update effect on the set of salient expressions; they render discourse entities that were not salient, salient. Furthermore, as Asher (2006) argues, this effect is absolute and not gradual. Probabilistic models of salience are all gradualist, but the data show sharp divisions that such gradualistic accounts cannot capture.

The above considerations raise a number of interesting questions, each of which revolves around whether there is a correlation between the uses of different types of referential expressions and the attachment of new information to the existing discourse structure. Ariel (1988), for example, estimates salience, or accessibility in her terminology, by measuring the distance between a referring expression and its antecedent. Low-salience referring expressions like definite descriptions and proper names, she finds, tend to find their antecedents more distantly than high-salience referring expressions. These findings suggest that definite descriptions and proper names may be good linguistic cues to the presence of a discourse pop, whereas the use of a pronoun signals a more local discourse attachment. The two questions that we address below are as follows. First, what types of referential expressions are used when a discourse pop occurs with a coreferential expression in the current discourse constituent? The analyses of Ariel (1988) and Gundel et al. (1993), for

[1] Prince (1981) includes indefinites in the hierarchy, which introduce novel discourse elements; they're a limiting case of the idea here.

instance, suggest that low salience referring expressions should be preferred, i.e., proper names, definite descriptions, etc. Second, what is the relationship between transitions in the referential hierarchy and the occurrence or non-occurrence of discourse pops? For example, is a shift in the referential hierarchy from a high-salience to a low-salience expression likely to be associated with a discourse pop? Conversely, the use of a high-salience referential expression such as a pronoun might be expected to indicate local discourse attachment. The example below illustrates this idea.

(1) a. That view was supported in a tender-offer document filed by USAir with the Securities and Exchange Commission and in another document filed in court. (π_{23})
...
b. "He doesn't give up easily (π_{47})
c. and one should never underestimate what he can or will do." (π_{48})
d. Mr. Icahn couldn't be reached for comment. (π_{49})

A model of text cohesion such as that proposed in Reinhart (1980) would predict a preference to attach π_{49} to π_{48}, since the pronouns in π_{47} and π_{48} are coreferential with the subject of π_{49}. Furthermore, the models of salience proposed by Ariel (1988) and Gundel et al. (1993) would predict the use of a high-salience referring expression such as a pronoun. We don't see this however. Rather, π_{49} uses a low-salience referring expression, i.e. the proper name *Mr. Icahn*, and attaches high in the discourse structure, viz. to the segment π_{23} The shift from a referring expression requiring a highly salient antecedent, i.e., a pronoun, to one that requires low salience, i.e., a short proper name, coincides in this case with a significant discourse pop.

2 Empirical Findings

In order to address the questions raised in the introduction, we examined a small corpus based on 25 WSJ stories from the MUC6 data. These documents, which come with coreference annotations, were enriched with discourse segmentation and SDRT-type discourse structures. In SDRT, a discourse structure can be represented as a graph, with a distinguished node $LAST$ for the last constituent incorporated into the discourse structure (Asher and Lascarides, 2003). All attachments to available nodes in the graph that are superordinate to $LAST$ are considered discourse pops. Since the relation of superordination is linear with respect to the set of available attachment sites in SDRT, we can measure the distance of the pop by simply considering the linear order on available attachment sites. Thus SDRT provides a more sophisticated model of linguistic salience then the simple distance measurements proposed by Ariel (1988).

The corpus contains a total of 415 coreferential (non-singleton) chains, which altogether contain 1783 NP mentions, which means that there are in effect 1368 anaphoric NPs. On the discourse side, the corpus contains 844 discourse segments (or Elementary Discourse Units, EDUs).

2.1 Details of the corpus

Table 1 summarizes the distribution of the different types of anaphoric NP mention found in our corpus. The different referential types were obtained automatically based on the NP string and the parts of speech output for the NP by a POS tagger. The category *other* labels NPs that do not fall into the types present in the RH; they include indefinite NPs, bare nouns, etc.[2] Interestingly, Table 1 indicates that there are more referential uses of names (more than 33%) and definites (almost 27%), than of pronouns (only 14.5%).

Before we look at discourse structure and discourse pops, it is instructive to consider the *distance* between an anaphor and its antecedent. Locality is indeed often taken to be a good approximation for salience, and authors like Ariel (1988) have suggested a strong correlation between the RH and (sentence) distance. Table 2 shows the distribution of anaphoric NP types in terms the average distance to the closest EDU containing an antecedent.

[2]For simplicity, we defined *Short PNs* as PNs containing a single token; all other cases are either *full PNs* or *Full PNs+appositive*.

Referential type	Count	Ratio
Definite NPs	364	0.266
Short PNs	289	0.211
Pronouns	198	0.145
Full PNs	140	0.102
Full PNs+appositive	32	0.023
Demonstrative NPs	11	0.008
Other	334	0.244

Table 1: Distribution of anaphoric NPs wrt referential types

Referential type	EDU distance
Pronouns	1.101
Demonstrative NPs	2.091
Short PNs	3.356
Definite NPs	4.874
Full PNs	5.734
Other	5.927
Full PNs+appositive	7.656

Table 2: Average EDU distance between an anaphor and its closest antecedent

The figures in Table 2 conform closely to the RH, showing that pronouns and demonstratives have in general shorter distances to their antecedents than do low-salience referential expressions. However, there is one striking divergence between the ordering in the RH and the data on EDU distance between anaphoric expressions and their antecedents: proper names do not behave in a uniform fashion. Roughly, the shorter the proper name, the closer the antecedent in terms of EDUs. In particular short proper names (*Short PNs*) take closer antecedents than definite NPs in our corpus, whereas the RH predicts that all proper names should take more remote antecedents on average than definite NPs for instance.

Turning to a discussion of discourse pops, the corpus contains a total of 324 discourse pops occur in the 25 documents that comprise our corpus. This means that pops occur roughly at a rate of 4 pops every 10 attachments. Not surprinsingly, the frequency of discourse pops decreases proportionally to the length of the pop (i.e., the number of intervening EDUs): thus, 38% of these pops were pops to the penultimate EDU, and almost 75% with less than 5 intervening EDUs. However, we had pops that jumped back as many as 86 EDUs.

Finally, various types of transitions on the referential hierarchy occur in the corpus. We distinguish between three types of transitions between an anaphoric expression and its linguistic antecedent with respect to the RH: (i) a *jump* is a rightward shift on the RH (e.g. from a pronoun to a PN), a *plunge* is a leftward shift on the RH (e.g. from a full PN to a short PN), and a *stay* corresponds to the situation where there is no shift in the RH (e.g., from pronoun to pronoun).[3] It is worth noting here that we only consider transitions within coreferential chains: that is, transitions between the anaphor and its closest antecedent. The number and frequency of the various referential transitions in our corpus are as follows:

[3]The additional *unknown* transition in the table above corresponds to cases where the referential type of one of the NP, either the anaphor or the antecedent, was "other".

Referential transition	Count	Ratio
Stay	453	0.331
Plunge	240	0.175
Jump	187	0.137
Unknown	483	0.353

Table 3: Distribution of the referential transition

2.2 Types of referential expressions in pops

The first question we want to address concerns the type of referring expresions that are found in discourse pops. Based on the work of Ariel (1988) and Gundel et al. (1993), the default hypothesis would seem to be that low salience referring expressions should be preferred in discourse pops. Roughly, the rationale is that the antecedent of an anaphor occurring in a pop is no longer salient.

The data in our corpus partially support this hypothesis, as shown in Table 4. Note that 40% of coreferential pops involve the use of a proper name of some sort, whereas only 25% involve definite descriptions and only 10% involve pronouns. A greater appreciation of these figures is obtained by comparing them to the overall scores reported in Table 1. The comparison indicates an increase in the use of proper names, and a decrease in the use of the other expressions. We conclude from these observations that names are a good signal of a pop, whereas demonstratives, definites and pronouns are more likely to suggest local attachments.

Status on RH	Count	Ratio
Short PNs	128	0.253
Definite NPs	126	0.25
Other	119	0.236
Full PNs	61	0.121
Pronouns	54	0.107
Full PNs+appositive	13	0.026
Demonstrative NPs	4	0.008

Table 4: Pops and the discourse new/old referents

2.3 Referential transitions in pops

The second question we address concerns the relationship between discourse pops and referential *transitions*. From the perspective of building discourse structures, the question can be recast as follows: are shifts in the referential hierarchy a good indication of a discourse pop or a local continuation? One natural hypothesis to explore is that a shift from a high-salience referring expression in $LAST$ to a low-salience anaphoric expression on the Referential Hierarchy, i.e. a Jump, is a good indication of a discourse pop. Tables 5 and 6 summarize our findings on these issues.

Table 5 provides the proportion of each type of referential transition occuring in a discourse pop to the total number of pops. Conversely, Table 6 provides the proportion of each type of referential transition occuring in a discourse pop to the type of transition. The former shows that Jumps, while only representing less than 14% of all transitions overall (cf. Table 3), make up more than 18% of the total pops. The latter in turn shows that almost 50% of the Jumps appear in discourse pops. These facts suggest that the presence of a Jump is a rather strong signal for the presence of a discourse pop. When looking at particular transitions, we find that transitions from

Transition type	# trans.& pops	#(trans&pop)/ # pops
Stay	171	0.339
Jump	93	0.184
Plunge	69	0.137
Unk	172	0.341

Table 5: Distribution of coreferential pops wrt RH transitions

Transition type	# trans.& pops	#(trans&pop)/ # trans
Jump	93	0.497
Stay	171	0.377
Unk	172	0.356
Plunge	69	0.288

Table 6: Distribution of RH transitions wrt to pops

a pronoun to a(ny sort of) PN (i.e., the largest jumps) are the most reliable for signalling a pop: they co-occur with pops at 60%-68% (the largest number is for full PNs).[4]

3 Formal Reflections

We presented two measures suggesting that the coreferential use of names in discourse is suggestive of a discrouse pop. Several measures also suggest a distinction between, short names and longer names, as well as between names and definite descriptions and pronouns.[5]

We now suggest a formalization for some of these observations involving names, building on Hunter and Asher (2005) who propose a formalism within which to treat the preferential tendencies of presuppositions. They argued that indexicals had presuppositions that preferred highest possible attachment. In particular they looked at the behavior of *actual*. Formally, *actual* introduces an operator $\uparrow$ over material in its scope that affects the resolvability, $|\vdash$, of a presupposition, where this includes the pragmatic constraints on attachment of SDRT (Asher and Lascarides, 2003). To define $|\vdash$, they simplify binding to a notion of DRS satisfaction $\models$ and accommodation as incorporation of a presupposition ϕ into a DRS $K\rhd\phi$. The interpretation of a subDRS K depends on assignments to discourse referents declared in superordinate DRSs but free in K. Let $K_0, \ldots, K_n$ be a sequence beginning with the global DRS K_0 such that $K_0 \geq \ldots \geq K_n$, and $\geq$ is the immediate superordination relation on DRSs. Then, $K_0, \ldots K_{l-1}, [K_l, \ldots K_n] \models \phi$ iff ϕ is a dynamic consequence of $K_l, \ldots K_n$ relative to any assignment to free variables occurring in $K_l, \ldots K_n, \phi$ that are declared in $K_0, \ldots K_{l-1}$ and satisfy the conditions in $K_0, \ldots K_{l-1}$. Then:

- $K_0, \ldots, K_i|\vdash\phi$ iff $\exists j \leq i$ and $\exists l \geq 0$ such that $K_0, \ldots K_{l-1}, [K_l, \ldots K_j] \models \phi$ or for some k, $0 \leq k \leq j, K_k\rhd\phi$, for ϕ a normal DRS or DRS condition.
- $K_0, \ldots, K_i|\vdash \uparrow \phi$ iff there is some $j \leq i$ such that $K_0, \ldots K_j|\vdash\phi$ and there is no $k < j$ such that $K_0, \ldots, K_k|\vdash\phi$

[4]We need to make one caveat about PNs and jumps. Newspaper articles like those of the MUC corpora tend to contain short PNs, and the main topic constituent (typically, the first sentence in the text) often contains a full PN. Here of course there is no pop, but we take this exception to our general claim to derive from the particularities of titles—namely, that they have to be short. We have not taken these examples into account in our survey.

[5]See also Asher (2006) for an involved discussion of the differences between ellipsis or zero-anaphors, pronouns and definite descriptions.

Informally, a presupposition is resolvable in a sequence of contexts just in case some subsequence entails the presupposition or it is accommodated at some element in the sequence. Resolving the presupposition means choosing some witness for the existential quantifier. The clause for $\uparrow \phi$ then forces the binding or accommodation of ϕ in the outermost context possible.

Hunter and Asher (2005)'s analysis entails that different presuppositions may require different resolution strategies depending on their environment or associated presupposition triggers. This view contrasts with the standard view of presuppositions, but is well supported by examples of presupposition triggers like *too*, where accommodation is not possible. They hypothesize that such presupposition triggers introduce a presupposition of the form $B\phi$, where the recursive clause for $|\vdash$ for a formula like this is:

- $K_0, K_1, \ldots, K_j | \vdash B\phi$ iff for some $k \leq j$ and some $l \geq 0$,
  $K_0, \ldots K_{l-1}, [K_l, \ldots K_k] \models \phi$

That is, such a presupposition will be resolvable at a context iff the material under its scope can be bound.

Hunter and Asher (2005) extend their view to proper names and definite descriptions as well. They observe that definite descriptions by themselves don't seem to have a predetermined resolution strategy, but certain *readings* of them do. For example, the attributive use of a definite description makes the presupposed content part of the local assertion typically. Thus, we could take $\downarrow$ (the converse operator to $\uparrow$) and assign the presuppositional content ϕ of a description read attributively as $\downarrow \phi$. Definites understood referentially have their presuppositions prefixed again by the operator $\uparrow$. The account also extends to proper names, which generate presuppositions that prefer a global binding or accommodation and thus have the form. $\uparrow \phi$ to account for their semantically rigid behavior.

Hunter and Asher (2005) provide a declarative formalism enabling one to state where presuppositions prefer to attach and are processed. Adapting their analysis to the present context, we have to keep two things distinct. First, there is the tendency for the presuppositional content of names to project outside the scope of various operators, ensuring their semantically rigid behavior. However, they also have a discourse dimension which we need to capture in a theory like SDRT. Using SDRT, we can take over the notions of resolvability as it stands, only reinterpreting the ordering over contexts to be the ordering over available constituents in an SDRS, and reinterpreting binding and accommodation in terms of attachment to the constituent via the discourse relations Consequence and Background Asher and Lascarides (1998).

To this end we complicate the language of presupposition operators to allow for discourse attachments while preserving the semantic behavior of presuppositions. We introduce dynamic sequence ; to handle semantic instructions first, e.g, $\uparrow, B$ and then discourse instructions. We will also introduce the *else* operator of dynamic logic to handle preferences. Finally, we add discourse structure instructions involving presupposition attachment preferences given by triggers: B^d (attach to a constituent via a discourse relation that captures binding—e.g., *Defeasible-Consequence*). More complex instructions like van der Sandt (1992) proposal (V^d) can also be encoded. As far as the information concerning proper names is concerned, our descriptive study indicates that short proper names prefer to be bound but at some constituent above $LAST$, the last entered EDU in the contextually given SDRS. Thus, the presupposition would look something like this in the Hunter-Asher notation (also understanding short names to have the same *semantic* behavior vis a vis operators that Hunter and Asher stipulate).[6]

- $(\uparrow^s; B^d_{>\text{LAST}}, else V^d)\phi$

The differences between the contexts with which Hunter and Asher concern themselves in DRT and those in SDRT are relatively minor. But there is a much bigger theoretical hurdle to cross. Our observations have to do with how presuppositions enable the asserted constituents to attach. This is something that no theory has looked at up to now, as far as we are aware. In fact there

[6]This analysis can not yet account for the differences observed above between the behavior of long vs. short proper names.

is nothing in any theory that tells us how this should be done. But now how are we to code up the fact that the presupposition *also* tells the asserted content to attach to the constituent to which it binds? This appears to be not a matter for the presupposition itself, but rather for the attachment logic, i.e. the Glue logic of SDRT. We need a rule to the effect that says that if we have a short proper name (at least that), then the asserted component must attach to where the presupposition binds. The corpus results are extremely interesting in this regard. In 14 stories surveyed, we found *no* case where presuppositions attached to a different constituent than the asserted component, at least where the presuppositions of referential expressions were concerned. Thus, it appears that the rule attach asserted content to the place where presuppositions bind has the promise of being a relatively robust default. This is not to be expected in a theory like SDRT. But it gives us an important clue to determining attachments that is more than what SDRT had previously: figure out where the presuppositions of a given EDU attach, and attach the asserted content there too. Previously, in SDRT attachments were only guided by the possibility of resolving anaphoric dependencies, but with presuppositional material acting as independent constituents this only really amounts to attaching asserted material somewhere underneath the presupposed material (so that the values assigned to variables in the presupposed material are available to the asserted material).[7]

Up until now, SDRT has lacked a developed theory of attachment. It does say that all attachments have to be made on the right frontier (with the exception of discourse subordinating phenomena of the sort discussed in Asher 1993), but it gives no other clue concerning attachment. Moreover, the observations about discourse subordination really remained somewhat detached from the rest of the theory. Discourse subordinations occur when certain linguistic constructions like definite descriptions make reference to a particular segment in the discourse. These mostly occur in dialogue but not always—*let me go back to the third thing you said*. With the sort of rules we're now proposing, we can put in preferences on attachment.

4 Summary

The empircal study described in this paper partially confirms the referential hierarchies proposed by Ariel (1988) and Gundel et al. (1993) for example. The one place where our findings diverged from these previous studies was in the behavior of short proper names, which we found to require a more salient antecedent than not only full names, but also definite descriptions. Another finding concerned the interaction between transitions over the referential hierarchy and facts about discourse attachment. In particular, we saw that referential jumps appear to be a good indicator of discourse pops. Furthermore, the presence of a proper name, regardless of whether or not there is a transition on the RH, is a good indicator of a discourse pop. In other words, we find a correspondence between the salience of a referring expression and how high the segment containing it attaches in the discourse structure. Building on these findings, we offered a preliminary formalization in terms of the presuppositional behavior of the various types of referring expression.

References

Ariel, Mira. 1988. Referring and accessibility. *Journal of Linguistics* 24:65–87.

Asher, Nicholas. 1993. *Reference to abstract objects in discourse*. Kluwer Academic Publishers.

Asher, Nicholas. 2006. Troubles on the right frontier. In *Proceedings of Constraints on Discourse 2005*.

Asher, Nicholas, and Alex Lascarides. 1998. The semantics and pragmatics of presupposition. *Journal of Semantics* 15:239–299.

[7] The content of presuppositions is always accessible to the asserted content, so it must be superordinate to the attachment point of the asserted content. This give rise to funny background and consequence relations that are distinguished from their assertoric counterparts, i.e., $Background_p$ and $Consequence_p$.

Asher, Nicholas, and Alex Lascarides. 2003. *Logics of conversation*. Cambridge University Press.

Gundel, Jeanette K., Nancy Hedberg, and Ron Zacharski. 1993. Cognitive status of referring expressions in discourse. *Language* 69:274–307.

Hunter, Julie, and Nicholas Asher. 2005. A presuppositional account of indexicals. In *Proceedings of the Amsterdam Colloquium*.

Prince, Ellen F. 1981. Toward a taxonomy of given-new information. In *Radical pragmatics*, ed. P. Cole, 223–255. Academic Press.

Reinhart, Tanya. 1980. Conditions for text coherence. *Poetics Today* 1:161–180.

van der Sandt, Rob. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9:333–377.