


Inferential Realization Constraints on Functional Anaphora in the Centering Model

Katja Markert, Michael Strube & Udo Hahn

 Computational Linguistics Research Group
Freiburg University
Europaplatz 1
D-79085 Freiburg, Germany
{markert, strube, hahn}@coling.uni-freiburg.de

Abstract

We present an inference-based text understanding methodology for the resolution of functional anaphora in the context of the centering model. A set of heuristic realization constraints is proposed, which incorporate language-independent conceptual criteria (based on the well-formedness and conceptual strength of role chains in a terminological knowledge base) and language-dependent information structure constraints (based on topic/comment or theme/rheme orderings). We state text-grammatical predicates for functional anaphora and then turn to the procedural aspects of their evaluation within the framework of an actor-based implementation of a lexically distributed text parser.

Introduction

Textual forms of anaphora are a challenging issue for the design of parsers for text understanding systems, since lacking recognition facilities either result in referentially incohesive or invalid text knowledge representations. At the conceptual level *functional* anaphora relates a quasi-anaphoric expression to its antecedent by conceptual attributes (or roles) associated with that antecedent (see, e.g., the relation between “Ladezeit” (*charge time*) and “Akku” (*accumulator*) in (3) and (2) below). Thus it complements the phenomenon of *nominal* anaphora, where an anaphoric expression is related to its antecedent in terms of conceptual generalization (as, e.g., “Rechner” (*computer*) refers to “316LT”, a particular notebook, in (2) and (1) below). The resolution of text-level nominal (and pronominal) anaphora contributes to the construction of referentially valid text knowledge bases, while the resolution of text-level functional anaphora yields referentially cohesive text knowledge bases.

1. Der 316LT wird mit einem Nickel-Metall-Hydride-Akku bestückt.
(The 316LT is – with a nickel-metal-hydride-accumulator – equipped.)
2. Der Rechner wird durch diesen neuartigen Akku für ca. 4 Stunden mit Strom versorgt.
(The computer is – because of this new type of accumulator – for approximately 4 hours – with power – provided.)
3. Darüberhinaus ist die Ladezeit mit 1,5 Stunden sehr kurz.
(Also, – is – the charge time of 1.5 hours quite short.)

In the case of functional anaphora, the conceptual entity that relates the topic of the current utterance to discourse elements mentioned in the preceding one is not explicitly mentioned in the surface expression. Hence, the appropriate conceptual link must be inferred to establish the local coherence of the discourse (for an early statement of that idea, cf. Clark (1975)). In sentence (3) the information is missing that “Ladezeit” (*charge time*) is a property of “Akku” (*accumulator*). This relation can only be established if conceptual knowledge about the domain, viz. the relation *property-of* between the concepts CHARGE-TIME and ACCUMULATOR, is available.

The solution we propose to account for functional anaphora is embedded in the framework of the centering model (Grosz et al., 1995). In this approach, discourse entities serving to link one utterance to other utterances in a particular discourse segment are organized in terms of centers. The crucial notion for establishing local coherence links in discourse is that of realization. Given a center element of the previous utterance, we say this element is *realized* if it is associated with an expression in the following utterance that has a valid interpretation in the underlying semantic/conceptual representation language. Functional anaphora has only been given insufficient treatment within the centering model in terms of rather sketchy realization conditions as opposed to the more elaborated “direct realization” constraints formulated for (pro)nominal anaphora (cf. Grosz et al. (1995)). As these criteria are overly vague, we intend to supply a more precise, formally grounded notion of realization for the analysis of functional anaphora in the centering framework by proposing a set of heuristic realization constraints to guide the underlying inference processes. These include language-independent conceptual criteria (based on the well-formedness and conceptual strength of role chains in a terminological knowledge base) and language-dependent information structure constraints (based on topic/comment or theme/rheme orderings). The criteria we postulate contribute additional restrictions on the search space of possible referents and also direct inference processes required to understand anaphoric utterances in the discourse. Thus, they can be considered a more adequate explanatory model for local coherence than the original centering model in that they further limit the resource demands for proper text understanding.

Conceptual Constraints

We assume a concept hierarchy to consist of a set of concept names $\mathcal{F} = \{\text{COMPUTER-SYSTEM, ACCUMULATOR, ...}\}$ and a subclass relation $isa_{\mathcal{F}} = \{(\text{NOTEBOOK, COMPUTER-SYSTEM}), (\text{NIMH-ACCUMULATOR, ACCUMULATOR}), \dots\} \subset \mathcal{F} \times \mathcal{F}$. The set of relation names $\mathcal{R} = \{\text{has-physical-part, has-accumulator, charge-time-of, ...}\}$ contains the labels of possible conceptual roles. These are organized into a hierarchy by the relation $isa_{\mathcal{R}} = \{(\text{has-accumulator, has-physical-part}), (\text{charge-time-of, property-of}), \dots\} \subset \mathcal{R} \times \mathcal{R}$. We also assume the common understanding of the terms range, domain and inverse of a relation.

For the identification and evaluation of suitable conceptual links between an antecedent and a functional anaphor, a *path finder* performs an extensive unidirectional search in the domain knowledge base, looking for *well-formed* paths between the two concepts, while a *path evaluator* selects the *strongests* of the ensuing paths. We will not go into the formal details of well-formedness criteria for a conceptual path $(r_1 \dots r_n)$ ($r_i \in \mathcal{R}$) linking two concepts $x, y \in \mathcal{F}$. Instead, we only briefly mention that we require complete connectivity (compatibility of domains and ranges of the included relations) and non-cyclicity (exclusions of inverses of relations) for a *conceptually well-formed* path. The latter criterion, though entirely formal, achieves the discrimination alluded to by Resnik's (1995) distinction between similar and semantically related concepts. Additionally, a path from x to y will be excluded from the path list iff it *properly includes* another path from x to y and thus is *conceptually longer*.

Our focus in this paper will be on empirical criteria of path evaluation, *viz.* those which mark certain paths as being preferred over others in terms of commonsense plausibility. Based on the analyses of approximately 60 product reviews from the information technology domain and evidences reported from several (psycho)linguistic studies (e.g., Chaffin, 1992), we stipulate certain predefined *path patterns*. From those general path patterns and by virtue of the hierarchical organization of conceptual relations, concrete conceptual role chains can *automatically* be derived by the knowledge base system based on the operation of a classifier (we assume a terminological reasoning framework). This allows us to distinguish between a subset \mathcal{P} of all types of well-formed paths, which is labeled "*plausible*", another subset \mathcal{M} which is labeled "*metonymic*", and all remaining paths which are labeled "*implausible*".

Plausible Paths. We now turn to the question what kinds of relation chains should be characterized as plausible ones (forming the set \mathcal{P}), *i.e.*, which compositions of relation types are likely to create reasonable role chains. All paths of unit length l are included in \mathcal{P} as they are explicitly supplied in the domain knowledge base and are therefore "*plausible*", by definition. Regarding longer role chains we incorporate observations about the transitivity of (*part-whole*) relations made by Chaffin (1992) and Winston et al. (1987). They distinguish several subtypes of *part-whole* rela-

tions, e.g., integral object-component (corresponding to what we call *has-physical-part*), collection-member, mass-portion, process-phase, event-feature, area-place. The major claim they make is that any of these *subrelations* are transitive, while the most general *part-whole* relation usually is not. In other words, a relation chain containing only relations of one of the above-mentioned subtypes induces a *relation of the same subtype*, whereas a relation chain containing different types of *part-whole* relations is, in general, not reasonable any more. Following this argument, we have included the path patterns (*has-physical-part**), (*collection-member**), (*mass-portion**), (*process-phase**), (*event-feature**), (*area-place**) and the corresponding inverses like (*physical-part-of**), (*member-of**), etc. in \mathcal{P} . We refer to the first six of these basic patterns as *transitive-part-whole patterns*, in short \mathcal{T} , and to the inverse patterns as \mathcal{T}^{-1} . Compositionality of relation types other than *part-whole* relations has not received that much attention in the literature (one of the rare exceptions is the study by Huhns & Stephens (1989)). We follow some of their suggestions and also include (*spatial containment**) and (*connection**) in \mathcal{P} .

Metonymic Paths. We also incorporate *whole-for-part*, *part-for-whole*, and *producer-for-product* metonymies (cf. Lakoff, 1987; Fass, 1988). To determine path patterns corresponding to these types of metonymies consider the conceptual link between an instance of the concept C_1 and an instance of the concept C_3 , which characterizes a metonymy and thus stands for another instance of a concept C_2 . A corresponding well-formed conceptual path $p = (r_1 \dots r_n)$ with $n \in \mathbb{N}$, $n > 1$, and $r_i \in \mathcal{R}$ ($i = 1, \dots, n$) must, first, link C_1 to C_2 via $p_1 = (r_1 \dots r_{j-1})$ for some $j \in \{2, \dots, n\}$. C_2 is then linked to C_3 via $p_2 = (r_j \dots r_n)$. We have restricted the first link p_1 to plausible paths to provide reasonable metonymic chains only. The second link p_2 must express one of the metonymic relations $\mathcal{MS} = \{\text{has-part, part-of, produced-by}\}$, depending on the specific metonymy¹. For a *producer-for-product* metonymy, e.g., $j = n$ and $r_n = \text{produced-by}$ must hold. For a *part-for-whole* or *whole-for-part* metonymy, $j < n$ may be possible, as all paths in \mathcal{T} and \mathcal{T}^{-1} (e.g., (*has-physical-part**)) also express a single *has-part* or *part-of* relation (see the explanations of plausible paths above). For notational convenience, we will consider the paths in \mathcal{T} and \mathcal{T}^{-1} as a single relation so that we may write (*has-physical-part**) $isa_{\mathcal{R}}$ *has-part* or (*physical-part-of**) $\in \mathcal{MS}$. Thus, we may restrict the above cases of well-formed metonymic paths to the pattern in Table 1 from which special path patterns for specific metonymies can be derived.

¹If the direction of search is reversed (searching from C_3 to C_1) the corresponding inverse relations, $\mathcal{MS}^{-1} = \{\text{part-of, has-part, produces}\}$, must be considered. This list of metonymic relations is by no means complete and can be supplemented, if necessary. We have, as yet, included only the most frequent types of metonymies that occur in our application domain. The incorporation of further metonymic relations does not affect the operation of the algorithm, whatsoever.

$$\begin{aligned} \text{Metonymic-Path } ((r_1 \dots r_n)) : \Leftrightarrow \\ (r_1 \dots r_n) \notin \mathcal{P} \\ \wedge ((n > 1 \wedge (r_1, r_2, \dots, r_{n-1}) \in \mathcal{P} \wedge r_n \in \mathcal{MS}) \\ \vee (n > 1 \wedge (r_2, r_3, \dots, r_n) \in \mathcal{P} \wedge r_1 \in \mathcal{MS}^{-1})) \end{aligned}$$

Table 1: Metonymic Path Patterns

The computation of paths between an antecedent x and a functional anaphor y may yield several alternative types of well-formed paths, *viz.* “plausible”, “metonymic” or “implausible”. In order to make a proper selection we define a ranking on those different path markers according to their intrinsic conceptual strength, which we denote by the relation “ $>_{str}$ ” (*conceptually stronger than*) (cf. Table 2). As a consequence of this ordering, metonymic paths will be excluded from a path list iff plausible paths already exist, while implausible paths will be excluded iff plausible or metonymic paths already exist. Hence, only paths of the strongest type are retained in the final path list for a given concept pair x and y .

$$\text{“plausible”} >_{str} \text{“metonymic”} >_{str} \text{“implausible”}$$

Table 2: Ordering of Path Markers by Conceptual Strength

To evaluate these conceptual strength criteria we selected 80 concept pairs at random from the underlying domain knowledge base (459 concepts, 334 relations). We submitted them to the path finder/evaluator and compared the automatically generated conceptual paths with introspective judgments about the kinds of relations linking each pair. The overall error rate was below 5%. The average number of connected paths between two concepts (41.8) was reduced by the non-cyclicity criterion to 10.4 well-formed paths, and by the inclusion criterion to 2.4. The criterion in Table 2 achieves a final reduction to merely 1.8 paths. Hence, the criteria achieve the desired discrimination. We plan a broader evaluation of our approach by running the algorithm on larger-sized knowledge bases in order to test the domain-independence and scalability of the above criteria.

All conceptual paths which meet the above linkage criteria for two concepts, x and y , are contained in a list denoted by $CP_{x,y}$. As, in the case of functional anaphora, we have to deal with paths leading from the anaphoric expression to several alternative antecedents, we usually have to compare pairs of path lists $CP_{x,y}$ and $CP_{x,z}$, where $x, y, z \in \mathcal{F}$. We do this by applying the same criteria we used for evaluating paths linking single concepts. As all paths in $CP_{x,y}$ and $CP_{x,z}$ were computed by the path finder, they already fulfill the connectivity and non-cyclicity condition. The inclusion criterion cannot be applied to any paths $p_1 \in CP_{x,y}$ and $p_2 \in CP_{x,z}$, as p_1 and p_2 do not lead to the same concept ($y \neq z$). However, the criterion which ranks conceptual paths according to their associated path markers is applicable as all paths in a single CP list have the same marker. A function, $\text{PathMarker}(CP_{i,j})$, yields either “plausible”, “metonymic” or “implausible” de-

pending on the type of paths the list contains. Hence, the same ordering of path markers as in Table 2 can be applied to compare two CP lists (cf. Table 3).

$$\begin{aligned} \text{StrongerThan}(CP_{x,y}, CP_{x,z}) : \Leftrightarrow \\ \text{PathMarker}(CP_{x,y}) >_{str} \text{PathMarker}(CP_{x,z}) \\ \text{asStrongAs}(CP_{x,y}, CP_{x,z}) : \Leftrightarrow \\ \text{PathMarker}(CP_{x,y}) = \text{PathMarker}(CP_{x,z}) \end{aligned}$$

Table 3: Comparison of Path Lists by Conceptual Strength

Centering Constraints

Conceptual criteria are of tremendous importance, but they are not sufficient for properly resolving functional anaphora. Additional criteria have to be supplied in the case of equal strength of conceptual path lists for several alternative antecedents. We therefore incorporate into our model various information structure criteria in terms of topic/comment or theme/rheme patterns which originate from (dependency) structure analyses of the underlying utterance. The framework for this type of information is provided by the well-known *centering* mechanism (Grosz et al. (1995)), for which psycholinguistic evidences are provided by Gordon et al. (1993) and Brennan (1995).

The theory of centering is intended to model the local coherence of discourse, *i.e.*, coherence among the utterances U_i in a particular discourse segment (say, a paragraph of a text). Local coherence is opposed to global coherence, *i.e.*, coherence with other segments in the discourse. Each utterance U_i in a discourse segment is assigned a set of *forward-looking centers*, $C_f(U_i)$, and a unique *backward-looking center*, $C_b(U_i)$. The forward-looking centers of U_i depend only on the expressions that constitute that utterance; previous utterances provide no constraints on $C_f(U_i)$. The elements of $C_f(U_i)$ are partially ordered to reflect relative prominence in U_i . The most highly ranked element of $C_f(U_i)$ that is *realized* in U_{i+1} (*i.e.*, is associated with an expression that has a valid semantic interpretation) is the $C_b(U_{i+1})$. The ranking imposed on the elements of the C_f reflects the assumption that the most highly ranked element of $C_f(U_i)$ is the most preferred antecedent of an anaphoric expression in U_{i+1} , while the remaining elements are (partially) ordered according to decreasing preference for establishing referential links.

The theory of centering, in addition, defines several transition relations across pairs of adjacent utterances (*e.g.*, continuation, retention, smooth and rough shift), which differ from each other according to the degree by which successive backward-looking centers are confirmed or rejected, and, if they are confirmed, whether they correspond to the most highly ranked element of the current forward-looking centers or not. The theory claims that to the extent a discourse adheres to all these centering constraints (*e.g.*, realization constraints on pronouns, preferences among types of center transitions), its local coherence will increase and the inference load placed upon the hearer will decrease. Therefore, the

tremendous importance of fleshing out the relevant and most restrictive, though still general centering constraints.

The main difference between Grosz et al.'s work and our proposal concerns the criteria for ranking the forward-looking centers. While Grosz assume (for the English language) that *grammatical* roles are the major determinant for the ranking on the C_f , we claim that for German and other languages with relatively free word order it is the *functional* information structure of the sentence in terms of topic/comment or theme/rheme patterns (cf. Strube & Hahn (1996) for a more detailed account). In this framework, the *topic* (*theme*) denotes the given information, while the *comment* (*rheme*) denotes the new information. This distinction can be rephrased in terms of the centering mechanism. The *theme* then corresponds to the $C_b(U_n)$, the most highly ranked element of $C_f(U_{n-1})$ which is realized in U_n . The theme/rheme hierarchy of U_n is determined by the $C_f(U_{n-1})$: elements of U_n which are contained in $C_f(U_{n-1})$ (*context-bound* discourse elements) are less rhematic than elements of U_n which are not contained in $C_f(U_{n-1})$ (*unbound* elements). The distinction between context-bound and unbound elements is important for the ranking on the C_f , since bound elements are generally ranked higher than any other nonanaphoric elements.

Grammar Predicates for Functional Anaphora

We build on a grammar model which employs default inheritance for lexical hierarchies. The grammar formalism (for a survey, cf. Hahn et al. (1994)) is based on dependency relations between lexical heads and modifiers. The dependency specifications allow a tight integration of linguistic (grammar) and conceptual knowledge (domain model), thus making powerful terminological reasoning facilities directly available for the parsing process.² The resolution of functional anaphora is based on two major criteria, a conceptual and a structural one. The conceptual strength criterion for role chains is already specified in Table 3. The structural condition is embodied in the predicate *isPotentialFuncAntecedent* (cf. Table 4). A functional anaphoric relation between two lexical items is here restricted to pairs of nouns. The anaphoric phrase which occurs in the n -th utterance is restricted to be a definite NP and the antecedent must be one of the forward-looking centers of the preceding utterance.

The predicate *PreferredConceptualBridge* (cf. Table 5) combines both criteria. A lexical item y is determined as the proper antecedent of the functional anaphoric expression x iff it is a potential antecedent and if there exists no alternative antecedent z whose conceptual strength relative to x ex-

²We assume the following conventions: $\mathcal{C} = \{\text{Word, Nominal, Noun, PronPersonal, ...}\}$ denotes the set of word classes, and $isa_{\mathcal{C}} = \{(\text{Nominal, Word}), (\text{Noun, Nominal}), (\text{PronPersonal, Nominal}), \dots\} \subset \mathcal{C} \times \mathcal{C}$ denotes the subclass relation which yields a hierarchical ordering among these classes. Furthermore, *object.r* refers to the instance in the text knowledge base denoted by the linguistic item *object* and *object.c* refers to the corresponding concept class c . *Head* denotes a structural relation within dependency trees, viz. x being the head of y .

$$\begin{aligned} & \text{isPotentialFuncAntecedent}(y, x, n) : \Leftrightarrow \\ & y \text{ isa}_{\mathcal{C}} \text{ Nominal} \wedge x \text{ isa}_{\mathcal{C}} \text{ Noun} \\ & \wedge \exists z : (x \text{ head } z \wedge z \text{ isa}_{\mathcal{C}} \text{ DetDefinite}) \\ & \wedge x \in U_n \wedge y.r \in C_f(U_{n-1}) \end{aligned}$$

Table 4: Potential Functional Antecedent

ceeds that of y relative to x or, if the conceptual strength is equal, whose strength of preference under the TC relation is higher than that of y . " $>_{TC}$ " defines a partial order on the conceptual/semantic items of C_f reflecting the functional information structure of the utterance U_n in which their linguistic counterparts, viz. z and y , occur.

$$\begin{aligned} & \text{PreferredConceptualBridge}(y, x, n) : \Leftrightarrow \\ & \text{isPotentialFuncAntecedent}(y, x, n) \\ & \wedge \neg \exists z : \text{isPotentialFuncAntecedent}(z, x, n) \\ & \wedge (\text{StrongerThan}(\text{CP}_{x.c,z.c}, \text{CP}_{x.c,y.c}) \\ & \vee (\text{asStrongAs}(\text{CP}_{x.c,z.c}, \text{CP}_{x.c,y.c}) \wedge z >_{TC} y)) \end{aligned}$$

Table 5: Preferred Conceptual Bridge

The Resolution of Functional Anaphora

The actor computation model (Agha & Hewitt, 1987) provides the background for the procedural interpretation of lexicalized grammar specifications, as those given in the previous section, in terms of so-called word actors. Word actors communicate via asynchronous message passing; an actor can only send messages to other actors it knows about, its so-called acquaintances. The arrival of a message at an actor triggers the execution of a method that is composed of grammatical predicates (for a survey, cf. Neuhaus & Hahn (1996)).

The resolution of functional anaphors within texts depends on the results of the preceding resolution of nominal anaphors (Strube & Hahn, 1995) and the termination of the semantic interpretation of the current utterance. It will only be triggered at the occurrence of the definite noun phrase NP when NP is not a (pro)nominal anaphor and NP is only connected via certain types of relations (e.g., *has-property*, *has-physical-part*)³ to referents denoted in the current utterance at the conceptual level.

³Associated with the set \mathcal{R} is the set of inverse roles \mathcal{R}^{-1} . This distinction becomes crucial for already established relations like *has-property* (subsuming *charge-time*, etc.) or *has-physical-part* (subsuming *has-accumulator*, etc.) insofar as they do not block the initialization of the resolution procedure for functional anaphora (e.g., ACCUMULATOR – *charge-time* – CHARGE-TIME), whereas the existence of their inverses, we here refer to as *POF-type relations*, e.g., *property-of* (subsuming *charge-time-of*, etc.) and *physical-part-of* (subsuming *accumulator-of*, etc.), does (e.g., ACCUMULATOR – *accumulator-of* – 316LT). This is simply due to the fact that the semantic interpretation of a phrase like "the charge time of the accumulator" already leads to the creation of the *POF-type relation* the resolution mechanism for functional anaphora is supposed to determine. This is opposed to the interpretation of its ellipitied counterpart "the charge time" in sentence (3), where the genitive object is zeroed.

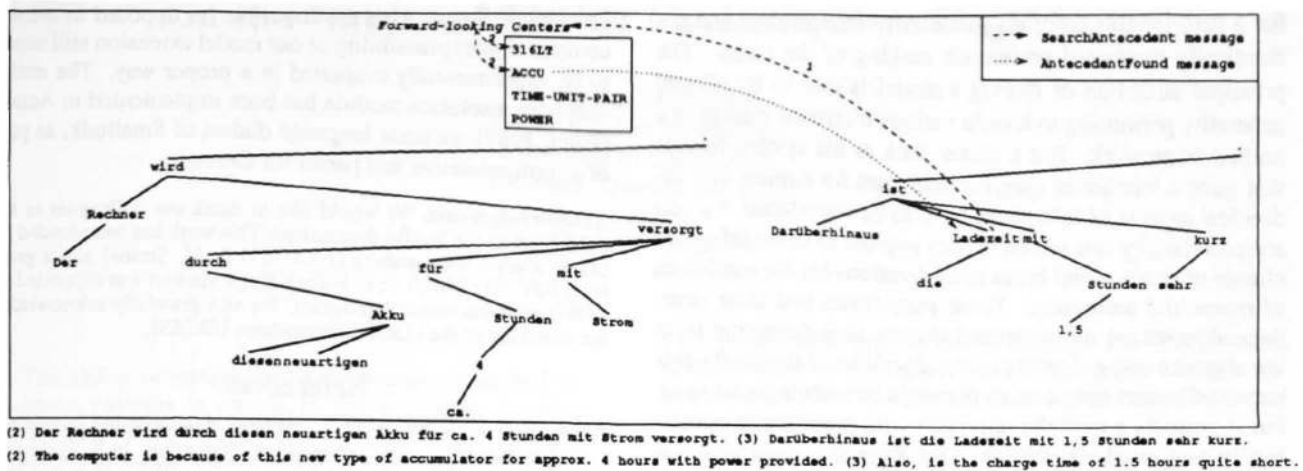


Figure 1: Sample Parse for the Resolution of a Functional Anaphor

The protocol level of actor-based text analysis encompasses the procedural interpretation of the grammatical predicates mentioned above. Fig. 1 illustrates the protocol for establishing proper conceptual relations, referring to the already introduced text fragment (1) - (3) which is partially repeated at the bottom line of Fig. 1. (3) contains the definite NP "die Ladezeit". Since "Ladezeit" (charge time) does not subsume any word at the conceptual level in the preceding utterance, the (pro)nominal anaphora test fails; the definite NP "die Ladezeit" has also not been integrated in terms of a functionally relevant conceptual relation as a result of its semantic interpretation. Hence, a *SearchTextFuncAntecedent* message is created by the word actor for "Ladezeit", which consists of two phases:

1. In *phase 1*, the message is forwarded from its initiator "Ladezeit" to the *forward-looking centers* of the *previous sentence*, where its state is set to *phase 2*.
2. In *phase 2*, the *forward-looking centers* of the *previous sentence* are tested for the predicate *PreferredConceptualBridge*, relative to the initiator of the *SearchAntecedent* message, viz. $x = \text{"Ladezeit"}$ (charge time).

The relevant knowledge base operations are performed on the four concepts associated with the current forward-looking centers, viz. 316LT, ACCUMULATOR, TIME-UNIT-PAIR (the conceptual representation for "Stunden"), and POWER⁴. In this case, the instance 316LT (the proper conceptual referent of the nominal anaphor "der Rechner" (the computer)) is related to CHARGE-TIME (the concept denoting "Ladezeit") via a *metonymic path*, viz. (charge-time-of accumulator-of) indicating a whole-for-part metonymy, while the concept ACCUMULATOR is related to CHARGE-TIME via a *plausible path* (viz. charge-time-of). As plausible paths are the

⁴Note that only nouns and pronouns are capable of responding to the *SearchTextFuncAntecedent* message and of being tested as to whether they fulfill the required criteria for a functional anaphoric relation.

strongest type of conceptual paths, only an element which is more highly ranked in the centering list and is linked via a plausible path to the functional anaphor could be preferred as the functional antecedent of "die Ladezeit" (the charge time) over "Akku" (accumulator) (according to the constraint from Table 5). As we know already that this is not the case, it is not necessary to test the remaining concepts associated with the current forward-looking centers (namely, TIME-UNIT-PAIR and POWER) and "Akku" can be selected as the proper functional antecedent. A *FuncAntecedentFound* message is sent from the word actor "Akku" to the initiator of the *SearchAntecedent* message, viz. "Ladezeit". An appropriate update links the corresponding instances via the role *charge-time-of* and, thus, local coherence is established at the conceptual level of the text knowledge base.

Comparison with Related Approaches

Searching links in a taxonomic hierarchy is a common application for spreading activation or marker passing techniques. The paradigm of path finding and evaluating they propose has obvious parallels to our approach. The criteria they employ, however, are mostly based on numerical restrictions, e.g., on weights (Charniak, 1986) or path lengths (Hirst, 1987). This is problematic as the foundation and derivation of these numbers is usually not made explicit (or it is *ad hoc*). We have tried to overcome this problem by stating structural and empirically plausible criteria which do not rely upon numerical restrictions in any way.

A pattern-based approach to inferencing closest in spirit to our approach has been put forward by Norvig (1989). The method he proposes can also be used to resolve (functional) anaphora. The main difference to our work lies in the fact that his path patterns are solely being defined in terms of "formal" link criteria in a knowledge base whose patterns are simply matched against the links being passed, whereas our definitions of path patterns take the semantic hierarchy of relations and their compositional properties into account. This allows

for a path-length- and thus granularity-independent and semantically motivated preference ranking of the paths. The principal attraction of Norvig's model is due to its alleged generality permitting to handle various inference classes in a unified framework. But a closer look at his system reveals that quite a number of specific-case rules for coping with individual aspects of inferences have to be introduced, e.g., an antipromiscuity rule which is only applied to some inference classes or recency and focus considerations for the resolution of referential ambiguity. These restrictions and their interdependencies are not expressed clearly, thus detracting from the elegance and generality of the algorithm. Admittedly, this paper addresses only one of Norvig's inferencing problems, but it presents a modular approach with precise and semantically motivated restrictions. Our algorithm combines two equally general, multi-purpose modules, viz. a path finder and a path evaluator, which are also used in the parsing process, and a centering mechanism which is applied to other forms of anaphora resolution problems as well. This has the advantage of a specific inference module with lucid triggering conditions.

The original centering model does not provide for methods for the resolution of functional anaphora. Grosz et al. rather sketchily point to the difference between the relations *directly realizes* and *realizes* whose precise definition they suggest depends on the semantic theory one adopts (Grosz et al., 1995, p.209). We have shown, however, that there are a lot of general constraints at the knowledge level which need not be covered by semantic theories at all.

Functional anaphora are also not an issue for standard grammar theories (e.g., HPSG, LFG, GB, CG, TAG). This is not at all surprising, as their advocates pay almost no attention to the text level (with the exception of several forms of pronominal anaphora) and also do not seriously take conceptual criteria as part of grammatical descriptions into account without which true text understanding seems infeasible.

Conclusions

The model of functional anaphora resolution we have outlined considers specific forms of conceptual inferences to be of primary importance. In order to constrain the realization of functional anaphora in the centering framework we propose conceptual well-formedness and strength criteria for role chains in a terminological knowledge base, by which the plausibility of various possible antecedents as proper bridges (Clark, 1975) to functional anaphora can be assessed. Information structure constraints on the underlying utterances in terms of topic/comment patterns contribute further inferential restrictions on proper antecedents for functional anaphora. Altogether, these extensions require a thorough revision of the original centering model. Our proposal has only been tested on moderately sized knowledge bases, with 800 and 500 concept/role specifications for the information technology and medicine domain, respectively, which are implemented in LOOM (MacGregor & Bates, 1987). So the scalability of the model still has to be demonstrated on larger sized

knowledge bases. Also the cognitive (as opposed to merely computational) plausibility of our model extension still needs to be experimentally evaluated in a proper way. The entire anaphora resolution module has been implemented in Actalk (Briot, 1989), an actor language dialect of Smalltalk, as part of a comprehensive text parser for German.

Acknowledgments. We would like to thank our colleagues in the *CLIF* group for fruitful discussions. This work has been funded by *LGFG Baden-Württemberg* (1.14-7631.0; M. Strube) and a grant from *DFG* (Ha 2907/1-3; U. Hahn). Katja Markert was supported by a *DFG* graduate research program. We also gratefully acknowledge the provision of the *LOOM* system from *USC/ISI*.

References

- Agha, G. & C. Hewitt (1987). Actors: A conceptual foundation for concurrent object-oriented programming. In B. Shriver & P. Wegner (Eds.), *Research Directions in Object-Oriented Programming*. Cambridge, Mass: MIT Press, pp.49-74.
- Brennan, Susan E. (1995). Centering attention in discourse. *Language and Cognitive Processes*, 10(2):137-167.
- Briot, J.-P. (1989). Actalk: A testbed for classifying and designing actor languages in the Smalltalk-80 environment. In *Proc. of ECOOP-89*, pp. 109-129.
- Chaffin, R. (1992). The concept of a semantic relation. In A. Lehrer & E.F. Kittay (Eds.), *Frames, Fields and Contrasts*. Hillsdale, NJ: Erlbaum, pp.253-288.
- Charniak, E. (1986). A neat theory of marker passing. In *Proc. of AAAI-86*, pp.584-588.
- Clark, H. (1975). Bridging. In *Proc. of TINLAP-1*, pp. 169-174.
- Fass, D (1988). An account of coherence, semantic relations, metonymy, and lexical ambiguity resolution. In S. Small, G. Cottrell & M. Tanenhaus (Eds.), *Lexical Ambiguity Resolution*. San Mateo, CA: Morgan Kaufmann, pp.151-178.
- Gordon, Peter C., Barbara J. Grosz & Laura A. Gilliom (1993). Pronouns, names, and the centering of attention in discourse. *Cognitive Science*, 17:311-347.
- Grosz, Barbara J., Aravind K. Joshi & Scott Weinstein (1995). Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21(2):203-225.
- Hahn, U., S. Schacht & N. Bröker (1994). Concurrent, object-oriented natural language parsing: The *ParseTalk* model. *International Journal of Human-Computer Studies*, 41(1/2):179-222.
- Hirst, G. (1987). *Semantic Interpretation and the Resolution of Ambiguity*. Cambridge: Cambridge University Press.
- Huhns, Michael & Larry Stephens (1989). Plausible inferencing using extended composition. In *Proc. of IJCAI-89*, Vol. 2, pp. 1420-1425.
- Lakoff, G. (1987). *Women, Fire and Dangerous Things*. Chicago: Chicago University Press.
- MacGregor, R. & R. Bates (1987). *The LOOM Knowledge Representation Language*. USC/ISI.
- Neuhaus, P. & U. Hahn (1996). Restricted parallelism in object-oriented lexical parsing. In *Proc. of COLING-96*.
- Norvig, P. (1989). Marker passing as a weak method for inferencing. *Cognitive Science*, 13(4):569-620.
- Resnik, P. (1995). Using information content to evaluate semantic similarity in a taxonomy. In *Proc. of IJCAI-95*. Vol.1., pp.448-453.
- Strube, M. & U. Hahn (1995). *ParseTalk* about sentence- and text-level anaphora. In *Proc. of EACL-95*, pp. 237-244.
- Strube, M. & U. Hahn (1996). Functional centering. In *Proc. of ACL-96*.
- Winston, M., R. Chaffin & D.J. Herrmann (1987). A taxonomy of part-whole-relations. *Cognitive Science*, 11(1):417-444.