A Dialogue System for Bare Fragments

Jonathan Ginzburg, Howard Gregory and Shalom Lappin
King's College London

Abstract

We present a dialogue system that interprets short questions and short answers within a typed feature structure grammar with a multidimensional construction hierarchy. The system provides a procedure for computing the content values of clausal fragments from contextual information contained in a discourse record of previously processed sentences. The feature system we use is based on an independently motivated model of dialogue interpretation.[1]

1 Introduction

A major task for a dialogue system is the resolution of fragments. Basic examples include **bare NP answers** (1), where the bare NP *John* is resolved as the assertion *John saw Mary*, and **sluicing** (2), where the *wh*-phrase *who* is interpreted as the question *Which student saw John.*

(1) A: Who saw Mary? B: John

(2) A: A student saw John. B. Who?

Either the antecedent or the fragment (or both) may be embedded (3).

(3)
1. A: Bill wonders who saw Mary. B: John. (*John saw Mary*)
2. A: Bill thinks a student saw John. B: Who? (*Which student does Bill think saw John?*)
3. A: Who saw Mary? B: John thinks Bill. (*John thinks Bill saw Mary*)
4. A: A student saw John. B: Bill wonders who. (*Bill wonders which student saw John*)

[1]Earlier versions of this paper were presented at the NLP colloquium of the Cambridge Computer Laboratory (March 2000), at the University College London Workshop on Syntax and Pragmatics (April 2000), and at the IBM Resarch NLP Colloquium in Hawthorne, NY (June, 2000). We are grateful to the participants of these forums for useful discussion of some of the ideas in this paper. In particular we would like to thank Steve Pulman and Richard Brehny for helpful comments on bare adjuncts. This research project is funded by grant number R00022269 from the Economic and Social Research Council of the United Kingdom.

Constructing a computational framework for interpreting such fragments in a principled and systematic way is important for the development of viable grammar-driven text understanding and dialogue management systems. The system presented here is implemented using [Erbach1995]'s ProFIT system, in which typed feature structures are compiled directly as Prolog terms. It employs a significantly modified version of the HPSG grammar and ellipsis resolution system of [Lappin and Gregory1997].

2 Theoretical Background

2.1 Ellipsis Resolution: a theory of context and parallelism

We view the task of accounting for an ellipsis phenomenon as involving: (a) locating an element in the context (the **source**) parallel to the ellipsis element (the **target**); and (b) computing from contextual information a property which composes with the target to yield the resolved content. In this respect we follow work on Higher Order Unification (**HOU**) [Dalrymple et al.1991, Pulman1997]. There are also similarities with the Dynamic Syntax approach of [Kempson et al.1999]. However, in contrast to existing HOU accounts (see section 4 below), we provide an explicit account of the relation between this parallelism and dialogue context.

We adapt the situation semantics based theory of dialogue context developed in the **Kos** framework [Ginzburg1996, Cooper et al.1999]. This combines a structuring of the propositional common ground of conversation [Webber1991, Asher1993] with a modelling of discourse topic based on [Carlson1983]. In [Ginzburg and Sag(forthcoming)] this framework is recast in the form of a typed feature structure grammar based on recent work in Head Driven Phrase Structure Grammar (HPSG) [Pollard and Sag1994, Sag1997, Ginzburg and Sag(forthcoming)]. Concretely, following [Ginzburg and Sag(forthcoming)], whose approach to describing ellipsis we adopt, we define two new attributes within the CONTEXT (CXT) feature structure: Maximal Question Under Discussion (MAX-QUD), whose value is of sort *question* and represents the question currently under discussion; and Salient Utterance (SAL UTT), whose value is a *sign*.

In this framework, questions are represented as semantic objects comprising a set of restricted indices and a SOA (4)[2]. This is the feature structure counterpart of the λ-abstract $\lambda\pi(\ldots\pi\ldots)$. In a *wh*-question the abstracted indices represent the INDEX values of the Wh phrases; for a polar question the INDICES set is empty.

(4) $$\begin{bmatrix} \textit{question} & \\ \text{INDICES} & \{\pi, \ldots\} \\ \text{SOA} & \begin{bmatrix} \textit{soa}(\ldots\pi\ldots) \end{bmatrix} \end{bmatrix}$$

In general a number of such questions may be available in a given discourse context, of which one is selected as the value of MAX-QUD. An algorithm is given below for the simple cases discussed in the present system, but it will be apparent that the system is flexible enough to be extended to more complicated dialogues.

[2] In [Ginzburg and Sag(forthcoming)] these are propositions rather than SOAs. A slightly simpler ontology is employed here. It is not difficult to restore G&S's distinction between propositions and SOAs if desired.

The feature SAL-UTT represents a distinguished (sub)utterance, a constituent of the utterance whose content is the current value of MAX-QUD. In information structure terms, it can be thought of as picking out and focalizing a potential parallel element. Since it is a *sign*, it enables us to encode syntactic *categorial* parallelism, thus ruling out cases of category mismatch and ensuring, for example, the correct case assignment for the fragment.[3]

Specifically, SAL-UTT is computed as the (sub)utterance associated with the role bearing widest scope within MAX-QUD:

- For *Wh*-questions, SAL-UTT is the *wh*-phrase associated with the INDICES set of the question.[4]
- If MAX-QUD is a question with an *empty* INDICES set, the context will be underspecified for SAL-UTT. Its possible values are either the empty set or the utterance of the widest scoping quantifier in the QUANTS list of MAX-QUD. This will be invoked to resolve sluicing.

2.2 A Grammatical Framework for Fragments

We adopt a version of HPSG, following [Sag1997, Ginzburg and Sag(forthcoming)], which encodes information about phrases by cross-classifying them in a multi-dimensional type hierarchy. Generalizations about the shared syntactic, semantic, and contextual properties of diverse types of phrases are captured by classifying phrases not only in terms of their phrase structure schema or X-bar type (e.g. whether they are headed or not; if they are headed, what kind of daughters are involved, etc.), but also with respect to a further informational dimension of CLAUSALITY. Each maximal phrasal type inherits both from a CLAUSALITY type and from a HEADEDNESS type. This classification allows us to specify systematic correlations between clausal construction types and types of semantic content.

[3] We invoke syntactic parallelism only for *matching conditions* between an elliptical form and a prior utterance. Thus, our approach is compatible with psycholinguistic work demonstrating the rapid decay of purely structural information [Fletcher1994]. Indeed, by relating MAX-QUD and SAL-UTT in the way we do, our approach enables us to make strong predictions about the range of possible relations of categorial parallelism.

[4] More generally, the utterance associated with the INDICES set when this is non-empty. This can arise when the antecedent is not an interrogative clause, for example in reprise or echo questions (mentioned in section 5; see [Ginzburg and Sag(forthcoming)] for detailed discussion). In such cases, SAL-UTT will be the utterance of the constituent to be clarified. Such cases illustrate that SAL-UTT need not be determined PRIOR to that utterance's taking place, but can be a consequence of how a conversationalist decides to structure the context.

Clausal Type Hierarchy

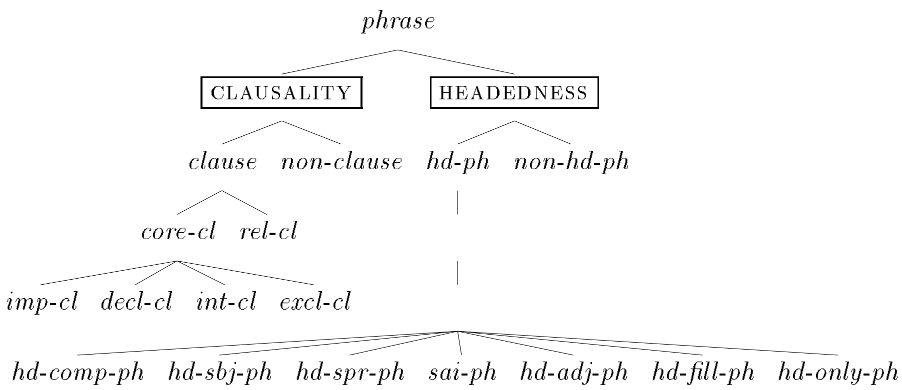

We restrict our attention here to two subtypes in the clausality dimension: declarative clauses (*decl-cl*) that denote SOAs and interrogative clauses (*inter-cl*) that denote questions.

At the same time, clauses are also typed as *full* or *bare*. In the case of elliptic sentences, instead of positing a phonologically null head, we simply assign fragments to some subtype of *bare-clause*. These are specified for a non-head daughter (the fragment), the remaining information for the interpretation being provided primarily by the CTXT features (5). The clausal content, notably the nucleus of the SOA, which in a headed clause would be inherited from the head daughter, is here constructed from the salient question (MAX-QUD). The (sub)utterance in the antecedent picked up by SAL-UTT enforces categorial identity (including case information) with the bare phrase and links it to the appropriate argument-role. The general constraints associated with the supertype *bare-clause* are shown in (5).

(5) *bare-ph* $\Rightarrow$

$$\begin{bmatrix} \text{STORE} & \{\} \\ \text{CONT} \mid \text{SOA} \mid \text{NUCL} & \boxed{1} \\ \text{CTXT} & \begin{bmatrix} \text{MAX-QUD} \mid \text{SOA} \mid \text{NUCL} & \boxed{1} \\ \text{SAL-UTT} & \begin{bmatrix} \text{CAT} & \boxed{2} \\ \text{CONT} \mid \text{INDEX} & \boxed{3} \end{bmatrix} \end{bmatrix} \\ \text{NHD-DTRS} & \left\langle [\begin{bmatrix} \text{CAT} & \boxed{2} \\ \text{CONT} \mid \text{INDEX} & \boxed{3} \end{bmatrix}] \right\rangle \end{bmatrix}$$

We define two subtypes of *bare-cl*: *bare-decl-cl* (for "short answers") and *bare-wh-cl* (for sluices). These subtypes are also subtypes of *decl-cl* and *wh-intr-cl* respectively, from which it follows that *bare-decl-cl* has a *soa* content, while *bare-wh-cl* denotes a question.

The resolved short answer thus contains the information in (6). The only features that remain to be specified are the scoping of quantifiers and the amalgamation of semantic restrictions. This is effected by the constraints in (7).

(6) *bare-decl-cl* $\Rightarrow$

$$
\begin{bmatrix}
\textsc{cont} & \begin{bmatrix} \textsc{quants} & \boxed{1} \oplus \boxed{2} \\ \textsc{restr} & \boxed{3} \cup \boxed{4} \end{bmatrix} \\
\textsc{ctxt} & \begin{bmatrix} \textsc{max-qud} & \begin{bmatrix} \textsc{indices} \mid \textsc{restr} & \boxed{3} \\ \textsc{soa} \mid \textsc{quants} & \boxed{2} \end{bmatrix} \end{bmatrix} \\
\textsc{nhd-dtrs} & \left\langle \left[\begin{bmatrix} \textsc{cont} \mid \textsc{index} \mid \textsc{restr} & \boxed{4} \\ \textsc{store} & \boxed{1}\ set(qf) \end{bmatrix} \right] \right\rangle
\end{bmatrix}
$$

(7)

1. **Quantifier scoping**. If the bare phrase is (or contains) a quantifier Q, then Q is scoped in wider than the existing quantifiers, if any, in the contextually salient question.
2. **Amalgamation of restrictions**. The (semantic) restrictions on the index I of the bare fragment are the union of the restrictions on I in MAX-QUD with the restrictions on I contributed by the bare phrase.

Turning to sluices (*bare-wh-cl*), it is noteworthy that existing linguistic work on sluicing has concentrated on embedded uses in monologue and have assumed (and indeed built into their accounts) the assumption that ONLY existentially quantified logical forms/propositions can serve as antecedents (see e.g. [Chung et al.1995, Reinhart1997]). However, as the following examples illustrate, the context for sluicing crucially involves the QUD-maximality of a question of the form WHETHER P, where *p* is a quantified proposition whose widest scoping quantifier is non-negative.[5]

(8) A: Many dissidents have been released. B: Do you know who?

(9) A: Did anyone show up for class today? B: Yes. A: Who?

(10) A: No student supported the proposal. B: hmm, # I wonder who (cf. I wonder who did.)

On our account, the type *bare-wh-cl* is defined for the information in (14), by virtue of type inheritance. Again, the only information that has to be added for this subtype concerns quantifiers and restrictions (15).[6]

[5] (Singular) definites are an exception to this, allowing as they do only reprise sluices:

(i) A: The murderer was obviously a vicious guy. B: #who?/WHO?/#You don't know who?/

(ii) # which murderer was obviously a vicious guy?

Our account will automatically associate the non-reprise sluice (i) with a content essentially synonymous with the non-elliptical (ii), which is indeed infelicitous in this context (i.e. as a non-reprise sluice).

[6] The account described here is concerned with bare answers and sluices corresponding to *arguments*. A fuller account must also cover *adjunct* fragments, such as the adjunct sluices below. Such forms of ellipsis, where there is no overt parallel element in the antecedent, pose a problem for all current frameworks (cf. the discussion in [Lappin and Gregory1997]).
(i) A: John saw Mary. B: When?
(ii) A: John likes Mary. B: Why?
While implementation of such cases is deferred to future work, in principle they can be handled as follows (cf. [Ginzburg1997]). Adjuncts in general have a CONT of type *soa* with a relation corresponding to a property of a *soa*, and a ROLE whose value is the *soa* content of the modified head.

(14) *bare-wh-cl* ⇒

$$\left[\begin{array}{ll}
\text{CONT} & \left[\begin{array}{ll}\text{INDICES} & \left[\begin{array}{ll}\text{INDEX} & \boxed{1}\\ \text{RESTR} & \boxed{2}\cup\boxed{3}\end{array}\right]\\ \text{SOA}\mid\text{QUANTS} & \boxed{4}\end{array}\right]\\
\text{CTXT} & \left[\begin{array}{ll}\text{MAX-QUD} & \left[\begin{array}{l}\text{INDICES}\quad \{\}\\ \text{SOA}\mid\text{QUANTS}\\ \langle\, \boxed{5}\left[\begin{array}{ll}\multicolumn{2}{l}{\textit{non-neg-qf-rel}}\\ \text{INDEX} & \boxed{1}\\ \text{RESTR} & \boxed{3}\end{array}\right] \oplus \boxed{4}\ \textit{list(qf)}\,\rangle\end{array}\right]\\ \text{SAL-UTT}\mid\text{STORE} & \boxed{5}\end{array}\right]\\
\text{NHD-DTRS} & \left\langle\ [\ \left[\begin{array}{ll}\text{STORE} & \left[\begin{array}{ll}\textit{index} & \boxed{1}\\ \text{RESTR} & \boxed{2}\end{array}\right]\end{array}\right]\ ]\ \right\rangle
\end{array}\right]$$

(15)
1. **Amalgamation of restrictions**. The (semantic) restrictions on the index I of the bare *wh*-phrase are the union of the restrictions on I in MAX-QUD with the restrictions on I contributed by the bare phrase.
2. **Quantifier replacement**. The index of the widest scoping quantifier Q in the antecedent's QUANTS list is replaced by the index of the non-head-dtr (the *wh*-phrase). This follows from the fact that Q is identified as the content of SAL-UTT, so that its index is unified with the index stored in the

---

(11)
$$\left[\begin{array}{ll}\text{CONT}\mid\text{SOA} & \left[\begin{array}{ll}\text{ADJ-REL} & \textit{at/because/manner/...}\\ \text{SOA-ROLE} & \textit{soa}\end{array}\right]\end{array}\right]$$

If the adjunct is interrogative, then it has an INDEX representing abstraction over the corresponding type of relation (e.g. temporal or causal).

(12) when:
$$\left[\begin{array}{ll}\text{CONT}\mid\text{SOA} & \left[\begin{array}{ll}\text{ADJ-REL} & \left[\begin{array}{ll}\text{INDEX} & \boxed{1}\\ \text{RESTR} & \textit{temporal-rel}(\boxed{1})\end{array}\right]\\ \text{SOA-ROLE} & \textit{soa}\end{array}\right]\end{array}\right]$$

(13) When did John see Mary?
$$\left[\begin{array}{ll}\text{INDICES} & \left\{\left[\begin{array}{ll}\text{INDEX} & \boxed{1}\\ \text{RESTR} & \textit{temporal-rel}(\boxed{1})\end{array}\right]\right\}\\ \text{SOA} & \left[\begin{array}{ll}\text{ADJ-REL} & \boxed{1}\\ \text{SOA-ROLE} & \left[\begin{array}{ll}\multicolumn{2}{l}{\textit{see-rel}}\\ \text{SEE-ER} & \boxed{2}\\ \text{SEE-EE} & \boxed{3}\end{array}\right]\end{array}\right]\end{array}\right]$$

The INDEX is put into store and retrieved for the INDICES set at the interrogative clause level.

A bare answer to a question with an interrogative adjunct would take that adjunct as its SAL-UTT value and substitute its own relation for the latter's INDEX. In the case of an adjunct sluice, however, SAL-UTT is empty (there is no parallel constituent to be picked up), and the non-head daughter contributes its INDEX value to store and hence to the INDICES of the bare clause. It also contributes restrictions on the index determining the semantic type of the relation that the index is to abstract over in the CONT of the recovered question. In the examples, the abstracted INDEX can be understood as (i) a temporal location bearing a time-of relation, or (ii) a soa bearing a cause-of relation, to the modified soa.

non-head-dtr.[7]

3 An Implemented System for Short Answers and Bare Sluiced Questions

Our fragment interpretation system consists of four main components:

I. An HPSG grammar. The grammar is a substantially modified version of that employed in the HPSG ellipsis resolution system of [Lappin and Gregory1997, Gregory and Lappin1999]. It is encoded in [Erbach1995]'s ProFIT system for representing typed feature structures as Prolog terms. A bottom-up chart parser is used to apply the grammar to input strings and generate AVM representations of typed feature structures.

II. A dialogue record. When a clause has been parsed and any ellipsis resolved by the procedures described in III and IV), its AVM is converted into a transitive list of Mother-Daughter-Relations (an **MDR** list). This format greatly simplifies the task of searching a feature structure for specific phrasal types and feature values. The MDR list is stored in a dialogue record paired with an index (counter) indicating the relative position of the AVM in the dialogue sequence. In the simple dialogue sequences implemented so far, this indexing corresponds to a linear sequence of utterances, but the format is designed so that in can be enriched to capture more comlicated dialogue structures. In addition to the MDR list, a list of MAX-QUD candidates is computed corresponding to the CONT of each subclause and stored as a component of the discourse record (the QUD-list or **candidate list**).

III. A CONTEXT resolution procedure. This assigns values from the dialogue record to the MAX-QUD and SAL-UTT features of the current clause C. The MAX-QUD value of C is an element selected from the candidate list, i.e. a question associated with the CONT value of an antecedent (sub-)clause A (in the default case, the most recent clause in the record). If A is a wh-question, its CONT value is simply taken as C's MAX-QUD. If A is a declarative clause, the value of MAX-QUD is constructed as a polar question whose nucleus is A's CONT value. The SAL-UTT feature value of C is taken as the AVM of a constituent of the antecedent, namely that constituent whose index corresponds to the INDICES list of MAX-QUD. If MAX-QUD is empty, SAL-UTT is identified with a quantified NP whose content is in the QUANTS list of MAX-QUD. These two cases capture bare answers and sluices respectively (see section 2 above).

IV. A bare clause resolution procedure. This computes the CONT of C by (i) unifying its NUCLEUS value N with the NUCLEUS of its MAX-QUD, and (ii) identifying the INDEX of its NON-HD-DTR with that of the SAL-UTT, causing this index to fill the appropriate role in N's relation. The index inherits restrictions both from its occurrence in SAL-UTT and its occurrence in the non-head daughter. If C is a *bare-wh-cl*, this index is further added to the INDICES list, removing the (quantified) **index** of SAL-UTT from STORE. Finally, it computes the syntactic category, including case, of the non-head daughter by unifying it with the CAT value of SAL-UTT.

[7] Thus, the widest scoping quantifier, if any, in the open proposition of the question after resolution will be whichever quantifier was previously scoped just narrower than Q.

In (1), repeated as (16), the antecedent clause has the CONT (17). The (abbreviated) AVM for the subject *wh*-phrase *who* is shown in (18). These are identified as the MAX-QUD and SAL-UTT respectively for the bare answer *john*, which is resolved as in (19), where $\boxed{1}$ is unified with the index of *John*', and $\boxed{2}$ is the index of *Mary*.

(16) A: Who saw Mary? B: John

(17)
$$\begin{bmatrix} \text{INDICES} & \begin{bmatrix} \text{INDEX} & \boxed{1} \\ \text{RESTR} & \left\{ \textit{person-rel}(\boxed{1}) \right\} \end{bmatrix} \\ \text{SOA} \mid \text{NUCL} & \begin{bmatrix} \textit{see-rel} & \\ \text{see-er} & \boxed{1} \\ \text{seen} & \boxed{2} \end{bmatrix} \end{bmatrix}$$

(18)
$$\begin{bmatrix} \text{CAT} & \text{NP[+nom]} \\ \text{CONTENT} & \boxed{1} \\ \text{STORE} & \begin{bmatrix} \text{INDEX} & \boxed{1} \\ \text{RESTR} & \left\{ \text{person-rel}(\boxed{1}) \right\} \end{bmatrix} \end{bmatrix}$$

(19)
$$\begin{bmatrix} \text{SOA} & \begin{bmatrix} \text{NUCL} & \begin{bmatrix} \textit{see-rel} & \\ \text{see-er} & \boxed{1} \\ \text{seen} & \boxed{2} \end{bmatrix} \\ \text{RESTR} & \left\{ \textit{person-rel}(\boxed{1}) \right\} \end{bmatrix} \end{bmatrix}$$

The content generated for *who* as a sluice in (2), repeated as (20), is as in (21), where [2] is the INDEX of *John*.

(20) A: A student saw John. B. Who?

(21)
$$\begin{bmatrix} \text{INDICES} & \begin{bmatrix} \text{INDEX} & \boxed{1} \\ \text{RESTR} & \left\{ \begin{array}{l} \textit{person-rel}(\boxed{1}), \\ \textit{student-rel}(\boxed{1}) \end{array} \right\} \end{bmatrix} \\ \text{SOA} \mid \text{NUCL} & \begin{bmatrix} \textit{see-rel} & \\ \text{see-er} & \boxed{1} \\ \text{seen} & \boxed{2} \end{bmatrix} \end{bmatrix}$$

The system recognizes the embedded fragments in (3), repeated as (22), as bare clauses of the appropriate type, and assigns them the AVM's in (19) and (21) respectively (embedded in a matrix AVM if appropriate).

(22)
1. A: Bill wonders who saw Mary. B: John. (*John saw Mary*)
2. A: Bill thinks a student saw John. B: Who? (*Which student does Bill think saw John?*)
3. A: Who saw Mary? B: John thinks Bill. (*John thinks Bill saw Mary*)

4. A: A student saw John. B: Bill wonders who. (*Bill wonders which student saw John*)

It is, furthermore, able to deal correctly with the dialogue sequence in (23), where "cascaded" instances of bare answers and sluices interact to give the specified interpretion of the final bare answer *Mary*. The correct result is equally obtained if any of the utterances are arbitrarily embedded. It is worth noting the effect of the rules given for quantifiers and restrictions (section 2), summarized in (24). The simplified AVM for the final bare answer is shown in (25).

(23)
1. A: Who saw John?
2. B: A girl. (= *A girl saw John.*)
3. A: Who? (= *Which girl saw John?*)
4. B: Mary. (= *A girl called Mary saw John*)

(24)
1. Index I1 of "who" in list of interrogative indices. Quants: empty
   Restrictions on I1: {human(I1)}
2. I1 removed from interrogative indices, identified with index of *a girl*. Quants: content C1 of *a girl*
   Restrictions on I1: {human(I1), girl(I1)}
3. Quants: empty (C1 removed). Interrogative indices: I1 (= index of C1)
   Restrictions on I1: {human(I1), girl(I1)}
4. I1 removed from interrogative indices, identified with index of *Mary*. Quants: empty
   Restrictions on I1: {human(I1), girl(I1), named(I1, mary)}

(25)

$$
\begin{bmatrix}
\text{PHON} & \text{mary}\\
\text{CAT} & \text{S[fin]}\\
\text{CONT} & \begin{bmatrix}\text{SOA} & \boxed{3}\\ \text{RESTR} & \boxed{4}\cup\boxed{6}\end{bmatrix}\\
\text{CTXT} & \begin{bmatrix}\text{MAX-QUD} & \begin{bmatrix} \textit{question} & \\ \text{INDICES} & \{\boxed{1}\}\\ \text{SOA} & \boxed{3}:\begin{bmatrix}\textit{see-rel} & \\ \text{see-er} & \boxed{1}\\ \text{seen} & \boxed{2}\end{bmatrix}\\ \text{RESTR} & \boxed{4}\end{bmatrix}\\ \text{SAL-UTT} & \begin{bmatrix}\text{CAT} & \boxed{5}:\textit{NP[nom]}\\ \text{INDEX} & \boxed{1}\end{bmatrix}\end{bmatrix}\\
\text{NHD-DTRS} & \left\langle \left[\begin{bmatrix} \text{PHON} & \text{mary}\\ \text{CAT} & \boxed{5}\\ \text{CONT} & \begin{bmatrix}\text{INDEX} & \boxed{1}\\ \text{RESTR} & \boxed{6}:\begin{bmatrix}\text{NAME} & \text{mary}\\ \text{NAMED} & \boxed{1}\end{bmatrix}\end{bmatrix}\end{bmatrix} \right] \right\rangle
\end{bmatrix}
$$

At the same time the requirement that the CAT feature of the non-head-dtr of the bare clause unify with that of the SAL-UTT assigns the correct case (nominative) to each of the fragments. It also rules out cases of category mismatch. For example the PP *to Mary* would not be accepted as a possible answer to the final question in the sequence, as the value of SAL-UTT is a (nominative) NP (*who*).

Similarly the accumulation of restrictions from the SAL-UTT and the CONT of the bare phrase on the index which fills the relation in the NUCLEUS of the bare clause's CONT value results in the identification of, for example, a non-human NP such as *A book*, as a *semantically* ill-formed bare answer.

4 Comparison with Other Approaches

Our system for interpreting bare clauses shares many features with the Higher Order Unification (HOU) account of ellipsis resolution [Dalrymple et al.1991, Pulman1997]. Like HOU the system generates the propositional content of a bare clause by removing an argument from the semantic representation of an antecedent clause, where this argument is treated as the parallel of the fragment phrase, in order to obtain a property term that is applied to the fragment. However, it differs from HOU in three important respects.

First, while HOU does not indicate how the relation of parallelism between the fragment and a counterpart term in the antecedent clause is specified, we provide an explicit definition of the counterpart term as the phrase that supplies the salient utterance (the value of SAL-UTT for the bare clause).[8]

Second, unlike HOU, we impose a syntactic matching condition on the category of the fragment and its counterpart term. This permits us to rule out cases of categorial mismatch, as discussed above. Moreover, because the entire AVM of an antecedent is recoverable from the discourse record, we can invoke additional syntactic constraints on fragment interpretation if these are required, as has been suggested for other cases of ellipsis resolution in, for example, [Lappin and Gregory1997].

Third, HOU is not possible when the semantic type of the target is distinct from that of the source, as when the target denotes a question and the source is a proposition or vice versa. [Pulman1997] attempts to bypass this problem by associating propositional contents with interrogatives. However, as the formal semantic literature on interrogatives suggests [Groenendijk and Stokhof1997], such a move is not semantically viable. Our system can handle such mismatches in semantic type by using the SOA and QUANTS features of different clause types present in the discourse record to specify the MAX-QUD and SAL-UTT values of a bare clause.

[Lappin and Shih1996] propose a generalized algorithm for ellipsis resolution, which is implemented in the system described in [Lappin and Gregory1997]. In this system bare NP fragments are treated as the NON-HD-DTRS of clauses with empty heads, which are replaced in resolution by an antecedent head. Counterpart arguments are replaced by the overt fragment, while non-counterpart arguments (and adjuncts) are copied over. The present model avoids the need for empty heads and syntactic reconstruction. Only the

[8] [Gardent and Kohlhase1997] combine a version of HOU with an abductive calculus for computing parallel elements. This invokes relations defined over a hierarchical classification of the lexical semantics of the subconstituents, following the work of [Hobbs1990] on parallelism in discourse. To the best of our knowledge this system has not yet been implemented.

semantics of the fragment clause and the category information of its non-head daughter are restored. Fragment resolution is not construed as the full reconstruction of missing material in a bare clause, but as the generation of a clausal content and limited syntactic properties from contextual information obtained from antecedents in the discourse record.

5 Conclusions and Future Work

We have presented a dialogue system that interprets short questions and short answers within a typed feature structure grammar. The system provides a procedure for computing the content values of clausal fragments from contextual information contained in a discourse record of previously processed sentences. The feature system we use is based on an independently motivated model of dialogue interpretation.

In our continuing research we are extending our system to handle a wide range of fragment types, including echo questions (26), verb phrase ellipsis (27), gapping (28), and pseudo–gapping (29). We will also apply our framework to fragments which occur without an overt linguistic antecedent. As part of this research we are investigating large scale corpora in order to construct a representative typology of actually occurring fragments.

(26) A: Mary sang. B: WHO?

(27) A: Mary wants vodka. B: Rosa does too.

(28) A: John gave flowers to Mary. B: And chocolates to Rosa.

(29) A: John sent flowers to Mary. B: And Bill did chocolates.

References

[Asher1993] Nicholas Asher. 1993. *Reference to Abstract Objects in English: a Philosophical Semantics for Natural Language Metaphysics*. Studies in Linguistics and Philosophy. Kluwer, Dordrecht.

[Carlson1983] Lauri Carlson. 1983. *Dialogue Games*. Synthese Language Library. D. Reidel, Dordrecht.

[Chung et al.1995] Sandy Chung, Bill Ladusaw, and James McCloskey. 1995. Sluicing and Logical Form. *Natural Language Semantics*.

[Cooper et al.1999] Robin Cooper, Staffan Larsson, Massimo Poesio, David Traum, and Colin Matheson. 1999. Coding Instructional Dialogue for Information States. In *Task Oriented Instructional Dialogue (TRINDI)*. University of Gothenburg, Gothenburg.

[Dalrymple et al.1991] Mary Dalrymple, Fernando Pereira, and Stuart Shieber. 1991. Ellipsis and Higher Order Unification. *Linguistics and Philosophy*, 14:399–452.

[Erbach1995] Gregor Erbach. 1995. ProFIT: Prolog with Features, Inheritance and Templates. In *Proceedings of the Seventh European Conference of the Association for Computational Linguistics*, pages 180–187.

[Fletcher1994] Charles Fletcher. 1994. Levels of Representation in Memory for Discourse. In Morton Ann Gernsbacher, editor, *Handbook of Psycholinguistics*. Academic Press.

[Gardent and Kohlhase1997] Claire Gardent and Michael Kohlhase. 1997. Computing Parallelism in Discourse. In IJCAI 2:1016-1021.

[Ginzburg and Sag(forthcoming)] Jonathan Ginzburg and Ivan Sag. (forthcoming) *English Interrogative Constructions*. CSLI Publications, Stanford. (An earlier version appeared as Lecture Notes for LSA Linguistic Institute. University of Illinois at Urbana Champaign.)

[Ginzburg1996] Jonathan Ginzburg. 1996. Interrogatives: Questions, Facts, and Dialogue. In Shalom Lappin, editor, *Handbook of Contemporary Semantic Theory*. Blackwell, Oxford.

[Ginzburg1997] Jonathan Ginzburg. 1997 Structural Mismatch in Dialogue. In Jaeger, G. and Benz, A., editors, *Proceedings of MunDial 97 (Technical Report 97-106*, pages p. 59–80. Universitaet Muenchen Centrum fuer Informations- und Sprachverarbeitung, Muenchen.

[Gregory and Lappin1999] Howard Gregory and Shalom Lappin. 1999. Antecedent Contained Ellipsis in HPSG. In Gert Webelhuth, Jean-Pierre Koenig and Andreas Kathol, editors, *Lexical and Constructional Aspects of Linguistic Explanation*. CSLI publications.

[Groenendijk and Stokhof1997] Jeroen Groenendijk and Martin Stokhof. 1997. Questions. In Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Linguistics*. North Holland, Amsterdam.

[Hobbs1990] Jerry Hobbs. 1990. *Literature and Cognition*. CSLI Lecture Notes, Number 21.

[Kempson et al.1999] Ruth Kempson, Wilfried Meyer-Viol and Dov Gabbay. 1999. VP Ellipsis: Towards a Dynamic, Structural Account. In Shalom Lappin and Elabbas Benmanoun, editors, *Fragments*. Oxford University Press, Oxford.

[Lappin and Gregory1997] Shalom Lappin and Howard Gregory. 1997. A Computational Model of Ellipsis Resolution. In Richard Oehrle, editor, *Proceedings of the Conference on Formal Grammar*. ESSLLI, Aix en Provence.

[Lappin and Shih1996] Shalom Lappin and Hsue-Hueh Shih. 1996. A Generalized Reconstruction Algorithm for Ellipsis Resolution. In *Proceedings of COLING-96*, pages 687–692.

[Pollard and Sag1994] Carl Pollard and Ivan Sag. 1994. *Head Driven Phrase Structure Grammar*. University of Chicago Press and CSLI, Chicago.

[Pulman1997] Stephen Pulman. 1997. Focus and Higher Order Unification. *Linguistics and Philosophy*, 20.

[Reinhart1997] Tanya Reinhart. 1997. Quantifier scope: How labor is divided between QR and choice functions. *Linguistics and Philosophy*, 20:335–397.

[Sag1997] Ivan Sag. 1997. English Relative Clause Constructions. *Journal of Linguistics*, 33:431–484.

[Webber1991] Bonnie Webber. 1991. Structure and Ostension in the Interpretation of Discourse Deixis. *Language and Cognitive Processes*, 14:107–135.