Bridging inferences in a dynamic frame theory

Ralf Naumann and Wiebke Petersen

Institut für Sprache und Information
Universität Düsseldorf
Germany

Abstract. In this article we develop a theory of bridging inferences in a dynamic frame theory that is an extension of Incremental Dynamics. In contrast to previous approaches bridging is seen as based on predictions/expectations that are triggered by discourse referents in a particular context where predictions are (more specific) instances of Questions under Discussion. In our frame theory each discourse referent is associated with a frame f that contains the information known about it in the current context. Predictions/QuDs are modelled as sets F of extensions of this frame relative to a (possibly complex) attribute about whose value no information is given so far. A continuation of the current context answers a question if it introduces a frame f' that contains information about the value of the attribute corresponding to the question. The set F is constrained by a probability distribution on the domain of frames. Only those extensions are considered whose conditional probability in the current context is high. The relation between f and f' can be restricted in several ways. Bridging inferences correspond to those restrictions in which (i) the frames belong to the semantic representations of two clauses and (ii) the relation is established by a separate update operation.[1]

1 Introduction

It is by now a well-known fact that the semantic processing of an utterance usually involves different sources of information which are used in parallel to arrive at a coherent interpretation of this utterance in the given context. Four principle sources must be distinguished: (i) the (linguistic) meaning of the lexical items; (ii) (non-linguistic) world and situational knowledge, (iii) the prior linguistic context and (iv) the information structure of the text, i.e. the way sentences are related by coherence relations and questions under discussion. A prime example of this interplay between different sources of information are bridging inferences. [AL98, 83p.] take bridging to be 'an inference that two objects or events that are introduced in a text are related in a particular way that isn't explicitly stated, and yet the relation is an essential part of the content of the text in the sense that without this information, the lack of connection between the sentences would make the text incoherent.' Examples of bridging inferences are given in (1) and (2).

(1) a. Lizzy met a dog yesterday. The dog was very friendly. [AL98, 86p.]

[1] The research was supported by the German Science Foundation (DFG) funding the Collaborative Research Center 991. We would like to thank the reviewers and audience of TbiLLC 2017 for helpful comments and suggestions.

b. John unpacked the picnic. The beer was warm. [CH77]
c. I was at a wedding last week. The mock turtle soup was a dream. [Geu11]
d. I've just arrived. The camel is outside and needs water. [AL98, 86p.]

(2) a. In the group there was one person missing. It was Mary who left.
b. John partied all night. He's going to get drunk again today.
c. Jack was going to commit suicide. He bought a rope. [Cha83]

Bridging inferences are most prominently related to definite descriptions, witness the examples in (1). This is, however, not the only possibility. They can also be triggered by 'it'-clefts, (2-a), temporal adverbials like 'again', (2-b), and indefinites like 'a rope' in (2-c), as shown by the examples in (2).

Common to all bridging inferences is (i) a new discourse referent is introduced (see [Bur06] for neurophysiological evidence) and (ii) a dependency (bridging) relation between this discourse referent (corresponding to the bridged expression) and a discourse referent that has already been introduced in the linguistic context (denoting an antecedent object) is established. Bridging inferences are often related to a presupposition. For example, the definite description 'the dog' in (1-a) triggers the presupposition that there is a unique dog in the context. The bridging inference consists in establishing a link between 'a dog' and 'the dog'. In this case the dependency relation is the identity relation. The dog introduced in the first sentence by the indefinite is identical to the dog denoted by the definite description in the second sentence. In (2-b) 'again' can be used felicitously only if John got drunk before today. The bridging inference is the inference that the previous occasion of John getting drunk was concurrent with his partying all night the day before. As noted in [Cla77] and [AL98], bridging inferences may occur in the absence of presupposition triggers as well. An example is (2-c) with the bridged expressing 'a rope' and the bridging inference that the rope is related to the planned suicide. It was the instrument to be used by Jack in his plan. An important further aspect of bridging inferences is that they provide additional information about the antecedent object. For example, in (1-d) the additional information is about the arriving event. The means of transport used in this event, or the presupposed moving event leading to the arrival, was a camel. In (2-c) the rope is the instrument used in the planned suicide and in (1-a) a comprehender gets to know that the dog introduced in the first sentence was very friendly. (1-c) shows that the dependency can be indirect. The turtle soup is directly related as a part (starter) of the meal which was served at the wedding. The examples in (1) and (2) in addition show that the dependency relation can be instantiated by various forms of relations: (a) identity (1-a), (b) constituent part-of (e.g. (1-b)) or concurrency, (2-b) (cf. [Cla77] for a comprehensive taxonomy). Due to lack of space, we will restrict bridging inferences to those cases involving NPs of the form 'the N' and 'an N' and hence to examples like those in (1) and (2-c).

2 Frame theory

At their core frames are attribute-value structures. Their strength for an analysis of bridging inferences lies in the fact that they allow for a fine-grained analysis of individuals and events.

Consider the frame in Figure 1.

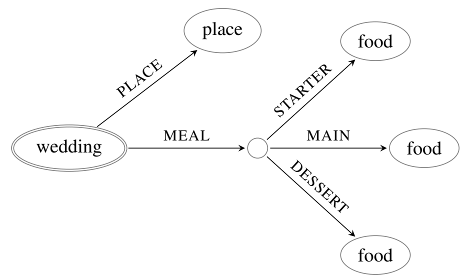

Fig. 1. wedding frame

This frame can be taken as a partial description of a wedding.[2] This wedding takes place at a particular location and the meal served had three parts: a starter, a main course and a dessert. This example shows that a frame contains two different kinds of information: relational information which links two objects in a frame via a chain of attributes and sortal information which classifies an object in the frame as belonging to a particular class (or sort) of objects. Relational information is represented by labeled arcs where the label indicates the arc. Sortal information is represented by circles with the sort being indicated by the label inside the circle. For example, the wedding is mapped to its location by the attribute PLACE and to the dessert served by the chain of attributes MEAL and DESSERT in that order. All three components of the meal are classified as being of sort **food**. Relational and sortal information are linked in a particular way. For each attribute, there is a source sort and a target sort. For example, the attribute PLACE has as source sort **physical object** including both individuals (human beings, engines, dogs etc.) and events (weddings, hittings, buying, eatings etc.) and as target sort objects of sort **place**. The target sort of the attributes STARTER, MAIN and DESSERT are all **food**. The same holds for the chains made up by the attributes MEAL and STARTER, MEAL and MAIN as well as MEAL and DESSERT.

Next, we will make the informal characterization given above more precise. One way of looking at the above figure is in terms of a relational model $\mathcal{M}$. Each chain of attributes is satisfied in a corresponding model relative to two objects and each sort formula is satisfied relative to a single object. This perspective on frames makes them similar to possible worlds which, too, are taken as relational models according to one formal representation. We will follow the lead of possible world semantics and two-sorted type theory in which possible worlds are objects of a domain D_w (and not relational models) and take frames as elements of a domain D_f of frames (and not as relational models). The link between a frame and the relational structure associated with it is defined

[2] Alternatively, it can be taken as a frame scheme or a frame type. In this case it refers to the set of weddings which have a location and in which the meal is made up by a starter, a main course and a dessert. In the text, a frame depicted is always meant as an instantiated frame in the sense that each node has a particular object as value.

indirectly, again similar to two-sorted type theory. Instead of interpreting attributes as functional relations on $D_o \times D_o$ with D_o the domain of objects comprising both individuals (human beings, chair, dogs etc.) and events (writings, pushings etc.), they are interpreted as ternary relations on $D_f \times D_o \times D_o$. For example, for ATTR an atomic attribute symbol like MEAL or STARTER, $[\![\text{ATTR}]\!]$ is a function that assigns to a frame f a binary relation on D_o s.t. $[\![\text{ATTR}]\!](f)(o)(o')$ is true if o and o' are related by ATTR in f.[3] Similarly, sort formulas are interpreted as (boolean combinations of) binary relations on $D_f \times D_o$. As already mentioned above, this way of relativizing the interpretation of expressions is similar to the way information is made world-dependent in two-sorted type theory. The formal definitions are given next. Let $\Sigma = \langle Sort, Attr \rangle$ be a frame signature of (atomic) sort and attribute symbols, respectively, with $Sort \cap Attr = \emptyset$. The frame language $\mathcal{L}$ based on Σ is defined in (3) and its interpretation is given in (4). The interpretation function $[\![\;]\!]$ assigns to each $\sigma \in Sort$ a binary relation on $D_f \times D_o$ and to each ATTR $\in Attr$ a ternary relation on $D_f \times D_o \times D_o$.

(3) a. $\phi ::= \sigma \,|\, \neg\phi \,|\, \phi_1 \wedge \phi_2$
    b. $\pi ::= \Delta \,|\, \text{ATTR} \,|\, \pi_1 \cap \pi_2 \,|\, \pi_1 \bullet \pi_2 \,|\, \uparrow \phi \,|\, \downarrow \phi$

(4) a. $[\![\sigma]\!](f)(o){=}1$ iff $\langle f, o\rangle \in I(\sigma)$.
    b. $[\![\neg\phi]\!](f)(o){=}1$ iff $[\![\phi]\!](f)(o){=}0$.
    c. $[\![\phi \wedge \psi]\!](f)(o){=}1$ iff $[\![\phi]\!](f)(o){=}1$ and $[\![\psi]\!](f)(o){=}1$.
    d. $[\![\text{ATTR}]\!](f)(o)(o'){=}1$ iff $\langle f, o, o'\rangle \in I(\text{ATTR})$.
    e. $[\![\pi \cap \pi']\!](f)(o)(o'){=}1$ iff $[\![\pi]\!](f)(o)(o'){=}1$ and $[\![\pi']\!](f)(o)(o'){=}1$.
    f. $[\![\pi \bullet \pi']\!](f)(o)(o'){=}1$ iff $\exists o''$:$[\![\pi]\!](f)(o)(o''){=}1$ and $[\![\pi']\!](f)(o'')(o'){=}1$.

The clauses for sort formulas are self-evident. Boolean operations besides $\neg$ and $\wedge$ are defined in the usual way. At the level of relational information, $\bullet$ is sequencing. It is used to built chains of attributes and is defined only if the target sort of the first attribute is a subsort of the source sort of the second attribute (details follow below). The intersection $\cap$ operator is similar to (boolean) conjunction at the level of sortal formulas. It requires that two objects in a frame satisfy both the relation formulas π and π'. Its main use in our frame theory is explained below.

So far, sortal and relational information are not connected with each other. However, as was said above, each attribute has both a source and a target sort. Therefore, one wants to say that the object at the end of a chain π satisfies the sortal information expressed by the sortal formula ϕ (3-a). Similarly, this information should also be expressible for the source sort. It is therefore necessary to go from the relational to the sortal level. This is achieved by two operators $\uparrow$ and $\downarrow$. Formula $\cap\!\uparrow\!\phi$ is true at a triple $\langle f, o, o'\rangle$ if o' satisfies the sortal information ϕ in f, i.e. one has ϕ is true for $\langle f, o'\rangle$. Hence, $\uparrow$ 'projects' a relation in a frame to the second object in this relation and classifies it by the sortal information expressed by its argument. By contrast, $\downarrow$ projects to the first object. The satisfaction clauses are given in (5).

(5) a. $[\![\uparrow \phi]\!](f)(o)(o'){=}1$ iff $[\![\phi]\!](f)(o'){=}1$.
    b. $[\![\downarrow \phi]\!](f)(o)(o'){=}1$ iff $[\![\phi]\!](f)(o){=}1$.

[3] Strictly speaking, it assigns to a frame f a 1-place function as attributes are required to be functional.

Having $\uparrow$ and $\downarrow$ together with $\cap$ allows to express the information that at the end (beginning) of a chain sortal information ϕ holds. This is achieved by relation formulas of the form $\pi \cap \uparrow\phi$ and $\pi \cap \downarrow\phi$. For example, in the wedding frame one has $\textsc{meal} \bullet \textsc{starter} \cap \uparrow\textbf{food}$. This formula expresses that the wedding is related to an object of sort **food** by the chain $\textsc{meal} \bullet \textsc{starter}$. Of course, this information can be made more specific by requiring that the starter is a subsort of **food**, e.g. a (mock turtle) soup: $\textsc{meal} \bullet \textsc{starter} \cap \uparrow\textbf{soup}$. As it stands, we also need to say that the object at the root of a frame satisfies the sortal information ϕ without any additional relational information. This case arises for instance for minimal frames which only contain sortal but no relational information. This kind of information is expressed by means of the null-ary operator Δ. Δ holds of a triple $\langle f, o, o' \rangle$ if one has $o = o'$. Hence, $\Delta \cap \downarrow\phi$ is true in a frame f and objects o and o' if ϕ is true in f for o and o' is identical to o. For example, in the wedding frame $\Delta \cap \downarrow\textbf{wedding}$ is true at the root of the frame. Note that this relation formula does not contain any (chain of) attributes. The satisfaction clause for Δ is given below.

(6) $[\![\Delta]\!](f)(o)(o')$=1 iff o=o'.

The three domains D_f (frames), D_o (objects) and D_w (possible worlds) are related in the following way. First, for each frame f, there is an object $o \in D_o$ about which f contains information. This relation is captured by a function $root$ which assigns to each $f \in D_f$ the object $root(f) \in D_o$. If $root(f) = o$, f is called a frame associated with o. Second, each frame belongs to a possible world $w \in D_w$. This relation is captured by a function *IN* that maps each $f \in D_f$ to the world $\mathit{IN}(f) \in D_w$ to which it belongs.[4] Given these functions, a frame can be taken as a partial description of its root in the world to which the frame belongs.

The relation between a frame and a particular relational structure is defined in terms of a function θ that maps a frame f to the set of relations about which it contains information relative to its referent $root(f)$. Elements of $\theta(f)$ are based on relation formulas $\textsc{attr}_1 \bullet \cdots \bullet \textsc{attr}_n \cap \uparrow\sigma$ for chains of length greater 0 and $\Delta \cap \downarrow\sigma$ for sortal information at the root of the frame. Hence, $\theta(f)$ contains for each chain π in the frame this chain together with sortal information at the end of the chain and sortal information about its root. For example, for the wedding frame above one has $\theta(f_{\mathit{wedding}}) = \{\Delta \cap \downarrow\textbf{wedding}, \textsc{place} \cap \uparrow\textbf{place}, \textsc{meal} \cap \uparrow\textbf{meal}, \textsc{meal} \bullet \textsc{starter} \cap \uparrow\textbf{food}, \textsc{meal} \bullet \textsc{main} \cap \uparrow\textbf{food}, \textsc{meal} \bullet \textsc{dessert} \cap \uparrow\textbf{food}\}$.[5] Due to the use of $\cap, \Delta, \uparrow$ and $\downarrow$ all elements of θ for a frame f are relation formulas and, hence, interpreted as functional relations on $D_f \times D_o \times D_o$. To underline that θ is based on chains in a frame, we write $\pi \in \theta(f)$ whenever there is a $\sigma \in Sort$ such that $\pi \cap \uparrow\sigma \in \theta(f)$. θ is closed both under prefixes

[4] If possible worlds and frames are taken as relational models, the relation between them can be made precise in the following way. Each frame is a particular submodel $\mathcal{M}$ of a possible world $\mathcal{M}_w$. $\mathcal{M}$ is constructed from $\mathcal{M}_w$ as follows. In a first step one forms the reduct $\mathcal{M}'$ of $\mathcal{M}_w$ to the language $\mathfrak{L}$ on which the frame is based. In a second step, one considers the set S of submodels $\mathcal{N}$ of $\mathcal{M}'$ that satisfy the axioms imposed on the frame. A frame is then any minimal model in S. See [NP17] for details.

[5] Though the elements are relations, we write for example $\pi \cap \uparrow\sigma$ instead of $[\![\pi \cap \uparrow\sigma]\!]$ to ease readability.

of attribute chains and supersorts. For closure under prefixes of attribute chains, one has: if $\text{ATTR}_1 \bullet \cdots \bullet \text{ATTR}_n \cap \uparrow\sigma \in \theta(f)$ then $\text{ATTR}_1 \bullet \cdots \bullet \text{ATTR}_{n-1} \cap \uparrow\sigma' \in \theta(f)$ for σ' the target sort (or one of its subsorts) of the attribute ATTR_{n-1}. Closure under supersorts says that if $\pi \cap \uparrow\sigma \in \theta(f)$ and σ' is a supersort of σ, then $\pi \cap \uparrow\sigma' \in \theta(f)$. In the sequel the chains with a supersort of a given sort will not be included if the value of θ is given for a frame f. Frames with the same referent (root) can be ordered according to the information they contain about their common referent. This is captured by the relation $\sqsubseteq$ on the domain of frames.

(7) $f \sqsubseteq f'$ iff $root(f) = root(f')$ and $IN(f) = IN(f')$ and $\forall\pi.(\pi \in \theta(f) \rightarrow \pi \in \theta(f'))$ and if $\pi \cap \uparrow\sigma \in \theta(f)$, and hence $\pi \in \theta(f')$, then there is some σ' with $\pi \cap \uparrow\sigma' \in \theta(f')$ such that $\forall o.\forall o'.[\![\pi \cap \uparrow\sigma']\!](f')(o)(o') \rightarrow [\![\pi \cap \uparrow\sigma]\!](f)(o)(o')))$.

$f \sqsubseteq f'$ holds if f and f' have the same root and belong to the same world. In addition, the information contained in f is a subset of the information contained in f'. This is the case if all chains belonging to $\theta(f)$ also belong to $\theta(f')$ and whenever a pair of objects satisfies a chain of f' that already belongs to f, the pair satisfies the same chain in f as well. The latter condition is necessary to account for the fact that the more specific frame f' may differ from the subsumed frame f by (a) the set of chains and (b) the specificity of the sortal restrictions added to the chains. Implicit in $\sqsubseteq$ is the fact that a frame is a partial description of an object. For example, the wedding frame $f_{wedding}$ at the beginning of this section is a particular element of the hierarchy for frames of sort **wedding**. It does not contain information about the bride or the broom. Adding this information yields a frame $f'_{wedding}$ with more information about the concept 'wedding'. This latter frame is higher in the frame hierarchy since one has $f_{wedding} \sqsubseteq f'_{wedding}$. $\theta(f'_{wedding})$ is $\theta(f_{wedding})$ augmented with chains for the bride and the broom. By contrast, leaving out the chain PLACE results in a less informative frame $f^*_{wedding}$ for which one has $\theta(f^*_{wedding}) = \theta(f_{wedding}) - \{\text{PLACE} \cap \uparrow\textbf{place}\}$. The minimal element in the 'wedding' frame hierarchy has $\theta(f^{min}_{wedding}) = \{\Delta \cap \downarrow\textbf{wedding}\}$. This information ($\Delta \cap \downarrow\textbf{wedding}$) can be further generalized to $\Delta \cap \downarrow\textbf{object}$. Though this information no longer classifies the wedding as a wedding and therefore does not, when taken in isolation, correspond to a frame in the 'wedding' frame hierarchy, it is the minimal frame in the 'object' frame hierarchy.

A second relation between two frames is that of one frame being a subframe of another. Let us illustrate this notion by some examples from the wedding frame. First, the whole wedding frame is a subframe of itself. The frame starting at the MEAL attribute is a subframe of the wedding frame. Let this subframe be f_{meal}. Its information is given by $\theta(f_{meal}) = \{\Delta \cap \downarrow\textbf{meal}, \text{STARTER} \cap \uparrow\textbf{food}, \text{MAIN} \cap \uparrow\textbf{food}, \text{DESSERT} \cap \uparrow\textbf{food}\}$. f_{meal} is the maximal subframe starting at the MEAL attribute. This subframe is a partial description of the meal that was served at the wedding. One of the subframes of this frame is the frame whose only attribute is STARTER which partially describes a meal by saying that it has a starter of sort **food**. These examples show that a subframe is always defined relative to a chain of attributes π corresponding to a relation formula $\pi \cap \uparrow\sigma$. For subframes starting at the root, one has $\pi = \Delta$ and for other subframes π always is of the form $\text{ATTR}_1 \bullet \cdots \bullet \text{ATTR}_n$. This notion, denoted by $\preceq_\pi$, is defined in (8).

(8) $f' \preceq_\pi f$ iff $\pi \in \theta(f)$ and f' satisfies the following conditions with respect to f and π: (a) $\mathit{IN}(f') = \mathit{IN}(f)$. (b) $root(f') = \iota o'.[\![\pi]\!](f)(root(f))(o')$. (c) Let $\theta_\pi(f) = \{\pi' \cap \uparrow\sigma' \mid \pi \bullet \pi' \cap \uparrow\sigma' \in \theta(f)\}$ and let S be a prefix-closed subset of $\theta_\pi(f)$ (i.e., $S \subseteq \theta_\pi(f)$ with if $\pi \bullet \text{ATTR} \in S$, then $\pi \in S$). Then $\theta(f') = S \cup \{\Delta \cap \downarrow\sigma \mid \pi \cap \uparrow\sigma \in \theta(f)\}$.

A frame f' can only be a subframe of f with respect to a chain π if π is a chain in f: $\pi \in \theta(f)$. In order to determine a subframe f' it is sufficient to specify the value of θ for f' and its root and the world it belongs to. A subframe is always required to be in the same world as its superframe: $\mathit{IN}(f') = \mathit{IN}(f)$. The root of f' is the object at the end of the chain π in f: $root(f') = \iota o'.[\![\pi]\!](f)(root(f))(o')$. The object o' is related to a set of o'-rooted frames f'. f' is a subframe of f only if f' contains a subset of the information about o' that f contains about o'. The information contained in f about o' is given by the suffixes π' with sortal information σ' of all chains $\pi \bullet \pi'$ such that $\pi \bullet \pi' \cap \uparrow\sigma' \in \theta(f)$. Let this set be $\theta_\pi(f) = \{\pi' \cap \uparrow\sigma' \mid \pi \bullet \pi' \cap \uparrow\sigma' \in \theta(f)\}$. The information about o' contained in f' is then a subset of $\theta_\pi(f)$ together with the sortal information at the root which is given by $\Delta \cap \downarrow\sigma$. The fact that $\theta(f')$ is required to be only a subset of the corresponding set in f accounts for the fact that the subframe relative to π and f is in general not unique as shown by the example of the two subframes starting at the MEAL attribute above. Since θ is required to be closed under chain-prefixes, it follows that if a chain $\pi' = \text{ATTR}_1 \bullet \cdots \bullet \text{ATTR}_n$ is in $\theta(f')$, then $\pi'' = \text{ATTR}_1 \bullet \cdots \bullet \text{ATTR}_{n-1}$ is in $\theta(f')$ too. The relation $\preceq$ is the union of the $\preceq_\pi$ for π an admissible chain for frames of the sort in question.

The relations $\sqsubseteq$ and $\preceq$ will play a central role in the analysis of bridging inferences below in sections 6 and 7. The interpretation of a bridged expression, for example 'the mock turtle soup' in a context in which a wedding was introduced previously, provides a frame f^* that is required to be a subframe of an extension of the frame for the wedding.

3 Combining Incremental Dynamics and Frame Theory

Our frame theory is integrated in *Incremental Dynamics*, [vE07, Nou03]. In this framework information states are defined as sets of stacks (also called 'contexts') and not as sets of (partial) variable assignments, as it is standardly done in model-theoretic semantics. A stack can be thought of as a function from an initial segment $\{0, \ldots, n-1\}$ of the natural numbers $\mathbb{N}$ to entities of a domain D_o that are stored in the stack. Hence, a stack can equivalently be taken as a sequence of objects $\{\langle 0, d_0\rangle, \ldots, \langle n-1, d_{n-1}\rangle\}$ of length n. If c is a stack, $|c|$ is the length of c. The objects stored in a stack are the *discourse objects*. By $c(i)$ we denote the object at position i at stack c. A link between stack positions and discourse objects that are stored at a position is established by two operations. First, there is a pushing operation:

(9) $c^\frown d := c \cup \{\langle |c|, d\rangle\}$.

Pushing an object d on the stack extends the stack by this element at position $|c|$. The pushing operation will be used in the interpretation of $\exists$, which, in turn is part of the interpretation of determiners (see below for details). The second operation retrieves an object from the stack.

(10) $ret := \lambda i.\lambda c.\iota d.c(i) = d.$

We write $c[i]$ for $ret(i)(c)$. The retrieval operation will become part of the interpretation of common nouns and verbs. For details on Incremental Dynamics without frames see [Nou03] and [vE07].
In the remainder of this section, we are going to incorporate frames into Incremental Dynamics. The first, and most important, modification is related to the type of objects that are stored in a stack. Storing only objects is insufficient to account for bridging inferences. Recall that a bridging inference basically consists in relating information about two objects in the discourse with each other. For example, a mock turtle soup becomes related to a wedding. Viewed from the perspective of our frame theory, in order to establish such an inference it is necessary to look both at the information got in a discourse about an object and, in addition, at possible ways of how this information can be extended. Consider again the following example from the introduction involving a wedding.

(11) I was at a wedding last week. The mock turtle soup was a dream.

After processing the first sentence, a new object o has been introduced in the stack about which one has got the information that it is a wedding. The frame f that contains this information is given by the conditions $\theta(f) = \{\Delta \cap \downarrow\textbf{\textit{wedding}}\}$ and $root(f) = o$. It is a minimal frame of sort **wedding** because there is no relational information linking the wedding to other objects. In order to relate the mock turtle soup to the wedding by a bridging inference, one uses both the information that it is a wedding, got from bottom-up processing, and the conceptual knowledge that weddings can be related to objects of sort **mock turtle soup** by the chain MEAL • STARTER. The former information is given by $\theta(f)$ whereas the latter is given by the frame hierarchy for objects of sort **wedding**. For example, given that a comprehender knows that an object is related to a frame f with $\theta(f) = \{\Delta \cap \downarrow\textbf{wedding}\}$ containing bottom-up information, he applies top-down conceptual knowledge about weddings to infer that the wedding can be related to a mock turtle soup in the way described above. In section 7 we will model these two kinds of information by assigning to a stack position pairs $\langle o, \langle f_o, F_o\rangle\rangle$ consisting of an object o and a pair consisting of a frame f_o containing the information about o got from bottom-up processing, and a set F_o, which is a set of frames each element of which extends f_o along $\sqsubseteq$ in a particular way. For the moment, we will stick to the simpler modelling and take a stack position to be a pair $\langle o, f_o\rangle$ consisting of an object, called the *object component* and an associated frame, called the *frame component*. Such pairs are called *discourse objects*.

Second, the notions of possibility and information state from Incremental Dynamics have to be adapted. We assume that an information state models the epistemic state of a comprehender, i.e. both his (factual) beliefs (knowledge) and his discourse information. This distinction will be represented by defining a possibility as a pair $\langle c, w\rangle$ consisting of a stack c (discourse component) and a world w (factual component). Possible worlds model epistemic uncertainty. An information state is a set of possibilities.

Finally, the lexicon has to be adapted. Since information expressed by common nouns and verbs is always sortal or relational, it is related to frames in our frame theory. This kind of information expresses either that in a frame two objects are related by a chain of attributes or that an object satisfies some sortal information. Hence, frames must become part of the interpretation of these lexical items. This is achieved in the following way.

Common nouns are translated as (atomic) sort expressions whereas the translation of verbs is based on a neo-Davidsonian decompositional analysis. Unary event predicate expressions are translated as (atomic) sort expressions and thematic relation expressions are translated as (atomic) attribute expressions, i.e. as elements of $Attr$. The interpretation of n-ary predicative expressions is lifted in a way similar to Incremental Dynamics without frames (see [vE07] for details). In particular, the type e is replaced by the type of indices ι with variables $i, j, \ldots$. In (12) the interpretations of $\exists$ and common nouns in terms of our discourse and factual components (i.e., stacks and possible worlds) are given. Let $A_w = \{\langle o, f_o\rangle \mid o \in D_o \wedge \mathit{root}(f) = o \wedge \mathit{IN}(f) = w \wedge \theta(f_o) = \{\Delta \cap \downarrow \textbf{object}\}\}$ for $w \in D_w$. That is A_w consists for each world $w \in D_w$ of all pairs $\langle o, f_o\rangle$ with $o \in D_o$ and f_o is the most general frame of o expressing only that o is of sort **object**.

(12) a. $\exists := \lambda s.\lambda s'.\exists\alpha(s = \langle c, w\rangle \wedge s' = \langle c', w\rangle \wedge c' = c^\frown\alpha \wedge \alpha \in A_w)$

b. $\lambda i.\lambda s.\lambda s'.\exists f'(s = \langle c, w\rangle \wedge s' = \langle c', w\rangle \wedge |c'| = |c| \wedge c'[j] = c[j]\ \mathit{for}\ (0 \leq j < |c| \wedge j \neq i) \wedge c[i] = \langle o, f_o\rangle \wedge f_o \sqsubseteq f' \wedge [\![cn]\!](f')(o) \wedge \theta(f') = \theta(f_o) \cup \{\Delta \cap \downarrow cn\} \wedge c'[i] = \langle o, f'\rangle)$.

In (12-a), $\exists$ introduces a new discourse object on the stack. The information associated with this object is the most general one since it is only required to be of sort **object**, which is true of all elements in D_o. This information is subsumed by any further information that is eventually added about the newly introduced object, for example by a head noun. According to (12-b), common nouns are not interpreted as pure tests, which would be their typical analysis in [vE07] but as operations on possibilities. The input and the output possibilities s and s' differ only with respect to position i of their respective discourse components. The semantic contribution of a common noun is to add sortal information. This is modelled by requiring that there is a frame f' that extends f_o ($f_o \sqsubseteq f'$) s.t. o satisfies the sortal information in $f' : [\![cn]\!](f')(o)$, and by adding this sortal information to $\theta(f_o)$ to yield $\theta(f') : \theta(f') = \theta(f_o) \cup \{\Delta \cap \downarrow cn\}$. Finally, f_o is replaced by f' in the output possibility: $c'[i] = \langle o, f'\rangle$. Thus, there is both a test and an update operation associated with the interpretation of a common noun.

4 The approach of Asher & Lascarides 1998

Since our approach is similar in spirit to that of Asher and Lascarides [AL98], we will begin by sketching their approach. Their analysis is based on Chierchia's [Chi95] analysis of definite descriptions as anaphoric. One way of analyzing 'the N' is given in (13-a). Chierchia enriches this meaning by adding a free $n+1$-ary relational constant that is functional in its last argument, (13-b). R links the argument x to an n-tuple of objects $y_1 \ldots y_n$. Functionality requires that given $y_1 \ldots y_n$ x is uniquely determined by

R: $R(y_1, \ldots, y_n, x_1) \wedge R(y_1, \ldots, y_n, x_2) \rightarrow x_1 = x_2$. On this analysis, 'the N' denotes a (unique) N which is related by some dependency relation to an n-tuple $y_1 \ldots y_n$.

(13) a. $\iota x.N(x)$.
b. $\iota x.[R(y_1, \ldots, y_n, x) \wedge N(x)]$.

[AL98] claims that in the case of a bridging inference, R is a binary functional relation B that has to hold between the antecedent object y and the denotation of the definite description x: $B(y, x)$. Hence, lexical semantics provides an underspecified relation B which functions as the bridge or the dependency relation and which must be determined by finding an appropriate value by connecting it to an object in the present discourse context.

How are B and y specified (resolved)? A common strategy is based on coherence relations. Consider e.g. (14).

(14) a. John took engine E1 from Avon to Dansville.
b. He picked up the boxcar (and took it to Broxburn).

Let K_α be a semantic representation of the first sentence, K_β a semantic representation of the second sentence and K_τ a semantic representation of the context in which the second sentence is interpreted so that K_α is a part (subrepresentation) of K_τ. K_α and K_β introduce two events of taking and picking up, respectively. This information is sufficient to defeasibly infer that the two sentences are related by the coherence relation *Narration*. Coherence relations are associated with (non-defeasible) rules that allow to infer additional information about the discourse referents introduced in the three semantic representations. For *Narration* one has: (i) e_α precedes e_β and (ii) if e_α and e_β have the same actor, the location of this actor at the end of e_α is the same as his location at the beginning of e_β. In addition, one has (iii): lexical semantic information about 'pick up' allows the inference that the theme of e_β is located at this location too. Together, (ii) and (iii) yield $In(boxcar, Dansville)$ since Dansville is the location of the actor (i.e. John) at the end of the taking event (= e_α) and the boxcar is the theme of the picking-up event (= e_β). The condition $In(boxcar, Dansville)$ is added to K_β. Resolving B to the function of containment In and assigning y the value $Dansville$, (which is part of an update operation) yields the required bridging inference because a relation between a discourse referent introduced in the first and a discourse referent introduced in the second sentence has been established.

In the above example the derivation of a coherence relation between the two sentences yielded the required bridging inference. However, as noted in [AL98, p.104], often there is not enough information in K_β to infer a particular coherence relation between it and the previous context because K_β contains non-resolved material (B and y) and is therefore underspecified. As a result, B and y must be resolved *before* a coherence relation can be established between the two sentences. The coherence relation is then used as a constraint on the resolution used. The resolution should be such that discourse coherence is maximized (principle 'Maximize Discourse Coherence'). Consider the following variant of example (1-d) from the introduction.

(15) a. John arrived yesterday at 3pm.

b. The camel was outside and needed water.

A possible coherence relation linking the two sentences is *Background*. However, this relation also applies if the second sentence is replaced by its present tense variant 'The camel is outside and needs water'. However, due to the tense shift a bridging inference should not be possible. Thus, a different strategy is needed. First, one uses lexical semantics to infer that 'arrive', being a motion verb, defines a thematic relation 'mode of transport' (besides the theme-relation that has already been introduced during processing the first sentence). Second, one uses world knowledge to infer that camels can be used as such a mode of transport. Both pieces of information are not yet elements of the semantic representations.

(16) a. $\forall e(\mathit{arrive}(e) \rightarrow \exists z.\mathit{Means\text{-}of\text{-}Transport}(e, z))$.
b. $\forall x(\mathit{camel}(x) \rightarrow \mathit{can\text{-}be\text{-}used\text{-}as\text{-}Means\text{-}of\text{-}Transport}(x))$.

Using the additional information in (16), a possible resolution is given by $y = e_{\mathit{arrive}}$ and $B = \mathit{Means\text{-}of\text{-}Transport}$. so that one has $\mathit{Means\text{-}of\text{-}Transport}(e_{\mathit{arrive}}, x)$ with x the camel. This is the required bridging inference because the referent of the definite description is linked to an object introduced in the first sentence. In addition, $\mathit{Means\text{-}of\text{-}Transport}(e_{\mathit{arrive}}, x)$ can be used to infer that the two sentences are related by the coherence relation *Result*. The state of the camel needing water described in the second sentence is the result (or was caused) by the arrival, or, more precisely, by the motion event presupposed by the arriving event.[6] When taken together, one gets a coherent interpretation of (15) because the two sentences are connected by the coherence relation *Result* so that the principle 'Maximize Discourse Coherence' is satisfied.

Let us make the following observations about the second strategy proposed by [AL98]: (a) B is part of the semantic representation of a definite description due to the familiarity constraint imposed by the determiner 'the'. It is therefore independent of any constraints that are imposed related to coherence considerations though it is used to establish a coherence relation in the above example. This strategy fails if the bridged expression is an indefinite like 'a rope' in (2-c) in the introduction since for indefinites a novelty condition rather than a familiarity condition applies. This raises the question where in the semantic representation B and y come from if bridging inferences are not triggered by definite descriptions. (b) A distinction is made between (lexical) semantic properties of an expression that are part of its current semantic representation and properties for which this does not hold. And (c) B is resolved to a property of the latter kind of properties. Observations (b) and (c) already contain one possible answer to the problem raised in the first observation. If B is ultimately (resolved to) a semantic property associated with a (candidate) antecedent object, it should be related to the semantic representation of this object instead of with the semantic representation of the bridged expression. The semantic contribution of the bridged expression to a bridging inference is to provide a value for this property relative to the antecedent object: $B(x) = y$. Definite descriptions are then the special case in which the existence of an appropriate B is required by the semantics of 'the'. On this perspective a bridging inference is trig-

[6] See [Pn97] for a formal analysis in which achievement verbs like 'arrive' are analyzed as boundary events of other, non-boundary events.

gered, in effect, by the antecedent object: there is a semantic property associated with this object the value of which is unknown for this object. Besides being directly applicable to bridged expressions that are not definite descriptions, a second advantage of this perspective is that it does not directly rely on the use of coherence relations and can therefore be applied across the board to bridging inferences. In the remainder of this article, we are going to work out this perspective on bridging inferences.

5 Bridging inferences and 'Questions under Discussion'

If bridging inferences are triggered by the antecedent object, a question to be raised is what happens if B is not already part of the semantic representation of a (candidate) antecedent object. As we have seen in the previous section, a link to coherence relations cannot be the answer because often a bridging inference needs to be done without relying on information provided by these relations. A second strategy to establish coherence between a context and its continuation is based on the notion of a *Question under Discussion* (QuD). According to [KR17], 'in QuD-models of discourse interpretation, clauses cohere with the preceding context by virtue of providing answers to (usually implicit) questions that are situated within a speaker's goal-driven strategy of inquiry.' If an object is introduced into a discourse, this introduction is in general not bare in the sense that no sortal and relational information is associated with it. Initial additional information is given by common nouns for individuals (e.g. it is a wedding) and verbs for events (e.g. 'it is a hitting'). This information can be extended in various ways in the subsequent discourse. However, such extensions are in general not arbitrary but are related to particular questions that are raised in relation to these objects and which depend on the context in which the object is introduced. More generally, one has: In a QuD-model of discourse every newly introduced object raises a set of questions (cf. [RR16]). For objects, i.e. individuals and events introduced by common nouns like 'suicide', these questions are related to possibly complex properties these objects have and, therefore, to sortal and relational information about them. The corresponding rhetorical relation is called *Entity-Elaboration*. If o is the object 'under discussion', the canonical form of an Entity-Elaboration is 'What about o?'. Events that are introduced in the interpretation of verbs raise questions that are related to a particular coherence relation. Examples of questions are 'And then?', 'Why?', 'So what?' and 'How were things like then?', (cf. [RR16]). The relation to coherence between sentences is the following. At each stage τ of a discourse there is a set of active questions related to the objects that have already been introduced into the discourse. An extension of τ with a sentence ϕ is coherent only if this continuation implicitly contains at least one answer to at least one active question raised in τ and thereby automatically links an object already introduced to information provided in the continuation ϕ.

Let's illustrate this with one example from the introduction. In the first sentence of (2-c) Jack and a (planned) suicide are introduced. One therefore gets QuDs that are related to Entity-Elaboration: *What about Jack?* and *What about the (planned) suicide?*. Possible answers are: $\exists y.\exists e.\textit{buy}(e, \textit{jack}, y) \wedge \textit{rope}(y)$, $\exists y.\textit{instrument}(e_s, y) \wedge \textit{rope}(y)$ $\wedge \textit{depressed}(\textit{jack})$. Note that the free variables in the answers refer to objects that are introduced in the first sentence whereas the existentially bound variables are objects that are introduced in the second sentence.

6 Bridging in our frame theory: Bridging relations and QuDs

Let us relate the results of the preceding section to our theory of frames. Information got about an object o by bottom-up processing is stored in the frame component f_o of the discourse object $\langle o, f_o \rangle$. The frame f_o contains at least sortal information which classifies o, e.g. as a wedding or a car. In our frame theory knowing the sort of an object is directly related to knowledge about the frame hierarchy for objects of this sort. The frame f_o is an element of this hierarchy, usually the minimal element (if only sortal information is known). Frames f for which $f_o \sqsubset f$ holds are extensions of f_o in which additional information about o is provided. A frame f that extends f_o by a chain π will be called a π-extension of f_o. This notion is defined in (17).

(17) A frame f^π is a π-extension of a frame f if (a) $\pi \notin \theta(f)$, (b) $f \sqsubset f^\pi$ and (c) for all π' that are not a prefix of π: if $\pi' \in \theta(f^\pi)$ then $\pi' \in \theta(f)$.

Two kinds of π-extensions must be distinguished. Let f^π be a π-extension of f with $\pi \cap {\uparrow}\sigma \in \theta(f^\pi)$, that is σ is the sortal information given in f^π at the end of chain π: (a) If σ is the target sort of π, this information already follows from conceptual knowledge. For example, knowing that o is a wedding, one knows that it took place at a particular location which is of sort **place**. This information is implied by knowledge of the frame hierarchy because if π is admissible for frames associated with objects of a particular sort, then its values are restricted by a particular sortal constraint expressed by its target sort. (b) If σ is not the target sort or if additional information beyond that sort is provided, the information contained in f_o is properly extended in the sense that it is neither implied by f_o nor does it follow from conceptual knowledge. Let us make this distinction between the two kinds of π-extensions explicit by defining π-extensions that do not introduce factual information as *non-factual* π-extensions.

(18) A non-factual π-extension f^π of a frame f is a π-extension of f with $\pi \cap {\uparrow}\sigma \in \theta(f^\pi)$ for which σ is the target sort of π and for each prefix π^p of π, one has $\pi^p \cap {\uparrow}\sigma' \in \theta(f^\pi)$ only if σ' is the target sort of π^p. This latter condition ensures that for prefixes too, no factual information is introduced.

This relationship between bottom-up information and top-down conceptual knowledge suggests the following strategy to model QuDs with frames. A QuD is always related (i) to a discourse object $\langle o, f_o \rangle \in c$ that is already on the stack and (ii) a non-factual π-extension f^π of f_o. Non-factual π-extensions with $\pi \cap {\uparrow}\sigma \in \theta(f^\pi)$ for which the sort σ is the target sort of π are *underspecified* answers to QuDs. A proper (or non-underspecified) answer related to π must provide additional information about this value and is therefore related to a (factual) π-extension f'^π in which new factual information about the value of π is provided so that one has $f^\pi \sqsubset f'^\pi$. How is f'^π related to the linguistic context? f'^π must be related to the bridged expression and therefore to a part of the semantic representation of a sentence ϕ that is a continuation of the stage τ of the current discourse. However, f'^{π} is in general not the frame that is introduced with the semantic representation of the bridged expression as a constituent of ϕ. Rather, this expression introduces a frame f^*. Since f^* provides information about the value of π, it follows that the relation has to be defined in terms of the subframe relation $\preceq$: $f^* \preceq f'^{\pi}$. When taken together, one gets (19).

(19) $\exists f'^{\pi}.f^{\pi} \sqsubset f'^{\pi} \wedge f^* \preceq f'^{\pi}$.

According to (17), a π-extension adds to a frame a chain π (together with its prefixes due to the definition of θ). This accounts for the fact that for a bridging relation of length ≥ 1 the antecedent object is related to a second object. A direct consequence of this definition is that it does not account for cases in which the bridging relation is identity. In this case the information used in the bridging inference is already used when the antecedent object was introduced in the first place. Hence, the bridged expression by itself does, at least in general, not provide new information about the antecedent object. For example, in (1-a) ('Lizzy met a dog yesterday. The dog was very friendly.') the sortal information **dog**, given by $\mathit{\Delta} \cap \downarrow\mathbf{dog}$, is already an element of $\theta(f)$, i.e. the frame information associated with the antecedent object 'a dog' from the first sentence. The information provided by the bridged expression is *given* information relative to the antecedent object so that the frame associated with the bridged expression does not give rise to a proper extension of the frame associated with the antecedent object. Identity has therefore to be treated in a different way. There are at least the following arguments for such a separate treatment. First, if the dependency relation is identity, the antecedent object is always related to itself. Second, the identity relation is possible only with bridged expressions of the form 'the N'. For example, in 'Lizzy met a dog yesterday. A dog was friendly' the two occurrences of 'a dog' cannot refer to the same dog. Third, and most importantly, there is empirical evidence that bridged expressions of the form 'the N' with the bridging relation being identity are processed differently in the brain. [Bur06] found a difference in the P600 effect, an ERP-component, during online semantic processing between the identity relation on the one hand and bridged DPs like 'the engine' and new DPs like 'a rope' on the other hand. Related to these arguments is the following observation. The additional information about the antecedent object is provided by a verbal expression, e.g. 'is friendly' in (1-a). Since the information is related to the same object, the relation between f and f^* can be defined by $\sqsubseteq$ alone (a reference to $\preceq_{\pi}$ is not needed):

(20) $\exists f' : f \sqsubseteq f' \wedge f^* \sqsubseteq f'$.

The above discussion has shown that a distinction has to be made between the bridging relation and a QuD. A QuD always involves a π-extension of the frame associated with the antecedent object. This is the case because new information about this object is provided. The case of a bridging relation is more complex. For establishing coherence, a π-extension is not necessary as shown by bridging inferences based on identity. This difference shows up in the way new information is added. In the case of identity (20) applies which does not require the $\preceq$ relation because no anaphoric relation to a second object is established. We are now ready to define QuDs in our frame theory. QuDs are represented as underspecified answers while answers to QuDs as (specified) answers.

(21) a. A QuD raised by a discourse object $\langle o, f_o \rangle$ is a set of frames F such that each element of F is a non-factual π-extension of f_o for some chain π.

b. An answer to a QuD raised by a discourse object $\langle o, f_o \rangle$ is a frame f^* which stands to f_o in the bridging relation given in (19).

A necessary condition for discourse coherence is defined in (22).

(22) A continuation ϕ of a stage τ of a discourse is coherent relative to τ only if semantic processing ϕ introduces a frame f^* that is an answer to a QuD raised by a discourse object belonging to the discourse component at stage τ.

From the above discussion it follows that in our frame theory the bridging relation B can be defined either at the level of objects or at the level of frames. Let $\langle o', f_{o'} \rangle$ be a discourse object introduced in the continuation ϕ. At the level of objects one gets $[\![\pi]\!](f_o)(o)(o')$; the antecedent object o and the 'bridged' object o' are related by π. At the level of frames, B is defined by a relation between f_o and $f_{o'}$; this relation $Rel_B(f_o, f_{o'})$ holds if there is a π-extension of f_o of which $f_{o'}$ is a subframe. In contrast to defining B at the level of objects, the definition of Rel_B at the frame level is not functional. For a given frame f_o, there are many frames that satisfy the definition. $f_{o'}$ depends on the information that is given about o' in the continuation.

Relating B to a QuD and modelling the latter as underspecified answers has the effect that the dependency relation itself is not underspecified in the sense that it is represented as a free variable as in [AL98]. Underspecification comes in because the value of the chain is constraint only by its target sort, information that is given to a comprehender independently of any discourse information. Having different π-extensions accounts for the fact that answers can involve various relations linking an antecedent object to an object to which it stands in a dependency relation. An example is given in (23).

(23) I took my car for a test drive. The engine/brakes/tyres made a weird noise.

This change of perspective on bridging inferences is made possible due to the shift in location of the bridging (dependency) relation. It is no longer related to the bridged expression but to the semantic representation of a (candidate) antecedent object.

Let us illustrate the above discussion by the following example and the two associated frames for the antecedent object denoted by 'my car' and the bridged object denoted by 'the engine'.

(24) I took my car for a test drive. The engine made a weird noise.

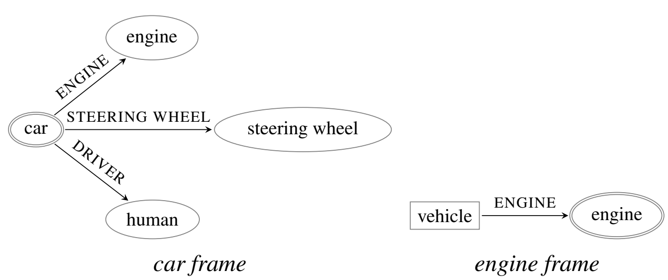

When the car is introduced in the first sentence, only sortal information is provided.[7] Hence, one has $\theta(f_{car}) = \{\Delta \cap {\downarrow}\mathbf{car}\}$. However, using his knowledge about the frame hierarchy, a comprehender knows that there are (frame-)extensions of f_{car} that provide additional information about the car. Three possible extensions are related to the attributes ENGINE, STEERING_WHEEL and DRIVER. One has: $f_{car} \sqsubseteq f_{car}^{\text{ENGINE}}$, $f_{car} \sqsubseteq f_{car}^{\text{STEERING\_WHEEL}}$ and $f_{car} \sqsubseteq f_{car}^{\text{DRIVER}}$. At the level of the function θ one has: $\theta(f_{car}^{\text{ENGINE}}) = \theta(f_{car}) \cup \{\text{ENGINE} \cap {\uparrow}\mathbf{engine}\}$, $\theta(f_{car}^{\text{STEERING\_WHEEL}}) = \theta(f_{car}) \cup \{\text{STEERING\_WHEEL} \cap {\uparrow}\mathbf{steering\_wheel}\}$ and $\theta(f_{car}^{\text{DRIVER}}) = \theta(f_{car}) \cup \{\text{DRIVER} \cap {\uparrow}\mathbf{person}\}$. Each extension corresponds to one possible underspecified answer to the QuD related to the Entity-Elaboration 'What about the car?'.

The frames f_{car}^{ATTR} only contain minimal information about the object o' to which the car is related by $\text{ATTR} \cap {\uparrow}\sigma$. The only information about o' is $\Delta \cap {\downarrow}\sigma$ where σ is the target sort of ATTR. However, the information provided about o' in a continuation will in general be richer because it contains all information got about it in this continuation. For instance, in the example at hand one gets to know that the car emitted a weird noise. Let the discourse object related to 'the engine' in the second sentence be $\langle o', f_{o'}\rangle$. A possible value for $\theta(f_{o'})$ is $\{\Delta \cap {\downarrow}\mathbf{engine}, \text{EMISSION} \cap {\uparrow}\mathbf{sound}, \text{EMISSION} \bullet \text{PITCH} \cap {\uparrow}\mathbf{weird}\}$. As a result, $f_{o'}$ contains more information about the engine of the car than f_{car}^{ENGINE}. $f_{o'}$ is depicted in the figure below.

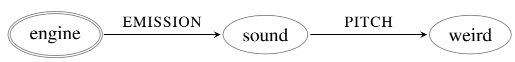

engine subframe

$f_{o'}$ is not a subframe of f_{car}^{ENGINE} because it contains more information about the engine than f_{car}^{ENGINE}. However, there is an extension $f_{car}^{\text{ENGINE}'}$ of f_{car}^{ENGINE} for which $f_{o'} \preceq f_{car}^{\text{ENGINE}'}$ holds: $\theta(f_{car}^{\text{ENGINE}'}) = \theta(f_{car}^{\text{ENGINE}}) \cup \{\text{ENGINE} \bullet \text{EMISSION} \cap {\uparrow}\mathbf{sound}, \text{ENGINE} \bullet \text{EMISSION} \bullet \text{PITCH} \cap {\uparrow}\mathbf{weird}\}$.

7 The formal account: extending frame-based Incremental Dynamic with QuDs

In order to implement the strategy for bridging inferences developed in the preceding section we first have to extend the theory from section 2 with QuDs. Recall that QuDs are always related to a particular discourse object $\langle o, f_o\rangle$ and are modelled as a set F_o of non-factual π-extensions of f_o. One way of integrating F_o into the theory is at the level of stack positions. Instead of storing pairs of the form $\langle o, f_o\rangle$ one stores pairs of the form $\langle o, \langle f_o, F_o\rangle\rangle$. The frame f_o will be called the factual frame component and the set F_o the set of non-factual π-extensions. There are at least two arguments in favour of this option. First, as will be shown below, the bridging constraint is integrated as a part of the normal process of semantic composition and not, say, as part of a more global pragmatic component that operates on given semantic structures. As an effect, F_o must be available locally during the process of semantic composition. The second argument is the context-sensitive character of QuDs. Though they are always raised relative to

[7] We leave out the information related to the verb 'take'.

discourse objects of a particular kind, what counts as an answer to them depends on the information available in the current context. Hence, F_o cannot be an element of the model but must be a component of an information state.

Adding F_o to the information stored about an object o on a stack c is only a first step in modelling the bridging constraint. The second step is to incorporate the bridging inference. Relative to this step two related questions have to be answered: (i) What kind of operation is associated with this step?, and (ii) Where in the process of semantic composition should this operation be placed? Let's begin with the first question for a definite description 'the N'. Recall that a bridging inference operates at two different levels. First, and foremost, it is related to making a text coherent by linking information in a continuation to an object that has already been introduced. Second, a comprehender gets additional information about that object by relating the new information to it. Hence, a bridging inference consists of two operations: establishing a dependency relation (bridging condition) and, if successful, an update operation on the frame component of the antecedent object.

Let's turn to the second question. Establishing the bridging condition (BC) amounts to testing either (19) or (20). This can be done as soon as the head noun has been semantically processed. By contrast, the update operation has to be executed after the verbal element Q has been processed. When taken together, one gets that the test of the bridging condition is done after processing the nominal element P and the update operation is executed after processing Q. In (25) the two test conditions are defined. (26) contains the corresponding update operations and (27) the interpretation of 'the'. The revised interpretation of $\exists$ is given in (27-d). One has: $A_w = \{\langle o, \langle f_o, F_o\rangle\rangle \mid o \in D_o \wedge \mathit{root}(f) = o \wedge \mathit{IN}(f) = w \wedge \theta(f_o) = \{\mathit{\Delta}\cap\downarrow\textbf{object}\} \wedge F_o \subseteq \{f \mid f \textit{ is a non-factual } \pi\textit{-extension of } f_o\}\}$ for $w \in D_w$.[8] π^1 and π^2 are the first and second projection functions for pairs.[9]

(25) a. $BC_1 := \lambda f.\lambda j.\lambda i.\lambda s.\lambda s'.\exists c.\exists o.\exists o'.\exists f_o.\exists F_o.\exists f_{o'}.\exists F_{o'}.$
$(s{=}s' \wedge \pi^1(s){=}c \wedge |\pi^1(s)| = i \wedge c[i]{=}\langle o', \langle f_{o'}, F_{o'}\rangle\rangle \wedge j < i \wedge$
$c[j]{=}\langle o, \langle f_o, F_o\rangle\rangle \wedge f_o \sqsubseteq f \wedge f_{o'} \sqsubseteq f)$.

b. $BC_2 := \lambda f.\lambda j.\lambda i.\lambda s.\lambda s'.\exists c.\exists o.\exists o'.\exists f_o.\exists F_o.\exists f_{o'}.\exists F_{o'}.\exists f^{\pi}_o.$
$(s{=}s' \wedge \pi^1(s){=}c \wedge c[i]{=}\langle o', \langle f_{o'}, F_{o'}\rangle\rangle \wedge j < i \wedge c[j]{=}\langle o, \langle f_o, F_o\rangle\rangle \wedge f^{\pi}_o \in F_o \wedge f^{\pi}_o \sqsubset f \wedge f_{o'} \preceq_{\pi} f)$.

(26) a. $\mathit{Update}_1 := \lambda f.\lambda j.\lambda i.\lambda s.\lambda s'.\exists c, c'.\exists o, o'.\exists f_o.\exists f^{\pi}_o.\exists F_o.\exists f_{o'}.\exists F_{o'}.\exists F'_o.$
$(\pi^2(s){=}\pi^2(s') \wedge \pi^1(s){=}c \wedge \pi^1(s'){=}c' \wedge c \approx_j c' \wedge c[i]{=}\langle o', \langle f_{o'}, F_{o'}\rangle\rangle \wedge$
$j < i \wedge c[j]{=}\langle o, \langle f_o, F_o\rangle\rangle \wedge f^{\pi}_o \in F_o \wedge f_o \sqsubseteq f^{\pi}_o \wedge f_o \sqsubseteq f \wedge f_{o'} \sqsubseteq$

[8] Requiring that F_o be a subset of the non-factual π-extensions of f_o raises the question of how this set can be further restricted. In general we taken the determination of the initial F_o to be context-specific, based on probabilities. We will come back to this question at the end of this section.

[9] Recall from section 2 and 3 that an information state is a set of possibilities and that a possibility is a pair $\langle c, w\rangle$ consisting of a stack c and a world w. In contrast to section 3 our stack elements are now pairs $\langle o, \langle f_o, F_o\rangle\rangle$ with object o, its frame f_o and a set of π-extensions F_o. In (25)-(27) o is used for objects, f for frames, F for sets of frames, c for stacks, i, j for stack indices and s for possibilities. Note that while π is used for chains of attributes, π^1 and π^2 denote the projection function.

$f \wedge c'[j]=\langle o, \langle f, F'_o\rangle\rangle \wedge \theta(f)=\theta(f_o) \cup \theta(f_{o'}) \wedge F'_o=F_o-\{f' \mid f' \in F_o \wedge f' \sqsubseteq f\})$.

b. *Update*$_2 := \lambda f.\lambda j.\lambda i.\lambda s.\lambda s'.\exists c, c'.\exists o, o'.\exists f_o.\exists f_o^\pi.\exists F_o.\exists f_{o'}.\exists F_{o'}.\exists F'_o.$ $(\pi^2(s)=\pi^2(s') \wedge \pi^1(s)=c \wedge \pi^1(s')=c' \wedge c \approx_j c' \wedge c[i]=\langle o', \langle f_{o'}, F_{o'}\rangle\rangle \wedge j < i \wedge c[j]=\langle o, \langle f_o, F_o\rangle\rangle \wedge f_o^\pi \in F_o \wedge f_o \sqsubseteq f_o^\pi \wedge f_o \sqsubseteq f \wedge f_{o'} \preceq_\pi f \wedge c'[j]=\langle o, \langle f, F'_o\rangle\rangle \wedge \theta(f)=\theta(f_o) \cup \{\pi \bullet \pi' \mid \pi' \in \theta(f_{o'})\} \wedge F'_o=F_o - \{f' \mid f' \in F_o \wedge f' \sqsubseteq f\})$.

c. $c \approx_i c' := |c|=|c'| \wedge \forall j(0 \leq j < |c| \wedge j \neq i \rightarrow c'[j]=c[j])$.[10]

(27) a. $[\![the]\!] := \lambda P.\lambda Q.\lambda s.\exists f.\exists j.(\exists \cdot P(|\pi^1(s)|) \cdot [BC_1(f)(j)(|\pi^1(s)|) \cdot Q(|\pi^1(s)|) \cdot \mathit{Update}_1(f)(j)(|\pi^1(s)|) \cup BC_2(f)(j)(|\pi^1(s)|) \cdot Q(|\pi^1(s)|) \cdot \mathit{Update}_2(f)(j)(|\pi^1(s)|)])(s)$.

b. $\phi \cdot \psi := \lambda s.\lambda s'.\exists s''(s'' \in \phi(s) \wedge s' \in \psi(s''))$.[11]

c. $\phi \cup \psi := \lambda s.\lambda s'.s' \in \phi(s) \vee s' \in \psi(s)$.

d. $\exists := \lambda s.\lambda s'.\exists \alpha(s = \langle c, w\rangle \wedge s' = \langle c', w\rangle \wedge c' = c^\frown \alpha \wedge \alpha \in A_w)$

The tests of the bridging conditions $BC_{1/2}$ in (25) are part of the interpretation of the determiners 'the' and 'a'. They therefore introduce the 'bridged' object if they are a constituent of a bridged expression. Their semantic function is to test for the bridging relations (19) and (20). BC_1 corresponds to (20) and therefore tests on the identity (or a sort subsumption) relation whereas BC_2 corresponds to (19) and is thus related to bridging involving a relation other than identity. For BC_1 the 'bridged' object o' is stored at position $i = |\pi^1(s)|$, i.e. at the last position of the stack since it has just been introduced. Its associated factual frame is $f_{o'}$ and the set of non-factual π-extensions is $F_{o'}$. In order for (20) to be satisfied, there has to be an antecedent object o that has already been introduced so that it is stored at a position j preceding $i = |\pi^1(s)|$. Recall that for (20) no non-factual π-extensions are used. Hence, the set F_o of non-factual π-extensions associated with o at j does play no role. All that is required is that there is a frame f that extends both the factual frame component associated with o' and that associated with o: $f_o \sqsubseteq f \wedge f_{o'} \sqsubseteq f$. The constraint that there is a factual π-extension providing new factual information is built into the update operation in (26). The difference between BC_1 and BC_2 consists in the bridging relation. Since for BC_2 this relation is (19), establishing this relation always involves a non-factual π-extension since the antecedent object o is related to another object o' by a chain of attributes π. Hence, one has that a non-factual π-extension f_o^π belonging to F_o must be extended by f and the factual frame component $f_{o'}$ associated with the 'bridged' object has to be a subframe of f : $f_o^\pi \sqsubset f \wedge f_{o'} \preceq_\pi f$.

The two update operations *Update*$_{1/2}$ in (26) operate on the frame component of the antecedent object because this component has to be updated due to the new information provided by the 'bridged' object. *Update*$_1$ is used for bridging inferences involving the identity relation. This operation therefore corresponds to BC_1. By contrast, *Update*$_2$ applies to bridging inferences where the antecedent object is related to a second object and, hence, BC_2 is used. In *Update*$_1$ the updated frame f for the an-

[10] $c \approx_i c'$ says that the stacks c and c' differ at most w.r.t. the value assigned to position i.

[11] In the definitions of $\cdot$ and $\cup$, ϕ and ψ map possibilities (i.e. pairs $\langle c, w\rangle$ consisting of a stack and a world) to sets of possibilities.

tecedent object has to satisfy three conditions: (i) it extends the 'old' frame f_o: $f_o \sqsubseteq f$; (ii) it extends the frame $f_{o'}$ associated with the bridged DP: $f_{o'} \sqsubseteq f$; and it extends a non-factual π-extension: $f_o^\pi \sqsubseteq f$. Let us illustrate these conditions with the example of Lizzy and the dog: 'Lizzy met a dog yesterday. The dog was very friendly.' One has: $\theta(f_o) = \{\Delta \cap \downarrow\mathbf{dog}\}$, $\theta(f_{o'}) = \{\Delta \cap \downarrow\mathbf{dog}, \text{BEHAVIOUR} \cap \uparrow\mathbf{friendly}\}$[12] and $\theta(f_o^{\text{BEHAVIOUR}}) = \{\Delta \cap \downarrow\mathbf{dog}, \text{BEHAVIOUR} \cap \uparrow\mathbf{behaviour}\}$. $f_{o'}$ is the frame that results after processing the VP. For the final updated frame f we have that $\theta(f)$ is the union of $\theta(f_o)$ and $\theta(f_{o'})$: $\theta(f) = \{\Delta \cap \downarrow\mathbf{dog}\} \cup \{\Delta \cap \downarrow\mathbf{dog}, \text{BEHAVIOUR} \cap \uparrow\mathbf{friendly}\} = \{\Delta \cap \downarrow\mathbf{dog}, \text{BEHAVIOUR} \cap \uparrow\mathbf{friendly}\}$. The way $\theta(f)$ is construed ensures that it is the minimal frame satisfying the three conditions. In this particular case the (updated) frame f is identical to f_o: This follows from the fact that the original information about the dog got in the first sentence is minimal, only sortal information is provided, and from the fact that the bridging relation is identity. As a result, the information got in the second sentence is still about o and repeats the sortal information from the first sentence. Furthermore, this example shows that the update operation does not simply take the frame f that passed the bridging constraints $BC_{1/2}$ as the new factual frame component of the antecedent object. This frame has in addition to comprise the information got about the dependent object o' from processing the verbal element. The new non-factual π-extensions component F_o' is the old one minus those non-factual π-extensions in this set that are subsumed by f because the corresponding QuDs have been answered.

Update$_2$ differs from *Update*$_1$ in the way the updated frame f is related to the new information provided by $f_{o'}$. Since the antecedent object and the object denoted by the bridged DP are not identical and are therefore related by a chain of attributes with length greater 0, $f_{o'}$ cannot be extended by f but has to be a subframe of f. Let us illustrate this with the car example: 'I took my car for a test drive. The engine made a weird noise'. One has: $\theta(f_o) = \{\Delta \cap \downarrow\mathbf{car}\}$, $\theta(f_{o'}) = \{\Delta \cap \downarrow\mathbf{engine}, \text{EMISSION} \cap \uparrow\mathbf{sound}, \text{EMISSION}\bullet\text{PITCH}\cap\uparrow\mathbf{weird}\}$ and $\theta(f_o^{\text{ENGINE}}) = \{\Delta\cap\downarrow\mathbf{car}, \text{ENGINE}\cap\uparrow\mathbf{engine}\}$. The frame $f_{o'}$ is the frame for the engine after processing the VP in the second sentence. For the updated frame f, $\theta(f)$ is construed as follows. The set $\theta(f_o) = \{\Delta \cap \downarrow\mathbf{car}\}$ is extended by chains $\pi \bullet \pi'$ where π is given by f_o^π: ENGINE and π' is an element from $\theta(f_{o'})$. Since there are three elements, one gets the chains $\text{ENGINE} \bullet (\Delta \cap \downarrow\mathbf{engine})$, $\text{ENGINE} \bullet \text{EMISSION} \cap \uparrow\mathbf{sound}$ and $\text{ENGINE} \bullet \text{EMISSION} \bullet \text{PITCH} \cap \uparrow\mathbf{weird}$. Since $\text{ENGINE} \bullet (\Delta \cap \downarrow\mathbf{engine})$ has the same satisfaction conditions as $\text{ENGINE} \cap \uparrow\mathbf{engine}$), one gets $\theta(f) = \{\Delta\cap\downarrow\mathbf{car}, \text{ENGINE}\cap\uparrow\mathbf{engine}, \text{ENGINE}\bullet\text{EMISSION}\cap\uparrow\mathbf{sound}, \text{ENGINE}\bullet \text{EMISSION}\bullet\text{PITCH}\cap\uparrow\mathbf{weird}\}$. Similar to the way f is construed in the case of an identity relation, the construction of f ensures that f is the minimal frame satisfying the three conditions. The new non-factual π-extensions component is the old one minus those non-factual π-extensions in this set that are subsumed by f because the corresponding QuDs have been answered. This is again similar to the case of the first update operation.

Finally, we turn to the interpretation of the definite and indefinite determiner. Processing the definite determiner 'the' consists of two branches (using the choice operation) after processing the head noun P (27). In the first branch BC_1 succeeds followed

[12] Recall that the value of θ for a frame f is closed under supersorts. Hence, $\theta(f_{o'})$ is, in effect the set $\{\Delta\cap\downarrow\mathbf{dog}, \text{BEHAVIOUR}\cap\uparrow\mathbf{friendly}, \text{BEHAVIOUR}\cap\uparrow\mathbf{behaviour}\}$. This set is a superset of the set $\theta(f_o^{\text{BEHAVIOUR}})$ given next.

by the interpretation of the verbal element Q and the update operation *Update*$_1$. In the second branch BC_2 succeeds followed by the interpretation of the verbal element and the update operation *Update*$_2$. On this interpretation of 'the', a definite description always is a bridging expression for which the bridging constraint has to be satisfied. One can therefore say that it 'signals' that there is a relation to the previous context.

For the determiner 'a', only bridging condition BC_2 applies, as shown above in section 6. Furthermore, this condition is only a *sufficient* condition to ensure discourse coherence. The bridging condition can equally be satisfied by another frame $f_{o''}$ introduced in the continuation. Hence, for 'a', both BC_2 and the update operation *Update*$_2$ must be optional. We model this by replacing the first branch in the interpretation of 'a' by a branch that only executes P and Q without any test or update operation.

(28) a. $[\![a]\!] := \lambda P.\lambda Q.\lambda s.\exists f.\exists j.(\exists \cdot P(|\pi^1(s)|) \cdot [Q(|\pi^1(s)|) \cup BC_2(f)(j)(|\pi^1(s)|) \cdot Q(|\pi^1(s)|) \cdot \mathit{Update}_2(f)(j)(|\pi^1(s)|)])(s).$

If the update operation *Update*$_2$ is not obligatory for indefinites, the following problem can arise. Processing a continuation can be successful, i.e. there is a (non-empty) output information state, without successfully checking the BC for at least one possibility. It is therefore necessary to explicitly test for this satisfaction. One way of doing this is during the combination of two sentences. There are at least two ways of how this testing can be done: at the level of the first (factual) frame component or on the second, QuD-related frame component. We will choose the first option. If a continuation contains information about a discourse object α that is already on the stack at some position i in some possibility s of the output of the first sentence, then the factual frame component f'_o at position i in a successor possibility s' of s must be a proper frame extension of the frame component f_o at position i in s: $f_o \sqsubset f'_o$. The notion of a successor possibility is defined in (29).

(29) a. $s \trianglelefteq s' := \exists c.\exists w.\exists c'.\exists w'.s = \langle c, w\rangle \wedge s' = \langle c', w'\rangle \wedge w = w' \wedge \exists c'' : c' = c^\frown c'' \wedge \forall i : 0 \leq i < |c| \rightarrow c[i] \lessdot_i c'[i].$

b. $c[i] \lessdot_i c'[i] := \exists o.\exists o'.\exists f.\exists f'.\exists F.\exists F'.c[i] = \langle o, \langle f, F\rangle\rangle \wedge c'[i] = \langle o', \langle f', F'\rangle\rangle \wedge o = o' \wedge f \sqsubseteq f' \wedge F \supseteq F').$

A possibility s' is a successor of a possibility s if they share the same world component and, therefore, contain information about objects and frames in the same world. Furthermore, s' possibly extends the discourse information of s in the following respects. First, it can contain information about more objects: $\exists c'' : c' = c^\frown c''$. Second, w.r.t. to the discourse objects in the discourse component c in s one has: the same objects are stored in the respected positions. For the frame components, one has that s' contains at least the information that s contains about the stored objects. The factual frame components are related by $\sqsubseteq$ and the QuD-component by $\supseteq$.It is not necessary that a successor possibility s' of s properly extends the information in s about a discourse objects. In this case one has $f = f'$ and $F = F'$. The bridging condition BC_test, defined in (30), captures the constraint of a proper extension for the factual frame component.

(30) $BC\_test(s, s') := s \trianglelefteq s' \wedge \exists i : 0 \leq i < |\pi^1(s)| \wedge \pi^1(\pi^2([\pi^1(s)](i))) \sqsubset \pi^1(\pi^2([\pi^1(s')](i))).$

Each successor s' of a possibility s in the input information state must properly extend the information associated with at least one discourse object that is an element of the stack in s. The requirement that *each* successor possibility has to satisfy the bridging constraint (for at least one position) is necessary because a comprehender does not know in advance which of these successors will eventually be eliminated. If BC_test is added to the definition of combining two sentences, one gets the required global test of discourse coherence.

(31) $\phi \cdot_D \psi := \lambda s.\lambda s'.\exists s''(\phi(s)(s'') \wedge \psi(s'')(s') \wedge BC\_test(s'', s'))$.

Note that $\cdot_D$ is different from $\cdot$ defined in (27-b) above. $\cdot$ is used at the lexical level to combine constituents of sentences that are built from $\exists$, dynamic properties, $BC_{1/2}$ and *Update*$_{1/2}$. By contrast, $\cdot_D$ is used at the discourse level to combine sentences.

A final question we need to address, is 'How is the set F determined?' Simply assuming that F consists of all (non-factual) π-extensions of the current factual frame component yields a set that is likely to be infinite. Relating F to the notion of prediction provides a possible way of analysing how F can be restricted to a proper subset of all possible non-factual π-extensions. In [NP17] we present an account that bases predictions on probabilities. The key idea is to define for each position i on a stack a probability measure Pr_i on subsets of the range of θ, i.e. relations on $D_f \times D_o \times D_o$. Expectations are ranked in such a way that pre-activation is restricted to those extensions whose probability exceeds a particular value.[13]

8 Conclusion

In this article we have developed a theory of bridging inference in frame theory. Using frames, bridging inferences can be modelled as update operations involving frames. In contrast to previous approaches, no 'incompleteness' in form of free variables is needed. Rather, incompleteness is replaced by underspecification. Following models of QuDs and results of neurophysiological research on predictions during semantic processing in the brain, each discourse object is related to a set of possible ways of how information about this object can be extended by a continuation of the discourse. Extensions are based on a particular chain of attributes and on knowledge of the frame hierarchy associated with objects of a particular sort. These extensions are underspecified in the sense that except for the constraint imposed by the target sort nothing is known in the discourse about the value of the chain. The bridging inference consists in relating one extension with a frame that is introduced in a continuation of the discourse. The implicit character of bridging inferences shows up in the fact that establishing and testing for them is modelled by separate update operations that, by themselves, are not needed in the process of semantically combining the constituents of a sentence and/or a discourse. The difference between definite descriptions and indefinites lies in the way they are related to QuDs/predictions. Whereas definite descriptions *always* discharge a bridging constraint and therefore ensure discourse coherence, this is only a possibility

[13] Therefore, the use of default logic in [AL98] and weighted abduction in [HSAM93] is replaced by probability measures on frame hierarchies.

for indefinites. They can but need not be related to a previously introduced discourse object by a bridging (dependency) relation.

References

[AL98] Nicholas Asher and Alex Lascarides. Bridging. *Journal of Semantics*, 15:83–113, 1998.

[Bur06] Petra Burkhardt. Inferential bridging relations reveal distinct neural mechanisms: Evidence from event-related brain potentials. *Brain and Language*, 98(2):159 – 168, 2006.

[CH77] H. H. Clark and S. E. Haviland. Comprehension and the Given-New Contract. In R. O. Freedle, editor, *Discourse Production and Comprehension*, pages 1–40. Ablex Publishing, Hillsdale, N. J., 1977.

[Cha83] Eugene Charniak. Passing markers: A theory of contextual influence in language comprehension. *Cognitive Science*, 7:171 – 190, 1983.

[Chi95] Gennaro Chierchia. *Dynamics of Meaning: Anaphora, Presupposition and the Theory of Grammar*. University of Chicago Press, Chicago, 1995.

[Cla77] Herbert Clark. Bridging. In Philip Johnson-Laird and Peter Wason, editors, *Thinking: Readings in Cognitive Science*, pages 411 – 420. Cambridge University Press, Cambridge, 1977.

[Geu11] Bart Geurts. Accessibility and anaphora. In Klaus von Heusinger, Claudia Maienborn, and Paul Portner, editors, *Semantics*, volume 2 of *Handbooks of Linguistics and Communication Science*, chapter 75, pages 1988–2011. DeGruyter, 2011.

[HSAM93] Jerry R. Hobbs, Mark E. Stickel, Douglas E. Appelt, and Paul A. Martin. Interpretation as abduction. *Artif. Intell.*, 63(1-2):69–142, 1993.

[KR17] Andrew Kehler and Hannah Rohde. Evaluating an expectation-driven question-under-discussion model of discourse interpretation. *Discourse Processes*, 54(3):219–238, 2017.

[Nou03] Rick Nouwen. *Plural pronominal anaphora in context*. PhD thesis, Netherlands Graduate School of Linguistics Dissertations, LOT, Utrecht, 2003.

[NP17] Ralf Naumann and Wiebke Petersen. A formal theory of frames with an application to plural quantification and cross-sentential anaphora. submitted to Journal of Semantics, 2017.

[Pn97] Christopher Piñón. Achievements in an event semantics. In Aaron Lawson, editor, *Proceedings SALT VII*, pages 276–293, Ithaca, NY: Cornell University, 1997.

[RR16] Uwe Reyle and Arndt Riester. Joint information structure and discourse structure analysis in an underspecified DRT framework. In Julie Hunter, Mandy Simons, and Matthew Stone, editors, *Proceedings of the 20th Workshop on the Semantics and Pragmatics of Dialogue (JerSem)*, pages 15 – 24, New Brunswick, NJ, USA, 2016. Rutgers University.

[vE07] Jan van Eijck. Context and the composition of meaning. In Harry Bunt and Reinhard Muskens, editors, *Computing Meaning*, pages 173–193. Springer Netherlands, Dordrecht, 2007.