Parsimonious Semantic Representations with Projection Pointers

Noortje J. Venhuizen
University of Groningen
n.j.venhuizen@rug.nl

Johan Bos
University of Groningen
johan.bos@rug.nl

Harm Brouwer
University of Groningen
harm.brouwer@rug.nl

Abstract

The influential idea by van der Sandt (1992) to treat presuppositions as anaphora in the framework of Discourse Representation Theory (DRT, Kamp and Reyle, 1993) has inspired a lot of debate as well as elaborations of his account. In this paper, we propose an extension of DRT, called Projective DRT, which adds pointers to all DRT referents and conditions, indicating their projection site. This means that projected content need not be moved from the context in which it is introduced, while it remains clearly discernible from asserted content. This approach inherits the attractive properties from van der Sandt's approach to presupposition, but precludes a two-step resolution algorithm by treating projection as variable binding, which increases compositionality and computational efficiency. The result is a flexible representational framework for a descriptive theory of projection phenomena.

1 Introduction

When it comes to presupposition projection, or more general 'projection phenomena', there seems to be some unpleasant friction between neat compositional approaches to discourse representation, and empirically driven theories. A case in point is Discourse Representation Theory (DRT), in which proper names are treated with a special procedure in order to account for their availability as antecedent for subsequent anaphora (Kamp and Reyle, 1993). This behaviour is due to the *projective* nature of proper names, that is, their existential indifference to logical operators such as negation and conditionals. In van der Sandt's (1992) empirically-driven theory of presupposition projection, formalized in the DRT framework, this discrepancy between compositionality and empirical soundness becomes very clear: presuppositions are only resolved in a second stage of processing by moving them from an embedded context to their context of interpretation. In purely compositional accounts of DRT, on the other hand, treatment of projection phenomena is usually simply left out (Muskens, 1996).

The goal of this paper is to investigate whether van der Sandt's idea to treat presuppositions in the same way as anaphora can be generalized to account for other projection phenomena, such as Potts's (2005) conventional implicatures, in a more compositional manner. To this purpose, we propose a representational extension of DRT, called Projective DRT (PDRT), that deals with presuppositions and other projection phenomena without moving semantic material within the representation. The approach is a simplification of Layered DRT, as proposed by Geurts and Maier (2003), since presuppositions and asserted content are treated on the same level. In PDRT, projection is represented by assigning variables ranging over DRSs, just as anaphora in dynamic frameworks are dealt with by assigning variables ranging over entities. This results in semantic representations that are close to the linguistic surface structure, while clearly distinguishing between asserted and projected content.

This paper is organized as follows. First, a theoretical background on projection phenomena in DRT is provided, focusing on van der Sandt's (1992) approach to presuppositions. In Section 3 we introduce Projective DRT, describing its preliminaries and how it deals with different types of (projective) content. The interpretation of PDRT is given via a translation to standard DRT, described in Section 4. Finally, Section 5 presents the conclusion and indicates directions for future work, describing an ongoing effort to implement PDRT into a large corpus of semantically annotated texts: the Groningen Meaning Bank (Basile et al., 2012).

2 Background

Presuppositions have a long history in the formal semantics and pragmatics literature (see, e.g., Beaver and Geurts, 2011, for an overview). In this paper, we focus on a specific representational theory of presuppositions based on Discourse Representation Theory (DRT; Kamp and Reyle, 1993).

2.1 Presuppositions as anaphora

In the theory of van der Sandt (1992), presupposition projection is treated on a par with anaphora resolution. This approach is motivated by the observation that presuppositions and anaphora display similar behaviour, since they both project their content from the scope of entailment-cancelling operators and show a preference for binding to an accessible antecedent. Unlike anaphora, however, presuppositions can occur felicitously in contexts where no suitable antecedent can be found. In these cases a new antecedent is created at an accessible discourse level, a process that has been called '*accommodation*'. The framework used by van der Sandt to implement his theory is DRT. In this account, each DRS is associated with a so-called A-structure, in which all presuppositions of that DRS are collected. In a second stage of processing, these presuppositions are resolved by either *binding* them to earlier introduced discourse referents or *accommodating* them at a suitable level of discourse. Presupposition resolution is secured by applying several constraints that determine relative preferences between alternative interpretations. These constraints include, for example, that binding is preferred over accommodation, and that global accommodation is preferred over local accommodation (see also Geurts, 1999).

An example of the working of van der Sandt's algorithm is shown in (1) (the A-structure introduced by the presuppositional content is indicated by a dashed box). In the unresolved representation in (1a) the presupposition triggered by the definite description "the cyclist" occurs in the A-structure at the introduction site. In the second stage of processing, this A-structure is resolved by accommodating the presupposition in the global DRS, resulting in the representation shown in (1b).

(1)   Someone did not see the cyclist.
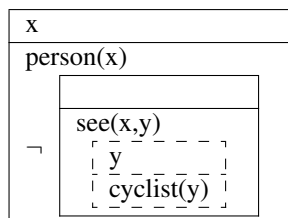
　　　a.　Unresolved DRS:　　　　　　　　　b.　Resolved DRS:
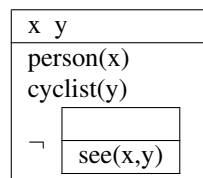
One of the main issues with van der Sandt's analysis of presuppositions in DRT is that, after presupposition projection, accommodated presuppositions and asserted content are indistinguishable. For example, in (1b) the accommodated presupposition "the cyclist" is added to the global context and therefore obtains the same status as the asserted content introduced by "someone". Krahmer (1998) argues, following Kracht (1994), that accommodated presuppositions should maintain their *presupposition-hood* because they are interpreted different from asserted content. For example, falsehood of a presupposition, also called *presupposition failure*, makes the sentence in which it occurs undefined (as in "The king of France is bald", where the existence of a king of France is presupposed), while in the case of falsely asserted content, the sentence is simply false (as in "France is a monarchy"). Moreover, given a compositional approach to semantics, we have to take into account that accommodated presuppositions may become bound later on, when more information of the surrounding context becomes available. This is not the case for asserted content, which implies that at each stage of processing these types of content should be distinguishable.

In order to resolve this issue, Krahmer (1998) introduces a marker for presuppositional content, such that presuppositions are accommodated at a higher discourse level *as presuppositions*, allowing for an

interpretation distinct from asserted content. While this increases compositionality, the presupposition is still moved away from its introduction site in case of accommodation, which makes it difficult to retrieve the linguistic surface structure. This is problematic for applications such as surface realisation – text generation from semantic representations – and for the treatment of phenomena that depend on this surface structure, such as factive constructions and VP-ellipsis. Introducing yet another marker to identify the introduction site of a presupposition would clutter the representation and severely reduce readability and computational efficiency. Another issue with this approach is that recently the property of projection has been associated with a wider range of linguistic expressions outside of presuppositions (see Simons et al., 2010, for an overview). An important example are conventional implicatures (CIs), as described by Potts (2005). An example of a CI is shown in (2) (adapted from Potts, 2005).

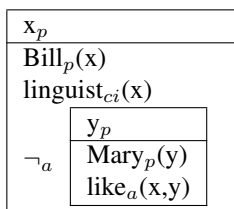(2)   It is not true that Lance Armstrong, an Arkansan, won the 2002 Tour de France.

The conventional implicature triggered in the appositive (that Lance Armstrong is an Arkansan), is projected from out of the scope of the negation, just like the presupposition triggered by the proper name. However, CIs show a different projective behaviour than presuppositions, since they have a strong resistance against binding to an antecedent. This is explained by the observation that they intuitively convey 'new' information, like asserted content. This preference for accommodation contrasts with the theoretical assumptions of van der Sandt (1992) and Krahmer (1998), who implement accommodation as a repair strategy.

In sum, we need a single representational framework that allows for a separate treatment of asserted and projected content. An important step in this direction is Layered DRT (Geurts and Maier, 2003) where different types of information are treated on different layers. We will show that although this representation accounts for the differences between asserted, presupposed and conventionally implied content, it fails to capture their similarities and interactions.

2.2   Layered DRT

In Layered DRT (LDRT), the distinction between different types of information is implemented by introducing different layers (Geurts and Maier, 2003). Each discourse referent and condition is associated with a set of labels that indicate the layers on which the information is interpreted. These layers allow for a distinction between asserted and presupposed content, but also for a separate interpretation of implicated, indexical and formal content. An example is shown in (3), where the label p indicates presupposed content, the label a implicates asserted content and ci indicates a conventional implicature.

(3)   Bill, a linguist, does not like Mary.

$$
\begin{array}{|l|}
\hline
x_p \\
\hline
\text{Bill}_p(x) \\
\text{linguist}_{ci}(x) \\
\neg_a \quad \begin{array}{|l|} \hline y_p \\ \hline \text{Mary}_p(y) \\ \text{like}_a(x,y) \\ \hline \end{array} \\
\hline
\end{array}
$$

This example shows that the different types of content are represented within a single framework, while being clearly distinguishable through the labels. The different layers are connected by sharing discourse referents, indicating the interaction between different types of content. Since all conditions are indexed with a label, projected material can remain at its introduction site, because it is interpreted at a separate layer and therefore it is not targeted by logical operators. The interpretation of LDRT is defined on the basis of the truth-conditional content of sets of layers. For example, the presupposed meaning of (3) is true in the set of worlds in which the individuals called Bill and Mary exist. The asserted content can only be defined in combination with the presupposed content, representing the set of worlds in which Bill does not like Mary.

Although LDRT nicely captures the differences and dependencies between the various types of information, the separation into different layers comes at a cost. Firstly, it is unclear under which conditions a new layer is created. According to Geurts and Maier (2003, pp.15–16), all information that has a "special status" may be put on a separate layer. However, this may result in abundance of layers that all have their specific interpretation, which would fail to account for any similarities between phenomena interpreted on different layers. In particular, the similar felicity conditions for anaphora and presuppositions described by van der Sandt (1992) and the strong correspondence between asserted content and conventional implicatures (see, e.g., Amaral et al., 2007) cannot be captured in a multi-dimensional (multi-layered) framework.

Secondly, not all material seems to strictly belong to a specific layer. For example, Maier (2009) adapts Layered DRT to account for the special behaviour of proper names and indexicals, which are taken to constitute a special layer for 'reference-fixing' content (Maier calls this the *'kripke-kaplan'* or kk-layer, separating its content from the *'fregian' fr*-layer). However, some expressions, such as proper names and third person pronouns must be allowed to 'hop' between layers in order to account for their different usages (e.g., third person pronouns are regularly used in both deictic and anaphoric constructions). This solution is criticized by Hunter (2012), who argues that a relaxation of the separation between layers seems to defeat their purpose, since apparently they do not represent strictly distinct parts of meaning. Hunter provides an alternative analysis in which she shows that no extra layer is needed for indexical content; the behaviour of reference-fixing expressions can be accounted for by adding an extra-linguistic context level to standard DRT, the content of which is determined by the actual state of the world. This context allows indexicals to pick out a unique object in the actual world, without the need for a separate layer of meaning.

The goal of the current paper is to apply a similar kind of dimension reduction for projection phenomena, and to show that their behaviour can be accounted for within a unidimensional framework. To this purpose, we develop Projective DRT, which extends standard DRT with a set of pointers to indicate the accommodation site of linguistic material. The framework can be seen as a refinement of Layered DRT, which integrates a subset of its layers into one and thereby accounts for the distinction, as well as the similarities between the different phenomena.
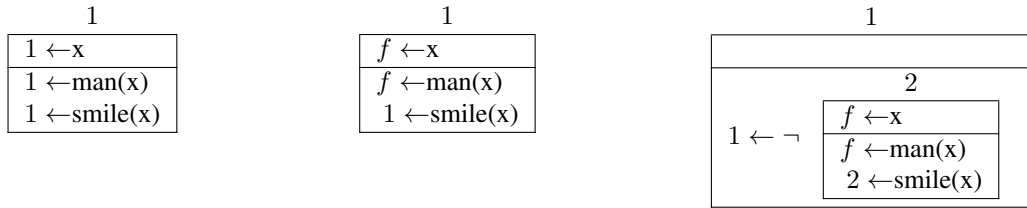
3 Projective Discourse Representation Theory

Projective DRT (PDRT) is an extension of standard DRT in which each referent and condition is associated with a pointer to indicate projection behaviour. The basic idea of PDRT is that all projected content is represented locally, i.e., at the introduction site, and that projection is signalled by means of pointers that indicate where the content is to be interpreted. This means that projection is not realised by physically moving semantic material in the resolution stage, but by setting a variable equation on pointers and PDRS labels. This representation stays closer to the linguistic surface structure, and reduces computational complexity born out of a two-stage resolution mechanism. Moreover, presupposed and asserted (i.e. non-projected) content are clearly discernible in the representation at each step of composition, while remaining subject to the same interpretation mechanism.

3.1 Projection as variable binding

In PDRT, asserted and projected material is treated in the same way, by associating the content with a pointer to its context of interpretation. The differences between asserted and projected material arise from the different contexts they point to. Asserted material gets as pointer the label of the PDRS in which it is introduced, and is thus interpreted locally. In the case of projected material, the pointer may refer to the label of an accessible PDRS (in van der Sandt's terminology: a PDRS on the projection path), or it may be a free variable. As a result, projected content is interpreted in the appointed PDRS or in the global PDRS in case the pointer is a free variable. An example is shown in (4), where we use integers to denote labels and bound pointers, and f for free pointers.

(4) a. A man smiles.   b. The man smiles.   c. It is not the case
                                                    that the man smiles.

1
1 ←x
1 ←man(x)
1 ←smile(x)

1
f ←x
f ←man(x)
1 ←smile(x)

1
1 ← ¬ [f ←x / f ←man(x) / 2 ←smile(x)] (label 2)

Each PDRS introduces a label, represented on top of the PDRS, and all referents and conditions associate with a label via a pointer, represented with an inverted arrow. If no material is projected, as in (4a), all material points to the PDRS in which it is introduced (the PDRS labeled '1'). In (4b) and (4c), on the other hand, the definite description 'the man' triggers a presupposition about the existence of its referent. In PDRT this is indicated by using a free variable as pointer for the presupposed material (here, 'f'). Free pointers are interpreted as pointing to the outermost PDRS (representing the discourse context), which both in (4b) and (4c) is the PDRS labeled '1'. As a result, the interpretations of (4a) and (4b) are equivalent, as desired, but on the representational level they are clearly distinguishable in order to account for their different compositional properties.

3.2 Preliminaries

The vocabulary of PDRT extends the standard DRT language with labels for DRSs and pointers for referents and conditions. A structure in PDRT (a PDRS) consists of a label ϕ, a set of projected referents D and a set of projected conditions C, resulting in a triple: $\langle \phi, D, C \rangle$. The projected referents and conditions are defined as follows:

Definition 1 (Projected referents)**.**
If p is a pointer and d is a discourse referent, then $\langle p, d \rangle$ is a projected discourse referent.

Definition 2 (Projected conditions)**.**

- If p is a pointer and P is an n-place predicate and $u_1, ..., u_n$ are discourse referents, then $\langle p, P(u_1, ..., u_n) \rangle$ is a projected condition.

- If p is a pointer and ϕ and ψ are PDRSs, then $\langle p, \neg\phi \rangle$, $\langle p, \phi \vee \psi \rangle$, $\langle p, \phi \rightarrow \psi \rangle$ are projected conditions.

Furthermore, accessibility between PDRSs and free variables are defined just as in standard DRT (Kamp and Reyle, 1993). Below, when possible, we will simply refer to the referents and conditions of PDRSs, instead of projected referents and projected conditions.

In the current implementation, the semantics of a PDRS is provided via a translation to standard DRT (see Section 4). This is computationally advantageous because of the model-theoretic properties of standard DRT, which are interpretable via first order logic (Muskens, 1996). This means that although in PDRT the movement of projected material is precluded at the representational level, in the interpretation it will be moved in order to obtain equivalence to DRT. This way, the theory inherits some attractive properties from the DRT account to presupposition, such as its inference mechanisms and predictions with respect to, for example, the proviso problem (cf. Geurts, 1999). However, the approach can easily be adapted to incorporate other interpretative models, for example a three-valued logic to account for presupposition failure in terms of undefinedness (see, e.g., Krahmer, 1998).

3.3 PDRS composition

Most presuppositional theories are lexically driven, i.e., based on the assumption that specific lexical items give rise to presuppositions (so-called 'presupposition triggers'). Therefore, projected material

will be manifested in the lexical semantics of projection triggers. Various authors have proposed a compositional treatment of DRT using basic tools from Montague Grammar and lambda calculus (Muskens, 1996; Bos, 2003; de Groote, 2006). Compositionality in PDRT is realised by providing every lexical item with an (unresolved) semantics in the form of a typed lambda term. In order to combine these unresolved semantics, a merge operation can be applied that combines two PDRSs into one by means of *merge-reduction* (see, e.g., Bos, 2003). In the current framework, we use different types of merge for asserted, presupposed and conventionally implied material in order to account for their different compositional properties.

In PDRT, projected material is not interpreted on a different level than asserted material, it only contributes to the context in a different way. This is realised by implementing distinct types of merge for asserted and presupposed material: assertive merge $(+)$ and projective merge $(*)$. Assertive merge between two (unresolved) PDRSs can be defined in the usual way by the union of the referents and conditions. Additionally, however, the pointers that refer to the merged PDRSs (i.e., the bound pointers) must be unified with the label of the resulting PDRS, in order to secure that asserted material is interpreted locally. The definition of assertive merge operations is shown below. For the renaming of pointers we use the notation '$A[x/y]$', which is taken to represent the set resulting from replacing every instance of y in the set A by x.

Definition 3 (Assertive merge).
$$
\begin{array}{c} i \\ \boxed{\begin{array}{c} D_i \\ \hline C_i \end{array}} \end{array}
\; + \;
\begin{array}{c} j \\ \boxed{\begin{array}{c} D_j \\ \hline C_j \end{array}} \end{array}
\; := \;
\begin{array}{c} j \\ \boxed{\begin{array}{c} D_i[j/i] \cup D_j \\ \hline C_i[j/i] \cup C_j \end{array}} \end{array}
$$

In words, the definition for assertive merge defines the merge of two asserted PDRSs as the union of the domains and conditions of the PDRSs, with the local pointers of the PDRS in the first argument of the merge (labeled i) replaced by the label of the second argument of the merge (labeled j).

Projected material, on the other hand, is not affected by the local context, but keeps its own pointer, which either refers to its accommodation site or is a free variable. Therefore, projective merge only involves adding the projected referents and conditions to the resulting DRS, without affecting their interpretation. This results in the following definition:

Definition 4 (Projective merge).
$$
\begin{array}{c} i \\ \boxed{\begin{array}{c} D_i \\ \hline C_i \end{array}} \end{array}
\; * \;
\begin{array}{c} j \\ \boxed{\begin{array}{c} D_j \\ \hline C_j \end{array}} \end{array}
\; := \;
\begin{array}{c} j \\ \boxed{\begin{array}{c} D_i \cup D_j \\ \hline C_i \cup C_j \end{array}} \end{array}
$$

Conventional implicatures, in turn, exhibit yet a different type of compositional behaviour (Potts, 2005). Like presuppositions, CIs project out of their local context. Unlike presuppositions, however, they cannot bind to an antecedent, nor accommodate locally (i.e., non-globally). In PDRT, this is realised by always projecting conventionally implied content to the outermost context (the "global" PDRS). This way, conventional implicatures receive an interpretation that is in some way between that of presuppositions and assertions: CIs accommodate at the highest possible context, while assertions accommodate locally and presuppositions remain free to indicate binding possibilities. In the definition for implicative merge, this means that all (bound) pointers of the conventionally implied content are replaced by a constant, say '0', which always refers to the outermost discourse context. This results in the following definition:

Definition 5 (Implicative merge).
$$
\begin{array}{c} i \\ \boxed{\begin{array}{c} D_i \\ \hline C_i \end{array}} \end{array}
\; \bullet \;
\begin{array}{c} j \\ \boxed{\begin{array}{c} D_j \\ \hline C_j \end{array}} \end{array}
\; := \;
\begin{array}{c} j \\ \boxed{\begin{array}{c} D_i[0/i] \cup D_j \\ \hline C_i[0/i] \cup C_j \end{array}} \end{array}
$$

3.4 Projection in PDRT

Next we will show how the different merge definitions are implemented in the lexical semantics of the linguistic material, resulting in a unified compositional framework for the representation of asserted content, presuppositions and conventional implicatures.

3.4.1 Asserted versus projected content

The distinction between asserted content and projected content is achieved by making use of different merge operations, reflecting the different ways in how the content is added to the discourse context. As an example, we look at the lexical semantics of definite descriptions and indefinites. In order to obtain the representations shown in (4), the indefinite should be added to the local context and the definite description should project using a free variable as pointer. This can be achieved by using different types of merge in the lexical semantics of "a" and "the". An indefinite description combines with the local context using an assertive merge, which means that the referent inherits the label from the merged PDRS and thus becomes asserted content. Definite descriptions, on the other hand, project out of their local context, which can be achieved using projective merge. The resulting lexical semantics for the determiners "a" and "the" are shown in (5).

(5)   a.   "a": $\lambda p.\lambda q.((\ \boxed{\begin{array}{c} i \\ \hline i \leftarrow x \\ \hline \end{array}}\ + p(x)) + q(x))$

      b.   "the": $\lambda p.\lambda q.((\ \boxed{\begin{array}{c} i \\ \hline i \leftarrow x \\ \hline \end{array}}\ + p(x)) * q(x))$

The lexical semantics of the indefinite article "a" introduces a discourse referent in a local PDRS. This PDRS is first combined with a predicate (e.g. a noun like "man") using assertive merge. The result of this merge operation is then combined with another predicate (e.g. a verb like "smiles"), again using *assertive* merge. This results in a representation where the indefinite description ("a man") is interpreted locally in the PDRS introduced by the rest of the context ("smiles"). For the definite article "the", on the other hand, the *projective* merge is used to combine the result of the first, assertive merge with the rest of the context. This means that the definite description keeps its own pointer, which will either be bound by an accessible PDRS, or become a free variable in the final representation, indicating accommodation.

     Other presupposition triggers, such as pronouns and proper names, receive a lexical semantics similar to definite descriptions, using projective merge. In case a presupposition gets bound, the standard DRT analysis can be used, introducing an equality relation between the referent and the antecedent (Kamp and Reyle, 1993). Alternatively, we can unify the referent with the antecedent, as in van der Sandt (1992).

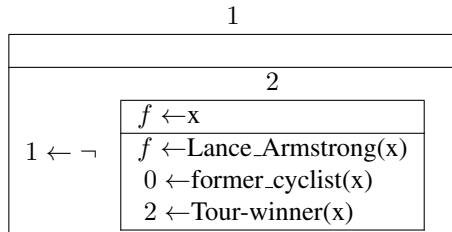### 3.4.2 Conventional Implicatures

Potts (2005) defined the class of conventional implicatures on the basis of a set of specific criteria, including non- cancellability, not at-issueness, scopelessness and speaker orientation. He roughly categorizes CIs into two groups: supplemental expressions (including appositives, non-restrictive relative clauses – NRRCs– and parenthetical adverbs) and expressives (including expressive attributive adjectives, epithets and honorifics). Potts (2005) presents a multi-dimensional framework in order capture the distinction between CIs and asserted content. However, there is strong evidence against such a multi-dimensional approach, as Amaral et al. (2007) argue that there is a strong interaction between CIs and asserted content and Simons et al. (2010) unify presuppositions and CIs as projection phenomena. Therefore, in Projective DRT conventional implicatures are treated in the same way as presuppositions and asserted content, with the peculiarity that CIs always accommodate to the global discourse level. This is realised by projecting CIs using the implicative merge defined in Section 3.3.

     Conventional implicatures are often triggered by constructions rather than lexical items, for example the subordinating constructions of appositives and NRRCs. In PDRT this is reflected by creating a special semantics for the subordinating comma, which projects its second argument. Because of the directionality of the merge operator, this means that the subordinating comma must reorder its arguments, such that the subordinated content is projected. The resulting semantics is shown in (6).

(6)   subordinating comma ",": $\lambda p.\lambda q.(q \bullet p)$

An example of the PDRT representation of an appositive is shown in (7). Note that the pointer of the appositive is '0', which is a constant referring to the label of the current global context, here '1'. Thus, both the presupposition introduced by the proper name and the CI introduced by the appositive accommodate to the global discourse context. The difference is that the pointer of the presupposition (indicated with 'f') remains available for binding, while the pointer of the appositive will always refer to the most global context.

(7) It is not true that Lance Armstrong, a former cyclist, is a Tour-winner.

$$
\begin{array}{|c|}
\hline
1 \\
\hline
\\
1 \leftarrow \neg\ \begin{array}{|l|}
\hline
2 \\
\hline
f \leftarrow x \\
\hline
f \leftarrow \text{Lance\_Armstrong}(x) \\
0 \leftarrow \text{former\_cyclist}(x) \\
2 \leftarrow \text{Tour-winner}(x) \\
\hline
\end{array} \\
\hline
\end{array}
$$

3.5 Comparison with related approaches

As described above, Layered DRT, as proposed in Geurts and Maier (2003), is a multi-dimensional framework that can account for different linguistic phenomena within a single representation. Projective DRT provides a unidimensional treatment for a subset of the phenomena covered in LDRT, including asserted content, presuppositions and conventional implicatures. The advantage of treating these different phenomena on a single 'layer' is that they are not treated as different kinds of meaning; they merely contribute their content to the context in a different way. A similar endeavour was taken by Hunter (2012), who argues for a unidimensional account of indexicals and asserted content. She proposes a DRT-style analysis in which an extra context is created for reference-fixing content, which is interpreted relative to the actual state of the world. This fits neatly within the idea of Projective DRT, where linguistic expressions are differentiated on the basis of the context the project ('point') to, and thus allows for a straightforward extension along these lines. We will leave an implementation of this and other extensions of PDRT for future work.

The account presented here is also related to the work of Schlenker (2011), who proposes a DRT account in the spirit of Heim (1983). In his representation, presupposed propositions are indexed with context variables that explicitly represent local contexts in the logical form. In this sense, his analysis is in line with approaches that use update semantics (e.g., Zeevat, 1992), because the context variable defines the context in which the presupposition is interpreted. The anaphoric aspect is therefore not in the presupposition itself, but in the context variables, which can anaphorically refer to accessible contexts. The consequence of this analysis is that accommodation does not imply adding the presuppositional content to a higher context, but rather interpreting it within this higher context. So, the interpretation of the presupposition itself, rather than that of the context in which it is accommodated is affected. In this respect Schlenker's approach crucially differs from Projective DRT, since in PDRT the traditional DRT strategy of adding presuppositions to their context of interpretation is applied. This allows for a straightforward analysis of cases of intermediate accommodation, which are difficult to capture in Schlenker's account. Moreover, PDRT allows for a more fine-grained analysis, since each referent and condition is associated with an interpretation site, while Schlenker only projects complete propositions.

4 Translation PDRT to DRT

The semantics of PDRSs can be described via a translation to standard DRT (Kamp and Reyle, 1993). As described above, PDRT is not strictly limited to this interpretation and may be extended to incorporate other interpretation models. We implemented PDRT as part of the wide-coverage semantic parser Boxer (Bos, 2008), including an automatic translation to standard DRT. Below we only provide a sketch of the algorithmic translation to DRT, as space limitations do not permit a description of the full translation.

4.1 Translation procedure

For the translation to DRT we make use of PDRT's separation of logical structure and linguistic content. Since each referent and condition is associated with a pointer to its accommodation site, it is possible to first separate this content from the embedded PDRS structure and accordingly project each condition to its appointed site. We assume that α-conversion is applied to the PDRS in order to make sure that all labels, pointers and referents use unique variables.

For convenience, we here describe the algorithm for translating PDRSs to DRSs in three steps. In the first step, all accommodation sites referents and conditions are gathered in separate sets. In the second step, the referents and conditions are added to their appointed accommodation site. In the third and final step, the PDRSs in the set of accommodation sites are combined to form a DRS.

Step 1. We start by creating three empty sets: one for accommodation sites (Π), one for discourse referents (Δ) and one for conditions (Γ). Starting from a PDRS $\Phi = \langle \varphi, D, C \rangle$, we define the pointer of Φ to be a constant: $p(\Phi) = g$, and we add this pointer, together with an empty PDRS with the label of Φ to Π: $\Pi \cup \langle p(\Phi), \langle \varphi, \{\}, \{\} \rangle \rangle$. All referents $d \in D$ are added to Δ. For the conditions $c \in C$, the base case is that c contains no embedded PDRSs, i.e., $c = \langle p, R(x_1, ..., x_n) \rangle$. In this case c is added to Γ. If c does contain an embedded PDRS, e.g., $c = \langle p, \neg \langle l, D_l, C_l \rangle \rangle$, then a fresh label is created, say l_0. This label is used as a sort of 'trace' to indicate where the embedded PDRS was introduced. We add $\langle l_0, \langle l, \{\}, \{\} \rangle \rangle$ to Π and $\langle p, \neg \langle l_0, \{\}, \{\} \rangle \rangle$ to Γ. This way, the context introduced by the embedded PDRS becomes available as an accommodation site, and the condition containing the embedded PDRS is added to the list of conditions. Accordingly, the referents (D_l) and conditions (C_l) of the embedded PDRS are recursively resolved in the same manner as described above, with respect to the current Δ, Γ and Π. This procedure can also be applied for other complex conditions, such as disjunctions, implications, modal expressions or propositional PDRSs (e.g., $c = \langle p, v : \langle l, D_l, C_l \rangle \rangle$).

Step 2. In this step, all referents in Δ and all conditions in Γ are projected to an appropriate PDRS in the list of accommodation sites, Π. For each referent $\langle l, u \rangle \in \Delta$, this means that if $\langle p, \langle l, D_l, C_l \rangle \rangle \in \Gamma$, then u is added to the domain: $D_l \cup u$ (so without the pointer). Otherwise, the label occurs free, so u is added to the domain of the outermost PDRS, which has g as pointer: $\langle g, \langle m, D_m \cup u, C_m \rangle \rangle$. The same strategy can be applied for conditions and the process continues until Δ and Γ are empty.

Step 3. The last step is to put the accommodation sites in Π (which now contain all the accommodated material) back together in order to form a translated DRS. We start with the DRS $\Phi = \langle D_1, C_1 \rangle$, such that: $\langle g, \langle l_1, D_1, C_1 \rangle \rangle \in \Pi$. This accommodation site is accordingly removed from Π. Then we check the conditions of Φ for embedded PDRSs. If such a complex condition is found, e.g. $c = \neg \langle l_c, D_c, C_c \rangle$, then the embedded PDRS is replaced by the DRS $\Psi = \langle D_m, C_m \rangle$, such that: $\langle l_c, \langle l_m, D_m, C_m \rangle \rangle \in \Pi$, which is accordingly removed from Π. Then, the set of conditions C_m of Ψ is again checked for embedded PDRSs. Once no embedded PDRSs remain, the remainder of the conditions of the dominating DRS (in this case, Φ) are checked. This recursive process goes on until Π is empty. At that point we will have a DRS with all the projected (and asserted) material at its accommodation site.

4.2 Example translation

We now provide an example of the translation procedure explained in the last subsection. The PDRS is shown in (8a), the desired DRS translation is shown in (8b) and its first-order logic equivalent in (8c).

(8) a.

b.

c. $\exists x \exists y (P(x) \wedge \neg Q(y))$

Step 1. We start with three empty sets: Δ, Γ and Π. First, we add an empty PDRS with the label of the outermost PDRS Φ and a fixed pointer, say 0, to the set of accommodation sites: $\Pi = \{\langle 0, \langle 1, \emptyset, \emptyset \rangle \rangle\}$. We add the referents and simple conditions of Φ to the correct sets: $\Delta = \{\langle 1,x \rangle\}$; $\Gamma = \{\langle f, P(x) \rangle\}$. Then, we create a fresh label, say 3, and add an empty PDRS with the label of the embedded PDRS and the fresh label as pointer to Π: $\Pi = \{\langle 0, \langle 1, \emptyset, \emptyset \rangle \rangle, \langle 3, \langle 2, \emptyset, \emptyset \rangle \rangle\}$. The condition with the operator and an empty PDRS with the fresh label is then added to Γ: $\Gamma = \{\langle f, P(x) \rangle, \langle 1, \neg \langle 3, \emptyset, \emptyset \rangle \rangle\}$. Finally, we add the content of the embedded PDRS to the corresponding sets: $\Delta = \{\langle 1,x \rangle, \langle 1,y \rangle\}$; $\Gamma = \{\langle f, P(x) \rangle, \langle 1, \neg \langle 3, \emptyset, \emptyset \rangle \rangle, \langle 2, Q(y) \rangle\}$.

Step 2. Now, we simply project each of the elements of Δ and Γ to the correctly labeled PDRS in Π, i.e., to the PDRS that has the pointer of the referent/condition as label, or to the PDRS with the pointer 0 in case of a free variable: $\Pi = \{\langle 0, \langle 1, \{x, y\}, \{P(x), \neg \langle 3, \emptyset, \emptyset \rangle\} \rangle \rangle, \langle 3, \langle 2, \emptyset, \{Q(y)\} \rangle \rangle\}$.

Step 3. Finally, we create a DRS Ψ from the accommodation site in Π that has 0 as pointer: $\Psi = \langle \{x, y\}, \{P(x), \neg \langle 3, \emptyset, \emptyset \rangle\} \rangle$. We check for embedded PDRSs in the conditions of Ψ and replace them with the DRS from the corresponding element in Π (matching the pointer to the label). The result is the following DRS: $\langle \{x, y\}, \{P(x), \neg \langle \emptyset, \{Q(y)\} \rangle\} \rangle$, which is exactly the desired DRS shown in (8b).

5   Conclusions and future work

In this paper we presented Projective DRT, and extension of DRT in which all linguistic material is associated with a pointer to indicate its accommodation site. This way, semantic material does not need to be moved or copied at the representational level, as projection is secured by using free variables as pointers, or by binding the pointers of projected material to labels introduced by higher level PDRSs. This is in line with van der Sandt's (1992) idea to treat presuppositions as anaphora, since in DRT anaphora resolution is also based on variable binding. The theory results in a simple and parsimonious representation of different linguistic phenomena, with a unified treatment of asserted content, presuppositions and conventional implicatures. Moreover, it allows for compositional construction of discourse structures with projected content while precluding a two-step resolution algorithm. The resulting representation structures have a straightforward interpretation via translation to standard DRT.

Projective DRT can be extended to account for other phenomena, as well as other interpretation models. For example, we above mentioned a possible extension with a special context for indexical content, as described by Hunter (2012). Other directions for future work include the incorporation of phenomena such as factive constructions and VP-ellipsis with presupposed content in PDRT. A proper treatment of such phenomena may ask for an extension of the PDRT syntax (for example, allowing multiple pointers for one condition) or a more elaborate semantics that is not necessarily interpretable via a translation to standard DRT.

All in all, PDRT provides a transparent and flexible compositional framework for investigating projection phenomena. The robustness of the framework has already been put to test through an implementation into Bos's (2008) wide-coverage semantic parser: Boxer. Future work will aim at evaluating and refining the PDRSs produced by Boxer via an integration into the Groningen Meaning Bank, a large-scale corpus of semantically annotated texts (Basile et al., 2012). PDRT allows for a coherent and easy-to-read representation of projection phenomena, since all content appears locally and the representation is therefore closer to the linguistic surface structure. This is important for a proper evaluation of semantic representations, as well as for studying the behaviour of linguistic phenomena. Implementation of PDRT into a large resource of semantically annotated texts will make an important contribution to corpus-based investigations into the behaviour of projection phenomena in discourse.

References

Amaral, P., C. Roberts, and E. Smith (2007). Review of "The Logic of Conventional Implicatures" by Chris Potts. *Linguistics and Philosophy 30*(6), 707–749.

Basile, V., J. Bos, K. Evang, and N. J. Venhuizen (2012). Developing a large semantically annotated corpus. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey, pp. 3196–3200. European Language Resources Association (ELRA).

Beaver, D. I. and B. Geurts (2011). Presupposition. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2011 ed.). Metaphysics Research Lab, CSLI, Stanford University.

Bos, J. (2003). Implementing the binding and accommodation theory for anaphora resolution and presupposition projection. *Computational Linguistics 29*(2), 179–210.

Bos, J. (2008). Wide-coverage semantic analysis with boxer. In *Semantics in Text Processing. STEP 2008 Conference Proceedings*, Volume 1, pp. 277–286.

de Groote, P. (2006). Towards a montagovian account of dynamics. In *Proceedings of SALT*, Volume 16, pp. 1–16.

Geurts, B. (1999). *Presuppositions and pronouns*. Elsevier.

Geurts, B. and E. Maier (2003). Layered DRT. Ms. To appear as: Layered Discourse Representation Theory. In Alessandro Capone (Ed.), *Perspectives on pragmatics and philosophy*. Springer, Berlin.

Heim, I. (1983). On the projection problem for presuppositions. In *Proceedings of the West Coast Conference on Formal Linguistics*, Volume 2, pp. 144–226.

Hunter, J. (2012). Presuppositional indexicals. *Journal of Semantics 0*, 1–41.

Kamp, H. and U. Reyle (1993). *From discourse to logic: Introduction to modeltheoretic semantics of natural language, formal logic and Discourse Representation Theory*. Kluwer Academic Dordrecht, The Netherlands.

Kracht, M. (1994). Logic and control: How they determine the behaviour of presuppositions. In J. van Eijck and A. Visser (Eds.), *Logic and Information Flow*, pp. 89–111. Cambridge, MA: MIT press.

Krahmer, E. (1998). *Presupposition and anaphora*. CSLI Publications.

Maier, E. (2009). Proper names and indexicals trigger rigid presuppositions. *Journal of Semantics 26*(3), 253–315.

Muskens, R. (1996). Combining Montague Semantics and Discourse Representation. *Linguistics and Philosophy 19*(2), 143–186.

Potts, C. (2005). *The Logic of Conventional Implicatures*. Oxford University Press, USA.

Schlenker, P. (2011). DRT with local contexts. *Natural language semantics 19*(4), 373–392.

Simons, M., J. Tonhauser, D. Beaver, and C. Roberts (2010). What projects and why. In *Proceedings of SALT*, Volume 20, pp. 309–327.

van der Sandt, R. (1992). Presupposition projection as anaphora resolution. *Journal of Semantics 9*, 333–377.

Zeevat, H. (1992). Presupposition and accommodation in update semantics. *Journal of Semantics 9*(4), 379–412.