Denial and correction in Layered DRT

Emar Maier
Department of Philosophy
University of Nijmegen
e.maier@phil.kun.nl

Rob van der Sandt
Department of Philosophy
University of Nijmegen
rob@phil.kun.nl

Abstract

The central characteristic of denials is that they perform a non-monotonic correction operation on discourse structure. A second characteristic is that they may be used to object to various kinds of information including presuppositions and implicatures. In this paper we first use standard DRT to capture these features, implement an earlier proposal of van der Sandt (1991) in DRT and point out a shortcoming of that approach. We then adopt *Layered DRT*. LDRT is an extension of standard DRT designed to represent and interpret different types of information conveyed in a conversation by distributing them over separate layers of the same LDRS. We will then show how LDRT allows us to solve the problems of the classic monostratal system. The resulting system makes use of a *directed reverse anaphora* mechanism to locate, remove and negate the material objected to.

1 Introduction

In the literature on denial and correction we can roughly distinguish two views. The dominant view has it that these phenomena are best captured within a general view of the semantics and pragmatics of negation. In terms of speech act theory this means that negatory force is located in the negative morpheme. A particular influential statement of this view is found in Horn (1985; 1989), who distinguishes between an ordinary 'descriptive' negation operator and a so-called metalinguistic negation. The former is the standard logical connective contributing to truth-conditional content, the latter is a non-truth-functional device that is meant to account for rejection of non-propositional material. This comprises objections to implicatural and presuppositional information, and to infelicities arising from style and register. Horn characterizes this operator as a metalinguistic device which can be used to signal an objection to an utterance on whatever grounds. This view has a number of drawbacks. It implies that some instances of of natural language negation are not interpretable by semantic means. It also implies an ambiguity in natural language negation which is not realized in any known language.

We defend a second view according to which the semantic notion of negation and the speech act notion of denial are fully independent. Denial patterns with assertion in that its nature and function should be accounted for in terms of the discourse effects it gives rise to. Just as the primary function of assertion is to convey new information, the primary function of a denial is to object to information which has been entered before and to remove it from the discourse record. The central characteristic of denial then is that it performs a non-monotonic correction operation on contextual information. This approach has several advantages. It yields a uniform account of denial, which comprises the standard 'propositional' and the marked cases, it does not force us to postulate an ambiguity in the natural language negation, and by incorporating ideas from Levinson (2000), van Leusen (1994; to appear) and van der Sandt (1991; 2003) it can be implemented in a natural way in a dynamic theory of discourse interpretation.

2 Denials in DRT

2.1 Dialogue

Van der Sandt (1991; 2003) gives a mechanism to locate and remove the information conveyed by a previous utterance in a discourse in order to account for the non-incremental behavior of denials in discourse. An implementation of this account in DRT requires some minor extensions to standard DRT; we need to keep track of who said what and when in a dialogue in order to locate and remove a previous speaker's contribution. We let a discourse be a sequence of sentences, $\sigma_1, \ldots, \sigma_n$.[1] The first addition to the syntax of the DRS language is harmless: we index every DRS condition and reference marker with a natural number specifying the sentence σ_i it originated from.[2]

Our goal is now to incrementally build DRSs, $\varphi_1, \ldots, \varphi_n$, representing the various stages of the ongoing discourse. We assume a given background representation φ_0 as starting point for the incrementation process.

2.2 Star conditions and reverse anaphora

We will assume that whenever an utterance is analyzed as a denial of a previous utterance, this denial is not further parsed, but leaves a simple negated $\star$-condition ($\neg[\ \star\ :\ \star\ ]$) in the DRS. DRS-construction now proceeds as follows. We first construct a preliminary sentence DRS from the parse of the sentence and merge it with the background DRS. The standard presuppositional resolution mechanism will then bind or accommodate the unresolved presuppositions and other anaphoric expressions (van der Sandt, 1992; Geurts, 1999). In the case of an assertoric utterance the output of this process is a new DRS which monotonically accumulates the new information conveyed by the utterance.

However, if the utterance under consideration gives rise to a $\star$-condition, we resolve it by a mechanism we call *reverse anaphora*. This mechanism collects all material originating from the previous utterance and *moves* it to the position of the $\star$. Reverse anaphora thus has a non-monotonic effect on the discourse structure established: the contribution of the previous utterance is removed from the main DRS and entered under the scope of the negation introduced by the denial.[3]

2.3 Examples I

The following example of a presupposition denial illustrates the procedure in some detail:

(1) σ_1 The King of France walks in the park.
σ_2 No, he doesn't,
σ_3 France doesn't have a king.

Let the background φ_0 be empty. The presuppositional expression of σ_1 (represented here by underlining) thus has to be accommodated. We denote the algorithms responsible for building preliminary sentence DRSs, resolution, and reverse anaphora, as **Prel**, **Res**, and **RA** respectively. $\oplus$ is the merge operator.

(2) a. $\mathbf{Prel}(\sigma_1)$ =
$[\ \underline{\mathrm{x}_1 : \text{King of France}_1(\mathrm{x})},\ \text{walk in park}_1(\mathrm{x})\ ]$

b. $\varphi_1 = \mathbf{Res}(\varphi_0 \oplus \mathbf{Prel}(\sigma_1))$ =
$[\ \mathrm{x}_1 : \text{King of France}_1(\mathrm{x}),\ \text{walk in park}_1(\mathrm{x})\ ]$

c. $\varphi_1 \oplus \mathbf{Prel}(\sigma_2)$ =
$\big[\ \mathrm{x}_1 : \text{King of France}_1(\mathrm{x}),\ \text{walk in park}_1(\mathrm{x}),\ \neg_2[\ \star : \star\ ]\ \big]$

d. $\varphi_2 = \mathbf{RA}\big(\mathbf{Res}(\varphi_1 \oplus \mathbf{Prel}(\sigma_2))\big)$ =
$\big[\ \cancel{\mathrm{x}_1} : \cancel{\text{King of France}_1(\mathrm{x})},\ \cancel{\text{walk in park}_1(\mathrm{x})},\ \neg_2[\ \mathrm{x}_1 : \text{King of France}_1(\mathrm{x}),\ \text{walk in park}_1(\mathrm{x})\ ]\ \big]$

[1] We will sometimes use the terms 'utterance' and 'sentence' rather loosely. The idea behind our terminology is that an utterance can be conceived of as a sentence combined with a representation of the *context*, where our notion of context comprises both a DRS representing the common ground and a Kaplanian context specifying the actual world, the speaker and the time of the sentence under discussion.

[2] Who said what and when can now be encoded in DRS conditions like *speaker(x,37)* and *time(t,37)*, meaning that x uttered σ_{37} at time t.

[3] Cf. the account of Levinson (2000) who calls this process *quasi-anaphora*

e. $\varphi_3 = \textbf{Res}(\varphi_2 \oplus \textbf{Prel}(\sigma_3)) =$
$\big[\ :\ \neg_2[\ \mathrm{x}_1 : \text{King of France}_1(\mathrm{x}),$
$\text{walk in park}_1(\mathrm{x})\ ],$
$\neg_3[\ \mathrm{x}_3 : \text{King of France}_3(\mathrm{x})\ ]\ \big]$

The upshot is that the final interpretation, φ_3, is equivalent to the predicate logical formula $\neg\exists x\big[\mathrm{KF}(x) \wedge \mathrm{walk}(x)\big] \wedge \neg\exists x\big[\mathrm{KF}(x)\big]$ which is easily seen to be equivalent to the intuitively correct $\neg\exists x[\mathrm{KF}(x)]$.

2.4 Problems

Strawson's famous (3) illustrates that we cannot always just remove whatever is conveyed by the offensive utterance:

(3) a. A man jumped off the bridge.
b. He didn't jump, he was pushed.
Strawson (1952)

when processing (4b) we need to retain the discourse referent introduced for *the man* in (4a) in order to bind the anaphoric pronoun.

Geurts (1998) argued against the strategy of removing the *full* contribution of the utterance objected to, pointing out that the problem is a general one. Consider an example where several presuppositions are involved:

(4) a. The King of France knows I quit smoking.
b. No he doesn't, France doesn't have a king.

The presupposition objected to is that France has a king and should be removed from the discourse record. The presupposition that I quit smoking seems to pass unharmed. We thus need to allow for the possibility that the final representation retains the latter information.

Another argument against removing the full contribution of a sentence is the fact that speakers sometimes overtly acknowledge parts of the last utterance's contribution while rejecting other parts. A case in point is (5) where the second speaker confirms that the person referred to is nice, but denies the implicature evoked by the use of the indefinite term 'a lady'.[4]

(5) a. Now, THAT's a nice lady
b. Yes, she is, but she's not a LAdy, she's my WIfe

In order to account for examples like (5) we need to represent and *interpret* different types of information conveyed by the very same utterance separately, thus enabling us to remove specific types information while retaining others.

3 Using layers

We will use layered representations in order to *direct* the reverse anaphora module at the *offensive* layers only. This requires a further extension of DRT which enables us to both encode and interpret the various kinds of information. We adopt Layered DRT (or LDRT for short).[5]

3.1 LDRT's syntax

The syntax of LDRT is like that of standard DRT, except that every discourse referent and DRS condition is paired with a layer label, specifying which kind of information it encodes. For the purpose of this paper we index every label with a number corresponding to the number of the sentence from which it originated (and 0 for background information). We will here restrict ourselves to the set of labels Λ_0 given in (6):

(6) $\Lambda_0 = \{0,\ fr_1, \ldots, fr_n,\ acc_1, \ldots, acc_n,\ imp_1, \ldots, imp_n,\ p_1, \ldots, p_n\}$
- $fr_i \approx$ the Frege layer, what is "said" in the contribution of σ_i
- $acc_i \approx$ accommodated material coming from p_i-layer
- $imp_i \approx$ implicature invoked by σ_i
- $p_i \approx$ presupposition triggered by σ_i

The primitive symbols of our LDRT language fragment are:

- a set $\mathcal{X}$ of reference markers

[4] Adapted from the classic:
A: That was a nice lady I saw you with last night.
B: That wasn't a LAdy, that was my WIfe
Note that by his utterance B certainly does not imply that his wife is not a lady as a purely truth-functional analysis would predict. Grice (1989: 37) remarks: 'Anyone who uses a sentence of the form *X is meeting a woman this evening* would normally implicate that the person to be met was someone other than X's wife, mother, sister, or perhaps even close platonic friend.'

[5] See Geurts and Maier (ms).

- (for some n) a set $\mathcal{P}red^{(n)}$ of n-place predicates
- the set Λ_0 of layer labels

The syntactic rules are as follows:

(7) a. if $x \in \mathcal{X}$, $L \subseteq \Lambda_0$, then $x_L = \langle x, L\rangle \in \mathcal{X} \times \wp(\Lambda)$ is a labeled reference marker

b. if $P \in \mathcal{P}red^n$, $L \subseteq \Lambda_0$, then P_L is a labeled predicate

c. if $x, y \in \mathcal{X}$, $L \subseteq \Lambda_0$, then $x =_L y$ is a labeled condition

d. if $x_1, \ldots, x_n \in \mathcal{X}$, P_L a labeled n-place predicate, then $P_L(x_1, \ldots, x_n)$ is a labeled condition

e. if φ and ψ are labeled conditions, $L \subseteq \Lambda_0$, then $\neg_L\varphi$, $\varphi \vee_L \psi$, and $\varphi \Rightarrow_L \psi$ are labeled conditions

f. if U is a set of labeled reference markers and Con a set of labeled conditions, then $\langle U, Con\rangle$ is an LDRS (notation: $\varphi = \langle U(\varphi), Con(\varphi)\rangle$)

The following is an example of a sentence and its preliminary LDRS representation:[6]

(8) a. σ_1 = It is possible the Pope is right.

b. $\big[x_{p_1}$: pope$_{p_1}(x)$, $\Diamond_{fr_1}$[: right$_{fr_1}(x)$], $\neg\Box_{imp_1}$[: right$_{imp_1}(x)$]$\big]$

This representation encodes the fact that the definite description 'the Pope' has a presuppositional status. Given an individual that satisfies this predicate the classical truthconditional (Fregean) contribution is that *he* is possibly right, and with a typical utterance of σ_1 the speaker furthermore conveys the implicature that *he* is not necessarily right. The crucial feature of the LDRT representation is that all conditions inhabit their own layer but nevertheless employ one and the same reference marker; the presuppositional, assertoric and implicature expressions are all linked to the same discourse referent, thus capturing the fact that they all attribute some property to the same individual.

[6]We will use $\neg\Box_{imp_1}$[: right$_{imp_1}$(x)] as a notational shorthand for $\neg_{imp_1}\big[$: $\Box_{imp_1}$[: right$_{imp_1}$(x)] $\big]$

3.2 LDRT's semantics

The idea of LDRT's semantics is that the truth definition is relativized to a set of labels, which means that we ignore all material labeled otherwise. The truth definition is thus only a minor extension of the standard definition of truth for DRSs in terms of verifying or truthful embeddings, requiring only two small adaptations: First, since ignoring certain layers often causes the remaining layers to form only an incomplete representation, missing perhaps some essential reference markers, the truth definition should be partial, i.e. for some choices of layer sets a truthvalue is undefined. In the following we will use $\uparrow$ and $\downarrow$ to abbreviate undefinedness and definedness respectively.

Secondly, since we are going to talk about propositions, we need at least an intensional version, i.e. we need possible worlds. An *intensional model* for our LDRT language fragment is a triple $\langle D, W, R\rangle$, where D is a set of individuals, W a set of 'possible worlds', which we take to be interpretation functions mapping basic predicates onto their extensions, and an an accessibility relation R ($\subseteq W^2$). With respect to such models the clauses below (9) give a partial definition of truth:

(9) For $L \subseteq \Lambda_0$, an LDRS φ ,and embedding functions f, g, define:

a. $U_L(\varphi) = \{x \in \mathcal{X} \mid \exists K[K \cap L \neq \emptyset \wedge x_K \in U(\varphi)]\}$

b. $Con_L(\varphi) = \{\psi \mid \psi \text{ is an LDRS condition } \wedge \exists K[K \cap L \neq \emptyset \wedge \psi_K \in Con(\varphi)]\}$

c. $f[X]g = f \subseteq g \wedge Dom(g) = Dom(f) \cup X$

Let $\mathcal{M} = \langle D, W\rangle$ be an intensional model. Let f be a partial embedding of the set of reference markers into D, $L \subseteq \Lambda_0$, $w \in W$, and φ an LDRS.

- $\|\varphi\|^f_{L,w} =$

$$\begin{cases} \{g \mid f[U_L(\varphi)]g\wedge \\ \quad \wedge\forall\psi \in Con(\varphi)\big[\\ \quad \|\psi\|^g_{L,w} = 1\big]\} & \text{if } \exists g\big[f[U_L(\varphi)]g\wedge \\ & \quad \wedge\forall\psi \in Con(\varphi)\big[\\ & \quad \|\psi\|^g_{L,w}\downarrow\big]\big] \\ \uparrow & \text{otherwise} \end{cases}$$

- $\|x =_K y\|^f_{L,w} =$
$$\begin{cases} 1 & \text{if } K \cap L = \emptyset \text{ or} \\ & \ x, y \in Dom(f) \wedge f(x) = f(y) \\ 0 & \text{if } K \cap L \neq \emptyset \text{ and} \\ & \ x, y \in Dom(f) \wedge f(x) \neq f(y) \\ \uparrow & \text{otherwise} \end{cases}$$
- $\|P_K(x_1, \ldots, x_n)\|^f_{L,w} =$
$$\begin{cases} 1 & \text{if } K \cap L = \emptyset \text{ or} \\ & \big(x_1, \ldots, x_n \in Dom(f) \wedge \\ & \wedge \langle f(x_1), \ldots, f(x_n)\rangle \in w(P)\big) \\ 0 & \text{if } K \cap L \neq \emptyset \text{ and} \\ & x_1, \ldots, x_n \in Dom(f) \wedge \\ & \wedge \langle f(x_1), \ldots, f(x_n)\rangle \notin w(P) \\ \uparrow & \text{otherwise} \end{cases}$$
- $\|\neg_K \psi\|^f_{L,w} =$
$$\begin{cases} 1 & \text{if } K \cap L = \emptyset \text{ or } \|\psi\|^f_{K\cap L,w} = \emptyset \\ 0 & \text{if } K \cap L \neq \emptyset \text{ and } \|\psi\|^f_{K\cap L,w} \downarrow \text{ and} \\ & \|\psi\|^f_{K\cap L,w} \neq \emptyset \\ \uparrow & \text{otherwise} \end{cases}$$
- $\|\psi \vee_K \chi\|^f_{L,w} =$
$$\begin{cases} 1 & \text{if } \|\psi\|^f_{K\cap L,w} \downarrow \wedge \|\chi\|^f_{K\cap L,w} \downarrow \text{ and} \\ & \|\psi\|^f_{K\cap L,w} \cup \|\chi\|^f_{K\cap L,w} \neq \emptyset \\ 0 & \text{if } \|\psi\|^f_{K\cap L,w} = \|\chi\|^f_{K\cap L,w} = \emptyset \\ \uparrow & \text{otherwise} \end{cases}$$
- $\|\psi \Rightarrow_K \chi\|^f_{L,w} =$
$$\begin{cases} 1 & \text{if } \|\psi\|^f_{K\cap L,w} \downarrow \wedge \forall g \in \|\psi\|^f_{K\cap L,w} : \\ & \quad \|\chi\|^g_{K\cap L,w} \downarrow \wedge \|\chi\|^g_{K\cap L,w} \neq \emptyset \\ 0 & \text{if } \exists g \in \|\psi\|^f_{K\cap L,w} :\ \|\chi\|^g_{K\cap L,w} = \emptyset \\ \uparrow & \text{otherwise} \end{cases}$$
- $\|\Box_K \psi\|^f_{L,w} =$
$$\begin{cases} 1 & \text{if } \forall w' Rw :\ \|\psi\|^f_{K\cap L,w'} \downarrow \wedge \neq \emptyset \\ 0 & \text{if } \exists w' Rw :\ \|\psi\|^f_{K\cap L,w'} = \emptyset \\ \uparrow & \text{otherwise} \end{cases}$$

Layered propositions or contents expressed by LDRSs are defined as follows:

(10) $\mathcal{C}^f_L(\varphi) =$
$$\begin{cases} \{w | \exists g \in \|\varphi\|^f_{L,w}\} & \text{if } \exists w[\|\varphi\|^f_{L,w} \downarrow] \\ \uparrow & \text{otherwise} \end{cases}$$
$\mathcal{C}_L(\varphi) = \mathcal{C}^\emptyset_L(\varphi)$

This defines a wide variety of content notions at once. For example $\|$(8b)$\|_{\{p_1,fr_1\}}$ is the proposition that the Pope is possibly right, and[7] $\|$(8b)$\|_{p_1,imp_1}$ the proposition that the Pope is not necessarily right.

In order to determine e.g. $\|$(8b)$\|_{imp_1}$, the proposition *implicated* by a sentence (a type of content we will often need), we might use a more sophisticated two-dimensional semantics, allowing us to determine such *open* propositions with respect to a backgroundable set of layers in addition to the layer under evaluation. See Geurts and Maier (ms) for such a two-dimensional account. Here we will follow a simpler strategy to determine a workable truthconditional substitute for *prima facie* undefined open propositions like $\|$(8b)$\|_{imp_1}$. The idea is to enlarge the set of evaluated layers until a proposition is expressed. Assuming that the layers we add, also express a (more general) proposition we get a substitute for the original undefined content by subtracting these two from each other.

(11) a. $\mathcal{C}^{[K]}_L(\varphi) =$
$$W - (\mathcal{C}_K(\varphi) - \mathcal{C}_{K\cup L}(\varphi)) = \{w \in W \mid w \in \mathcal{C}_K(\varphi) \rightarrow w \in \mathcal{C}_{K\cup L}(\varphi)\}$$
b. $\mathbf{Cl}(\varphi, L) =$ the smallest $K \subseteq \Lambda_0$ s.t. $\mathcal{C}_{K\cup L}(\varphi) \downarrow$ ('**Cl**osure set')
c. $\mathcal{C}^*_L(\varphi) = \mathcal{C}^{[\mathbf{Cl}(\varphi,L)]}_L(\varphi)$

It is easily verified that $\mathcal{C}^*_L(\varphi) = \mathcal{C}_L(\varphi)$ if both are defined. And if the plain layered proposition $\mathcal{C}_L(\varphi)$ is not defined, we can think of $\mathcal{C}^*_L(\varphi)$ as modeling the (weakest possible) truthconditional contribution of φ's L layers.

As an example, consider again the possibly right pope of (8): we saw that $\|$(8b)$\|_{imp_1}$ is undefined because the implicature layer lacks a reference marker. The *closure* of imp_1 is p_1, so:

(12) $\mathcal{C}^*_{imp_1}((8b)) =$
$W - (\mathcal{C}_{p_1}((8b)) - \mathcal{C}_{p_1 imp_1}((8b))) =$
the proposition that if there is a Pope, then there is a not necessarily right Pope

[7] leaving out subscripted and other obvious braces for convenience here and in the following

3.3 Directed reverse anaphora

The reverse anaphora algorithm depends on **Off**, a mechanism for singling out offensive material from an utterance that has been denied by a subsequent utterance. **Off** is defined as in (13), where L and K are sets of labels.

(13) $\mathbf{Off}(\varphi, K)$ = the smallest $L \subseteq \Lambda_0$ s.t. $\mathcal{C}^*_L(\varphi) \cap \mathcal{C}^*_K(\varphi) = \emptyset$

So $\mathbf{Off}(\varphi, K)$ gives us a set of labels that, in φ, clashes with a given set K, in other words, the smallest set of layers disjoint from K that cause φ as a whole to be contradictory.

To direct the reverse anaphora mechanism in an assertion-denial-correction sequence, σ_1-σ_2-σ_3 (as in (1)), we again represent the contribution of σ_2 by a condition of the form $\neg_{fr_2}[\ \star : \star\ ]$, as in 2.2. This however does not suffice for **Off** to determine which layer is offensive, since there is no contradiction yet. We thus first add the layered contribution of σ_3, causing a contradiction somewhere in our growing LDRS (say ψ). We then determine the cause of the contradiction on the basis of the fr-content of the correction: $\mathbf{Off}(\psi, \{fr_3\})$. The conflicting material in the offensive layer is now moved, by the modified reverse anaphora mechanism, to the place of the $\star$s thus ending up under the negation in layer fr_2.

Formally this comes down to the following redefinition of reverse anaphora, $\mathbf{RA}^*$

(14) $\mathbf{RA}^*_i(\varphi)$ = move $U_{\mathbf{Off}(\varphi, fr_i)}$ and $Con_{\mathbf{Off}(\varphi, fr_i)}$ from their original position in φ to where the $\star$'s are in φ (then erase the $\star$'s)[8,9]

[8] $\varphi \ominus \varphi_{\mathbf{Off}(\varphi,\{fr_{i+2}\})}\Big[[\ \star : \star\ ] \mapsto \varphi_{\mathbf{Off}(\varphi,\{fr_{i+2}\})}\Big]$

[9] Instead of moving the whole offensive layer determined by **Off** we may need an even more finegrained Reverse Anaphora algorithm moving only the smallest inconsistent subDRS within that layer. This is essentially what Van Leusen's (to appear) 'nonmonotonic updating' does in order to account for the discourse effects of denials and corrections in a compositional, monostratal variant of DRT. Something like this finer grained targetting of offensive stuff is needed to solve the problems posed by multiple accommodation (4b) and the Strawson example (3b). A definition of 'offensive material, conflicting with correction σ_i', depending on the notion of syntactic DRS inclusion, $\sqsubseteq$, would have to be added:

(i) $\mathbf{OFF}(\varphi, i)$ = the $\sqsubseteq$-smallest subDRS ψ, $\psi \sqsubseteq \langle U_{\mathbf{Off}(\varphi, fr_i)}, Con_{\mathbf{Off}(\varphi, fr_i)}\rangle$, s.t. $\mathcal{C}^*_{\mathbf{Off}(\varphi, fr_i)}(\psi) \cap \mathcal{C}^*_{fr_i}(\varphi) = \emptyset$

and $\mathbf{RA}^*$ would be modified accordingly:

(ii) $\mathbf{RA}^{*\prime}$ = move $\mathbf{OFF}(\varphi, i)$ to the position of the $\star$'s.

Interpretation now proceeds as follows:

- If $\langle \sigma_i, \sigma_{i+1}, \sigma_{i+2}\rangle$ is an assertion-denial-correction sequence, then $\varphi_{i+1} = \mathbf{Res}(\varphi_i \oplus \mathbf{Prel}(\sigma_{i+1})) = \varphi_i \oplus \big[\ : \neg_{\{fr_{i+1}\}}[\ \star : \star\ ]\ \big]$ and $\varphi_{i+2} = \mathbf{RA}^*_{i+2}\big(\mathbf{Res}(\varphi_{i+1} \oplus \mathbf{Prel}(\sigma_{i+2}))\big)$
- Otherwise, $\varphi_{i+1} = \mathbf{Res}(\varphi_i \oplus \mathbf{Prel}(\sigma_{i+1}))$ and similarly for φ_{i+2}

In the next section we will illustrate the workings of the definitions given above by applying them to some selected examples.

3.4 Examples II

Consider first an example of an implicature denial, cf. (8):

(15)
σ_1 It is possible the Pope is right.
σ_2 No, it is not POssible,
σ_3 it is NEcessary that he is right.

We assume that the denial σ_2 pertains only to part of the content of σ_1 and that the correction σ_3 indicates that only the implicature is at stake. The step-by-step interpretation process of this dialogue according to the Directed Reverse Anaphora algorithm is shown below:

- $\varphi_0 = [x_0 : \text{pope}_0(x)]$
- $\mathbf{Prel}(\sigma_1) =$
$\big[x_{p_1} : \text{pope}_{p_1}(x),\ \Diamond_{fr_1}[: \text{right}_{fr_1}(x)],\ \neg\Box_{imp_1}[: \text{right}_{imp_1}(x)]\big]$
- $\varphi_1 = \mathbf{Res}(\varphi_0 \oplus \mathbf{Prel}(\sigma_1)) =$
$\big[x_0 : \text{pope}_0(x),\ \Diamond_{fr_1}[: \text{right}_{fr_1}(x)],\ \neg\Box_{imp_1}[: \text{right}_{imp_1}(x)]\big]$
- $\mathbf{Prel}(\sigma_2) = \big[\ : \neg_{fr_2}[\star : \star]\big]$
- $\varphi_2 = \mathbf{Res}(\varphi_1 \oplus \mathbf{Prel}(\sigma_2)) =$
$\big[x_0 : \text{pope}_0(x),\ \Diamond_{fr_1}[: \text{right}_{fr_1}(x)],\ \neg\Box_{imp_1}[: \text{right}_{imp_1}(x)],\ \neg_{fr_2}[\star : \star]\big]$
- $\mathbf{Prel}(\sigma_3) =$
$[z_{p_3} : \text{masc}_{p_3}(z),\ \Box_{fr_3}[: \text{right}_{fr_3}(z)]]$

- $\psi = \mathbf{Res}(\varphi_2 \oplus \mathbf{Prel}(\sigma_3)) =$
  $\big[x_0 : \text{pope}_0(x),\ \Diamond_{fr_1}[: \text{right}_{fr_1}(x)],$
  $\neg\Box_{imp_1}[: \text{right}_{imp_1}(x)],\ \neg_{fr_2}[\star : \star],$
  $\Box_{fr_3}[: \text{right}_{fr_3}(x)]\big]$

- $\mathbf{Cl}(\psi, fr_3) = \{0\}$ so $\mathcal{C}^*_{fr_3}(\psi)$ is the proposition that if there is a pope, there is a necessarily right one (see (12)).

- $\mathbf{Off}(\psi, fr_3) = \{imp_1\}$, because $\mathcal{C}^*_{imp_1}(\psi)$ is the proposition that if there is a pope, there is a not necessarily right one, which contradicts $\mathcal{C}^*_{fr_3}(\psi)$, on the assumption that it is common ground between the participants that there exists some unique pope.

- $\varphi_3 = \mathbf{RA}^*(\psi) =$
  $\big[x_0 : \text{pope}_0(x),\ \Diamond_{fr_1}[: \text{right}_{fr_1}(x)],$
  $\cancel{\neg\Box_{imp_1}[: \text{right}_{imp_1}(x)]},$
  $\neg_{fr_2}[: \neg\Box_{imp_1}[: \text{right}_{imp_1}(x)]],$
  $\Box_{fr_3}[: \text{right}_{fr_3}(x)]\big]$

Note that van der Sandt (1991) and our implementation thereof was able to handle this example, too, giving an equivalent representation as the final output. However, in the course of the resolution process it temporarily threw out the whole contribution of the assertion objected to.

The example in (5), repeated below as (16), shows that such an omnilayered removal is problematic:

(16) σ_1 Now, THAT's a nice lady
σ_2 Yes, she is,
σ_3 but she's not a LAdy,
σ_4 she's my WIfe

The earlier monostratal proposal would in effect first have to somehow acknowledge that the person referred to is both nice and female with σ_2, but then by reverse anaphora remove that information from the discourse representation along with everything else conveyed by σ_1. The information that this person is female is then restored with the correction σ_4, but the suggestion that she is nice does not survive. To see how the layered version handles this case, note first that this example is a little more involved than the previous ones since it features an *affirmation* or acknowledgement, σ_2, of (part of the content of) σ_1, before the familiar denial (σ_3) and correction (σ_4).[10]

- $\varphi_0 = [x_0 : \text{woman}_0(x), \text{pointed\_at}_0(x)]$
  (we start with a representation of someone's pointing at a woman)

- $\mathbf{Prel}(\sigma_1) = [y_{p_1} : \text{pointed\_at}_{p_1}(y),$
  $\text{lady}_{fr_1}(y), \text{nice}_{fr_1}(y), \text{stranger}_{imp_1}(y)]$
  (assuming the use of 'a lady' invokes the implicature that the person referred to is not a close relative of the addressee)

- $\varphi_1 = [x_0 : \text{woman}_0(x), \text{pointed\_at}_0(x),$
  $\text{lady}_{fr_1}(x), \text{nice}_{fr_1}(x), \text{stranger}_{imp_1}(x)]$

- $\mathbf{Prel}(\sigma_2) = [z_{p_2} : \text{fem}_{p_2}(z), \text{nice}_{fr_2}(z)]$
  (affirmation is treated just like asserted content)

- $\varphi_2 = \mathbf{Res}(\varphi_1 \oplus \mathbf{Prel}(\sigma_2)) =$
  $[x_0 : \text{woman}_0(x), \text{pointed\_at}_0(x),$
  $\text{lady}_{fr_1}(x), \text{nice}_{fr_1 fr_2}(x), \text{stranger}_{imp_1}(x)]$

- $\mathbf{Prel}(\sigma_3) = \big[\ : \neg_{fr_3}[\star : \star]\big]$
  (the negative sentence σ_3 is marked as a denial)

- $\varphi_3 = \big[x_0 : \text{woman}_0(x), \text{pointed\_at}_0(x),$
  $\text{lady}_{fr_1}(x), \text{nice}_{fr_1 fr_2}(x), \text{stranger}_{imp_1}(x),$
  $\neg_{fr_3}[\star : \star]\big]$

- $\mathbf{Prel}(\sigma_4) = [w_{p_4} : \text{fem}_{p_4}(w), \text{wife}_{fr_4}(w)]$

- $\psi = \mathbf{Res}(\varphi_3 \oplus \mathbf{Prel}(\sigma_4)) =$
  $\big[x_0 : \text{woman}_0(x), \text{pointed\_at}_0(x),$
  $\text{lady}_{fr_1}(x), \text{nice}_{fr_1 fr_2}(x), \text{stranger}_{imp_1}(x),$
  $\neg_{fr_3}[\star : \star], \text{wife}_{fr_4}(x)\big]$

- $\mathbf{Off}(\psi, fr_4) = imp_1$, since $\mathcal{C}^*_{fr_4} \cap \mathcal{C}^*_{imp_1} = \emptyset$
  ('if there is a woman pointed at, a woman pointed at is a stranger' and 'if there is a woman pointed at, a woman pointed at is my wife' do not strictly speaking contradict each other, but they do once we add the by

[10] The formal handling of affirmations like this in the current framework necessitates among other things some minor revision of the assumed consecutive numberings in the formal definitions of section 3.3. We will not spell out the details.

all means reasonable assumption that the dialogue partners assume the object of the pointing to be uniquely determined)[11]

- $\varphi_4 = \mathbf{RA}^*(\psi) =$
$\big[x_0 : \text{woman}_0(x), \text{pointed\_at}_0(x), \text{lady}_{fr_1}(x), \text{nice}_{fr_1 fr_2}(x), \cancel{\text{stranger}_{imp_1}(x)}, \neg_{fr_3}[: \text{stranger}_{imp_1}(x)], \text{wife}_{fr_4}(x)\big]$

4 Conclusion

In this paper we presented a representational account of the discourse function of denial in an extension of DRT and discussed two variants of what we called reverse anaphora. The first version follows van der Sandt (1991) by removing the full contribution of a previous utterance from the discourse record. We then presented *directed reverse anaphora* which removes the main objection against the undirected version by allowing us to direct and limit the retraction procedure to the offensive part of a previous utterance. The mechanism is able to do so by relying on a DRT-extension that enables us to encode and interpret information that historically has been labeled pragmatic and non-truth-conditional in nature. The account has moreover the advantage of being more general than Horn (1989) and Geurts (1998) in retaining a unified account of the propositional, implicature and presuppositional denials, providing a general semantics and accounting for the function of denials in terms of their non-monotonic discourse effects.[12]

[11] The two-dimensional treatment of open propositions alluded to on page 5 and discussed in (Geurts and Maier, ms) will give an immediate contradiction here.

[12] One might think that the style and register cases, discussed in Horn (1989) would escape our treatment since the information objected too does not seem propositional in an intuitive sense.

(i) They didn't call the POlice—they called the poLIce.

We should remark that LDRT does allow us to account of denials like (i), by positing a layer for intonational and other surface features. However, the semantics of formal layers turns out to require heavier semantic apparatus than we have space here to develop, so we leave out examples like (i) in our treatment of denial, and refer the interested reader again to (Geurts and Maier, ms).

References

Bart Geurts and Emar Maier. ms. Layered DRT. (2003).

Bart Geurts. 1998. The mechanisms of denial. *Language*, 74:274–307.

Bart Geurts. 1999. *Presuppositions and Pronouns*. Elsevier, Amsterdam.

H. Paul Grice. 1989. *Studies in the Way of Words*. Harvard University Press, Cambridge, Massachusetts.

Laurence R. Horn. 1985. Metalinguistic negation and pragmatic ambiguity. *Language*, 61(1):121–174.

Laurence R. Horn. 1989. *A Natural History of Negation*. Chicago University Press, Chicago.

Hans Kamp and Uwe Reyle. 1993. *From Discourse to Logic: Introduction to Modeltheoretic Semantics in Natural Language, Formal Logic and Discourse Representation Theory, Vol. 1*. Kluwer Academic Publishers, Dordrecht.

Stephen C. Levinson. 2000. *Presumptive Meanings*. The MIT Press, Cambridge, Massachusetts.

Peter F. Strawson. 1952. *Introduction to Logical Theory*. Methuen, London.

Rob A. van der Sandt. 1991. Denial. In *Papers from CLS 27(2): the parasession on negation*, pages 331–344. CLS.

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

Rob A. van der Sandt. 2003. Denial and presupposition. In Hannes Rieser Peter Kühnlein and Henk Zeevat, editors, *Perspectives on Dialogue in the New Millennium*. John Benjamins, Amsterdam.

Noor van Leusen. 1994. The interpretation of corrections. In Peter Bosch and Rob van der Sandt, editors, *Focus and Natural Language Processing*, pages 523–532, Heidelberg. IBM Deutschland, Institute for Logic and Linguistics.

Noor van Leusen. to appear. *Nonmonotonic update in description grammar*. Ph.D. thesis, University of Amsterdam, Amsterdam.