Layered Discourse Representation Theory

Bart Geurts and Emar Maier

Abstract Layered Discourse Representation Theory (LDRT) is a general framework for representing linguistic content. Different types of content (e.g. asserted, presupposed, or implicated information) are separated by putting them on different layers, all of which have a model-theoretic interpretation, although not all layers are interpreted uniformly. It is shown how LDRT solves so-called 'binding problems', which tend to arise whenever different kinds of content are separated too strictly. The power of the framework is further illustrated by showing how various kinds of contextual information may be accommodated.

1 Introduction

The information conveyed by any utterance is a mixed bag. Utterances carry content about the world as it is according to the speaker, but also about speakers' attitudes, the way they speak, what has been said before, and so on. There are many kinds of information that are conveyed by way of language, and differences in kind correlate with differences in status. Presupposed information exhibits a distinctive projection behaviour; conversational implicatures are cancellable in a way that asserted information is not; in French or German, a pronoun's grammatical gender may help to determine a referent, but is otherwise truth-conditionally inert; and so on.

B. Geurts (✉)
Department of Philosophy, University of Nijmegen, Nijmegen, Netherlands
e-mail: brtgrts@gmail.com

E. Maier
Department of Philosophy, University of Groningen, Groningen, Netherlands
e-mail: emar.maier@gmail.com

A. Capone et al. (eds.), *Perspectives on Linguistic Pragmatics*, Perspectives in Pragmatics, Philosophy & Psychology 2, DOI: 10.1007/978-3-319-01014-4_12,

Interpreting utterances is as much a matter of integrating these various kinds of information as of keeping them apart. This much is uncontroversial. As far as we are aware, however, no attempts have been made thus far to devise a fully general framework within which processes of information integration can be modeled. There are partial theories, to be sure. For example, there are quite a few well-developed analyses of the interaction between presupposed and non-presupposed content. But to the best of our knowledge the problem of information integration as such has not been addressed before. So that is what this chapter is about: a general framework for representing and integrating all and sundry kinds of information that can be conveyed by linguistic means. This may seem like a grandiose project, and perhaps it is, but it is less ambitious than one might think. Our aim in this chapter is to develop a framework for *representing* different kinds of linguistic and para-linguistic information. How this information is processed is a different matter altogether, and not our main concern in the following.

2 Information Integration

In order to explain what we mean by information integration, we will discuss a few concrete cases. It will be seen that our examples are quite diverse, but this is to be expected in view of the broad aim of this chapter.

2.1 Presupposition

Our first example concerns the representation of presuppositions. In Discourse Representation Theory (DRT; Kamp 1981; Kamp and Reyle 1993), presuppositions are treated on a par with anaphoric expressions. Presuppositions prefer to link up to an antecedent, and if no suitable antecedent is available, they are interpreted by way of accommodation (van der Sandt 1992, Geurts 1999). Here is an example:

(1) Perhaps White met the Chinese Empress today.

The initial representation of this sentence is as follows:

(2) [y: White(y), ◇[x:Chinese-Empress(x), meet(y, x)]]

The underlining in this Discourse Representation Structure (DRS) reflects the fact that the definite NP 'the Chinese Empress' triggers the presupposition that there is a Chinese Empress. Assuming that this presupposition doesn't have a suitable antecedent, it may be construed by way of accommodation (provided the hearer is prepared to accept that there is a Chinese Empress), which means that the presupposition is added to the principal DRS, yielding the following representation:

(3) [x, y: Chinese-Empress(x), White(y), ◇[: meet(y, x)]]

This correctly captures what is expressed by (1), viz. that there is a Chinese Empress, and that White may have met her. Note that in the final representation of (1) the distinction between presupposed and asserted information is obliterated.

For many purposes this is fine, for we mainly need that distinction in order to account for the fact that presupposed material is processed in its own special way. But as it turns out, the distinction between presupposed and asserted information remains active after a sentence has been processed and the presupposition accommodated. To illustrate this, consider what would happen should another speaker object against (1) as follows:

(4) No, he had an encounter with the Japanese President.

Intuitively, this response only corrects what (1) asserts; the accommodated presupposition that there is a Chinese Empress remains unscathed. There are also ways to achieve the opposite, to deny the accommodated presupposition but leave the asserted content (Maier and der Sandt 2003; von Fintel 2004):

(5) Hey, wait a minute, China doesn't have an Empress!

It is not a good idea, apparently, to discard the division between presupposed and non-presupposed material once the mechanism of presupposition projection has performed its duty.

The moral of these observations is obvious and quite independent of the theory of presupposition we happen to prefer. It is simply that presuppositions will have to be separated from other types of information, because they have a special status: presuppositions are processed in their own way and once they have been accommodated they continue to be treated differently, as is shown by (5).[1]

2.2 *Implicatures*

What has been just said about presuppositions holds good for implicatures, too. By way of example, consider what is generally regarded a 'scalar implicature':

(6) The porridge is warm.

An utterance of this sentence presupposes that there is porridge, it asserts that the porridge is warm, and it implicates that the porridge is not hot; so the lexical meaning of 'warm' does not by itself rule out that the porridge is hot. According to this analysis, the information that the porridge is warm is of a different kind than

[1] Actually, there are two issues here. First, presuppositional material has to be separated from other types of content in order for the projection mechanism to perform its function. Secondly, once the presupposition has been processed, it must remain separated, as we have just argued. In the following, we confine our attention to the second issue. Whether or not presuppositions in preliminary DRSs must be interpreted, too, is a different matter, which we will not take a stance on here.

the information that the porridge is not hot, and one of the stock-in-trade arguments in favour of the distinction is that the implicature is cancellable in a way the assertion is not:

(7) a. The porridge is warm. As a matter of fact, it is hot.
b. ?The porridge is warm. As a matter of fact, it is cold.

(7a) shows that the implicature is cancellable, and the oddness of (7b) suggests rather strongly that asserted information is more robust.

The upshot of these observations is analogous to that of the presuppositional case. We need to separate implicated information from other information conveyed by an utterance, and it will not do to discriminate between presupposition, assertion, and implicature only for the duration of sentence processing (as in Gazdar 1979, for example); for the subsequent discourse may need these distinctions, too.

2.3 Non-Literal Meaning

The heading of this rubric is somewhat tentative. What we have in mind are such phenomena as metaphor, metonymy, irony, and so on: non-truth-conditional content that is clearly part of the speaker's message, but may be at odds with its literal meaning. Even if such phenomena are to be treated in terms of conversational implicature, we prefer to distinguish them from run-of-the-mill cases of implicature, which merely add to the literal meaning of an utterance. Especially stark cases in point are irony and sarcasm. Suppose a connoisseur of modern art volunteers (8), pointing at what is obviously a fumbled attempt at self-expression:

(8) That is a beautiful painting.

Under the circumstances, this statement is probably intended to convey the opposite from what it literally says.

Another example to bring out the need for information segregation is sentence (9), as said by a father to his 15-year-old son:

(9) Someone used my after shave this morning.

The use of an indefinite would normally implicate that the subject is thought not to be present in the context of discourse, but in this particular case the utterance may be understood as implying that the addressee is the culprit, and if it is, the implicature is cancelled.

It is, mildly put, something of a mystery how such 'double meanings' are computed, and we don't have anything new to offer in this regard. However, we do have a proposal as to how different levels of meanings can be represented in such a way that some bits of information are shared (e.g., the reference of the pronoun is shared between the literal and the non-literal meaning of the sentence), while others are segregated.

2.4 The Former and the Latter

Consider the following example:

(10) If a beggar meets a bishop, then the latter will bless the former.

On the face of it, anaphoric devices like 'the former' or 'the latter' do not seem particularly troublesome. On reflection, however, they add an interesting wrinkle to the problems posed by definite NPs. As the wrinkle will appear, in some way or other, no matter what our theoretical predilections concerning definites are, we will follow our own. According to the DRT treatment of definite NPs that we favour, definites are presuppositional expressions, which is to say that they prefer to link up to a contextually given antecedent. In this respect, everything is fine in the present example, since 'the latter' as well as 'the former' have suitable antecedents: the former refers back to 'a bishop', the latter to 'a beggar'. However, problems begin to emerge when we ask ourselves how exactly these expressions manage to link up to their antecedents. To see this, consider how the story about 'the prelate' in (11) would go:

(11) If a beggar meets a bishop, then the prelate will bless him.

In this sentence, 'the prelate' establishes an anaphoric link in much the same way as 'the latter' does in (10). In this respect there is little difference between the two expressions. However, the descriptive content of 'the prelate' is very different from that of 'the latter'—so different, in fact, that some people would say the information contained in 'the prelate' is part of the sentence's truth-conditional content, whereas nobody would want to claim that the descriptive content of 'the latter' enters into the truth conditions of (10).

So the wrinkle is this. Since they appear to be just a special case of presuppositional (or anaphoric) expressions, we would like to analyse 'the former' and 'the latter' as being on a par with any other definite NP, except of course that they constrain the process of interpretation by referring not to *what* has been said but *how* it was said. Qua presuppositional expression, 'the latter' presents its referent as given in the same way 'the prelate' does; it is just that we have two rather different modes of givenness, so to speak. The problem is, therefore, how we can distinguish two very different kinds of information—about the discourse and about the world—and have a uniform account of definites at the same time.

Note, incidentally, that in some cases ordinary pronouns may be used in the same way as English 'the former/latter', so a strict distinction between 'referring' and 'formal' definites becomes even less desirable:

(12) Am Ende besteht ein wesenhafter Unterschied zwischen [dem Erfassen des Ganzen des Seienden an sich]$_i$ und [dem Sichbefinden inmitten des Seienden im Ganzen]$_j$. Jenes$_i$ ist grundsätzlich unmöglich. Dieses$_j$ geschieht ständig in unserem Dasein. (Martin Heidegger, *Was ist Metaphysik?*)

We will not endeavour to render this passage in colloquial English. Suffice it to say that it illustrates how in German the distal and proximal demonstrative pronouns 'jenes' and 'dieses' are used precisely as 'the former' and 'the latter' would be used in English. The following quote from Somerset Maugham shows that English pronouns have the same meta-linguistic use, too, although this may be a more isolated example:

(13) For it was clear that the two were irreconcilable, the state and the individual conscious of himself. *That* uses the individual for its own ends, trampling upon him if he thwarts it, rewarding him with medals, pensions, honours, when he serves it faithfully; *this*, strong only in his independence, threads his way through the state, for convenience' sake, paying in money or service for certain benefits, but with no sense of obligation; and, indifferent to the rewards, asks only to be left alone. (W. Somerset Maugham, *Of Human Bondage*)

2.5 Grammatical Gender

Up to a point, grammatical-gender pronouns function not unlike 'the latter' and 'the former', as witness the following example from German:

(14) Braun hat {einen Wagen/ein Auto} gekauft. {Er/Es} ist grün.
Braun has bought {a car$_{neut}$/a car$_{masc}$}. {Pro$_{neut}$/Pro$_{masc}$} is green.

On pains of unintelligibility, the pronoun in the second sentence has to agree in gender with its antecedent in the first, and although the term 'grammatical gender' may suggest otherwise, this is not a grammatical phenomenon. Neither pronoun in (14) is bound syntactically: they are perfectly ordinary referential anaphors, whose duty it is to retrieve a discourse referent from the context. What makes these pronouns special is the requirement that, to a first approximation at least, the last mention of their referents should have employed an expression of the same gender.

That this is not quite right yet appears from the fact that gender pronouns may be used deictically, that is, without a linguistic antecedent. For instance, a Frenchman watching someone trying to get a table ('la table$_{fem}$') into his car might remark (Tasmowski-De Ryck and Verluyten 1982):

(15) Tu n'arriverais jamais à {la/*le} faire entrer dans la voiture.
You'll never manage to get {pro$_{fem}$/*pro$_{masc}$} into the car.

Here the pronoun has to agree in gender not with an earlier expression, for there was no previous mention of the table, but rather with the noun that would have been used by default to refer to the table. Observations like this highlight the fact that grammatical gender, too, depends on the non-linguistic context for its interpretation, although the information it carries is of a linguistic nature, and must therefore be represented on a different level.

2.6 Direct Reference

Due to the work of Kripke (1972) and Kaplan (1989) it has become widely accepted that certain types of singular terms, especially proper names and indexicals, 'refer directly'. This is not to deny that these terms have descriptive content. For it is obvious that, for example, 'I' carries the information that the speaker is referring to himself, and that the name 'Brown' refers to someone who is called 'Brown'. It is just that this sort of content is not truth-conditional content (Geurts 1997).

Direct reference poses a problem for theories of meaning that treat all descriptive content alike. Kripke and Kaplan have argued against such mono-semantic accounts, observing that a sentence like (16a) does not have the same truth-conditions as (16b); the first is a contingent truth, while (16b) is necessarily true.

(16) a. Brown is called 'Brown'.
b. Brown is Brown.

The same point can be made with indexicals:

(17) a. I am the speaker.
b. The speaker is the speaker.

Whereas the proposition expressed by (17a) might have been false, (17b) is, on one of its readings at least, necessarily true.

Kaplan's well-known analysis of direct reference involves dividing the Fregean notion of sense into two components, which he calls 'character' and 'content'. The character of an expression is its linguistic meaning, which in a given context determines the expression's truth-conditional content. The descriptive content of an indexical is unlike that of a definite description in that it remains at the level of linguistic meaning, and doesn't enter the truth-conditional level. We will see later on how this distinction can be captured in our representational framework.

3 Binding Problems

We have discussed a number of phenomena that illustrate the mundane truth that different kinds of linguistic and para-linguistic information need to be kept apart. But although we must separate between different kinds of content, the separation had better be not too strict. This is the lesson taught by a problem first noted by Karttunen and Peters (1979), which has come to be known as the 'binding problem' of presupposition projection (we will shortly see, however, that the problem is quite general). Karttunen and Peters' example is the following:

(18) ?Someone managed to succeed George V on the throne of England.

This sentence has a question mark because it is pragmatically infelicitous: it suggests that the person who succeeded George V found it difficult to do so, which

can hardly be the case (at least not in the sense intended here; George V's successor may have had problems adjusting to his new station, but he obtained it without effort). Apparently, the presupposition triggered by the verb 'manage' fails in this case. The problem is that many theories of presupposition (including Karttunen and Peters' own) cannot account for this kind of infelicity, because they strictly separate between asserted and presupposed information, as a consequence of which the content of (18) is predicted to have the following components:

(19) Assertion: Someone succeeded George V on the throne of England. Presupposition: It was difficult for someone to succeed George V on the throne of England.

Unfortunately, the presupposition triggered by 'manage', thus construed, comes out true: practically everybody would have had a hard time succeeding George V. This is a problem not only for Karttunen and Peters' own treatment of presupposition, but for theories of a younger vintage, too. The problem arises, obviously, because presupposed and asserted content are separated too strictly, and it is the opposite from the problem discussed above: the DRT treatment of presupposition does not run into the binding problem because it keeps presuppositions and assertions together, though for other reasons they should be differentiated more than they currently are, as we have seen in Sect. 2.1.

Although there has been much discussion of the binding problem in the literature (Krahmer 1998; Beaver 2001), it has rarely been noted that the problem is not confined to presuppositions. But as van der Sandt (1992) points out, binding problems are liable to crop up whenever a strict separation is made between different kinds of information with interdependencies between them. We should expect, therefore, that they also arise in connection with implicatures—and they do:

(20) Some years ago, a young Russian pianist recorded some of the Beethoven sonatas.

Applying the standard Gricean reasoning, we observe that a speaker who uttered (20) could just as easily have made a stronger statement:

(21) Some years ago, a young Russian pianist recorded all the Beethoven sonatas.

Why didn't the speaker utter (21) rather than (20)? Presumably, because he believes that (21) isn't true.[2] But if (21) isn't true, then the speaker is committing himself to the claim that no young Russian pianist ever recorded all the Beethoven sonatas—which in a normal run of events would not be implied by an utterance of (20).

[2] We're cutting a few corners here for dramatic effect. See Geurts (2010) for extensive discussion.

As discussed at length by Geurts (2006, 2010), the problem with this pseudo-implicature is caused by the assumption that pragmatic reasoning is conducted solely in terms of sentence-sized semantic units that are disconnected from each other; so the heart of the trouble is the same as in the case of Karttunen and Peters' binding problem. If instead of asking ourselves why the speaker didn't say (21) instead of (20), we would have asked why the speaker didn't say that *the pianist in question* recorded all the Beethoven sonatas, we would have obtained the implicature that, to the best of the speaker's knowledge, the pianist in question didn't record all the Beethoven sonatas—which is correct. But this requires that the implicature is not fully segregated from the assertion: they are about the same individual.

In this section and the last one we have discussed various phenomena illustrating that different kinds of information need to be kept apart, but in such a way that certain interdependencies between them are captured. In the remainder of this chapter we present a unified account that attempts to accomplish just this.

4 Layered DRT

The basic idea underlying Layered DRT (or LDRT for short) is straightforward enough. It is that a discourse representation should consist of more than one layer of information. All the information that is exchanged between speakers will go into the same representation, but within this representation we want to distinguish between information that is asserted, presupposed, implicated, and so on. So within a layered DRS (LDRS) there will be layers for assertions, presuppositions, implicatures, grammatical features of utterances, formal properties of the discourse, and so on. In many cases, information will reside on a single layer, but occasionally the same information can be on more than one layer. This holds, in particular, for discourse referents, which may be seen as inter-layer communication switches.

Formally, layers are implemented as sets of labels on discourse referents and conditions. Every layer has its own label, and as the same piece of content may be on several layers at once, discourse referents and conditions will be assigned sets of labels. In the following, we will show how to add layers to the standard DRT language (Kamp 1981; Kamp and Reyle 1993).

4.1 The LDRT Language

The vocabulary of LDRT extends the standard DRT language with a set of layer labels. We start, as usual, with sets of discourse referents, predicates, and logical constants. All conditions in an LDRS will bear zero or more labels; discourse referents will be labeled, too, but only when they are introduced, not when they

occur as arguments.[3] Taking as given inventories of discourse referents, predicates, logical constants, and layer labels, the following clauses simultaneously define the set of LDRSs, labeled discourse referents and labeled conditions:

(22) In the following clauses, L may be any set of layer labels:

a. An LDRS φ is a pair $\langle U(\varphi), Con(\varphi)\rangle$, where $U(\varphi)$ is a set of labeled discourse referents and $Con(\varphi)$ is a set of labeled conditions.
b. If u is a discourse referent, then u_L is an L-labeled discourse referent.
c. If P is an n-place predicate and $u_1, \ldots, u_n$ are discourse referents, then $P_L(u_1, \ldots, u_n)$ is an L-labeled condition.
d. If u and v are discourse referents, then $u =_L v$ is an L-labeled condition.
e. If φ and ψ are LDRSs, then $\neg_L\varphi$, $\varphi \vee_L \psi$, and $\varphi \Rightarrow_L \psi$ are L-labeled conditions.

Furthermore, to be able to focus on an L-part of an LDRS, we define $U_L(\varphi)$ as the subset of discourse referents in $U(\varphi)$ whose label sets overlap with L, and similarly for $Con_L(\varphi)$:

(23) a. $U_L(\varphi) := \{u_K \in U(\varphi) \mid K \cap L \neq \emptyset\}$
b. $Con_L(\varphi) := \{\psi_K \in Con(\varphi) \mid K \cap L \neq \emptyset\}$

Officially, LDRSs are set-theoretic constructs, but unofficially we will employ the following notation, which we find easier to read. First, instead of $\langle\{u_1, \ldots, u_m\}, \{\varphi_1, \ldots, \varphi_n\}\rangle$ we will write $[u_1 \ldots u_m : \varphi_1 \ldots \varphi_n]$. Secondly, if $\{\alpha_1, \ldots, \alpha_n\}$ is a non-empty set of layer labels we will write $\alpha_1 \ldots \alpha_n$; hence x_{abc}, $P_{abc}(x_1, \ldots, x_n)$, etc. Thirdly, if a discourse referent or condition resides on all layers, we will omit the label set; for example, if there are only three labels, x_{abc} will sometimes be shortened to x (note the difference between x and $x_\emptyset$).

To illustrate how the LDRS language may be used, the interpretation of example (6), repeated here as (24a), may be rendered in LDRT as (24b):

(24) a. The porridge is warm.
b. $[x_p$: $\text{porridge}_p(x)$ $\text{warm}_a(x)$ $\neg_i[:$ $\text{hot}_i(x)]]$

The LDRS in (24b) has three layers: a, p, and i, which contain asserted, presupposed, and implicated material, respectively. The intuitive interpretation of (24b) is that it is presupposed that there is some x that is porridge, that x is asserted to be warm, and that x is implicated not to be hot. Note that the asserted and implicated parts use a discourse referent that is presupposed: we simply cannot say what (24a) asserts or implicates without referring to whatever it is that is being presupposed

[3] Cf. Maier (2006, 2009) for a modification of the current version of LDRT, in which discourse referents are never labeled (except in preliminary DRS structures, where labels indicate layered resolution restrictions, which we will not discuss here). In the terms of this paper, Maier assumes that every discourse referent carries all available labels. Intuitively, this may seem to give rise to unwanted existence claims, but in fact these are quite harmless, as long as conditions are sensibly labeled.

by the subject term. This is to say that structures like (24b) cannot be interpreted as consisting of three fully independent layers of information: discourse referents serve to connect information on different layers.

4.2 Semantics of LDRT

The idea underlying our semantics for the LDRS language is simply that, instead of specifying what is the truth-conditional content of an LDRS φ, we have to define what is the truth-conditional content of a selection L of layers in φ. That is to say, if φ contains a condition ψ_K, where K is the layer set associated with ψ, ψ_K is to be ignored unless K and L overlap; i.e. unless $K \cap L \neq \emptyset$.

In the standard semantics of DRT, an embedding function f is said to be extended by another function g, with respect to a given DRS φ, iff $f \subseteq g$ and $dom(g) = dom(f) \cup U(\varphi)$. Here we extend f only with discourse referents with relevant labels:

(25) $f[^{\varphi}_{L}]g :=$ $f \subseteq g$ and $dom(g) = dom(f) \cup U_L(\varphi)$

Given a world w, a label set L, and an embedding function f, the principal semantic object associated with an LDRS φ is $\|\varphi\|^f_{L,\mathrm{w}}$, which, if defined, is the set of embedding functions g that extend f and make the L-part of φ true at world w. If $\|\varphi\|^f_{L,\mathrm{w}}$ is a non-empty set, the L-part of φ is true at w; if $\|\varphi\|^f_{L,\mathrm{w}} = \emptyset$, the L-part of φ is false at w; and otherwise $\|\varphi\|^f_{L,\mathrm{w}}$ is undefined. If φ is a labeled condition, $\|\varphi\|^f_{L,\mathrm{w}}$, if defined, is either $\top$ or $\bot$.

Let M = ⟨D, W⟩ be a model, where D is a domain of individuals and W is a set of interpretation functions ('worlds'); w ∈ W; f is a partial function from the set of discourse referents into D; and L is a set of layer labels:

LDRSs: definedness and interpretation

Let φ be an LDRS. Then:

a. $\|\varphi\|^f_{L,w}$ is defined iff $\exists g$: $f[^{\varphi}_{L}]g$ and $\forall \psi \in \mathrm{Con}_L(\varphi)$: $\|\psi\|^g_{L,w}$ is defined.

b. If defined, $\|\varphi\|^f_{L,w} = \{g \mid f[^{\varphi}_{L}]g$ and $\forall \psi \in \mathrm{Con}_L(\varphi)$: $\|\psi\|^g_{L,w} = \top\}$.

Labeled conditions: definedness

c. $\|\mathrm{P}_K(u_1, \ldots, u_n)\|^f_{L,w}$ is defined iff $\{u_1, \ldots, u_n\} \subseteq dom(f)$.

d. $\|u =_K v\|^f_{L,w}$ is defined iff $\{u, v\} \subseteq dom(f)$.

e. $\|\neg_K \varphi\|^f_{L,w}$ is defined iff $\|\varphi\|^f_{L,w}$ is defined.

f. $\|\varphi \vee_K \psi\|^f_{L,w}$ is defined iff $\|\varphi\|^f_{L,w}$ and $\|\psi\|^f_{L,w}$ are defined.

g. $\|\varphi \Rightarrow_K \psi\|^f_{L,w}$ is defined iff $\|\varphi\|^f_{L,w}$ and $\|\varphi \oplus \psi\|^f_{L,w}$ are defined, where $\varphi \oplus \psi = \langle U(\varphi) \cup U(\psi),\ Con(\varphi) \cup Con(\psi)\rangle$ (LDRS-merge).

Labeled conditions: interpretation

If φ_K is a labeled condition, $\|\varphi_K\|^f_{L,w} = \top$ iff $\|\varphi_K\|^f_{L,w}$ is defined and one of the following holds:

h. φ_K is of the form $P_K(u_1,\ldots,u_n)$ and $\langle f(u_1),\ldots,f(u_n)\rangle \in w(P)$.
i. φ_K is of the form $u =_K v$ and $f(u) = f(v)$.
j. φ_K is of the form $\neg_K\psi$ and $\|\psi\|^f_{L,w} = \varnothing$.
k. φ_K is of the form $\psi \vee_K \chi$ and $\|\psi\|^f_{L,w} \cup \|\chi\|^f_{L,w} \neq \varnothing$.
l. φ_K is of the form $\psi \Rightarrow_K \chi$ and $\forall g \in \|\psi\|^f_{L,w}$: $\|\chi\|^g_{L,w} \neq \varnothing$.

$\|\varphi_K\|^f_{L,w} = \bot$ iff $\|\varphi_K\|^f_{L,w}$ is defined and $\|\varphi_K\|^f_{L,w} \neq \top$.

If φ is an LDRS, $\|\varphi\|^f_{L,w}$ is the set of embedding …

If φ is an LDRS, $\|\varphi\|^f_{L,w}$ is the set of embedding functions that extend f and that make the L-part of φ come out true in w. Shifting to a more general notion of content, the following clauses define the set of worlds in which φ's L-layers are true:

(27) $\|\varphi\|^f_L = \{w \mid \|\varphi\|^f_{L,w} \neq \emptyset\}$, if $\exists w$: $\|\varphi\|^f_{L,w}$ is defined; undefined otherwise.
(28) $\|\varphi\|_L = \|\varphi\|^{f_0}_L$, where f_0 is the empty function.

In LDRT, every choice of labels engenders its own sort of information. For example, $\|$ (24b) $\|_{\{p\}}$ is the set of worlds that contain porridge; $\|$ (24b) $\|_{\{p,a\}}$ is the set of worlds with warm porridge; and $\|$ (24b) $\|_{\{p,a,i\}}$ is the set of worlds containing porridge that is warm but not hot. Note that, for example, $\|$ (24b) $\|_{\{a\}}$ and $\|$ (24b) $\|_{\{i\}}$ are undefined. This is because the a- and i-layers of (24b) use a discourse referent that is introduced in the p-layer. The undefinedness of $\|$ (24b) $\|_{\{a\}}$ and $\|$ (24b) $\|_{\{i\}}$, as opposed to the definedness of $\|$ (24b) $\|_{\{p\}}$, is due to the fact that the assertion and implicature of (24b) depend on what the sentence presupposes, but not vice versa. What (24b) asserts or implicates can only be specified relative to a given value (or range of possible values) of the discourse referent x.

If we collect all labels into one set L, then $\|\cdot\|_L$ captures everything that is somehow expressed by a sentence or discourse. We have seen that such content aggregates need not be consistent:

(29) a. That is a beautiful painting. (= (8))
b. $[x_k : \text{dem}_k(x)\ \text{beautiful-painting}_a(x)\ \neg_i[: \text{beautiful-painting}_i(x)]]$

On the intended reading of (29a), the sentence is meant to convey that the object under discussion is in fact not a beautiful painting at all. Assuming for the nonce that this is an implicature, and that demonstratives are represented on a special k-layer (more about which in Sect. 5), we obtain a representation along the lines of (29b). The literal meaning of (29a) is $\|$ (29b) $\|_{\{k,a\}}$; the implicit message is $\|$ (29b) $\|_{\{k,i\}}$; and $\|$ (29b) $\|_{\{k,a,i\}} = \emptyset$.

One application of the rich representations and flexible semantics of LDRT is in the analysis of denials. In Sect. 2.1, we saw how denials can target different layers, e.g. asserted or presupposed information, and they can also be directed at several layers at once (see Maier and der Sandt 2003 for discussion and an analysis of denial in LDRT). In the following, we will demonstrate the power of LDRT with two other applications, each of which will be seen to require some minor additions to the basic semantics presented above.

5 Indexical Content: The k-Layer

The LDRS-semantics given in (26) is uniform in the sense that, once a group of layers have been selected, all layers are treated alike. (That is, they are treated alike by the semantics. If layers were alike in every respect, there would be no point in having them in the first place.) In the remainder of this chapter we discuss two classes of phenomena which show that this is not quite right, and adjust our semantics accordingly.

Standard DRT has trouble with names, indexicals, and demonstratives because it has no way of separating descriptive content from contextual, 'reference fixing' content. In LDRT we can simply represent the two types of content on two different layers: 'a' is for asserted content, and 'k' for contextual, rigid content. We already used these layers in the previous section (cf. (29b)); a crucial example is the following, which we encountered before in Sect. 2.6:

(30) a. I am the speaker. (= (17a))
b. $[x_k : \text{speaker}_k(x)\ \text{speaker}_a(x)]$

However, it is not enough just to put indexical content on a layer of its own: $\| \text{(30b)} \|_{\{k,a\}}$ merely says that there is a speaker, whereas it should say of the individual who *in fact* is doing the talking that he is the speaker. In order to account for this, we follow Kaplan by making the content of an LDRS dependent on the context in which it occurs. A context may be regarded as a small world in the sense that it determines a unique speaker, hearer, etc. In our LDRS-semantics worlds are identified with interpretation functions, so if c is a context, then c(speaker), c(hearer), c(now), etc., are singleton sets. There are various ways of enforcing this restriction, two of which will be demonstrated in the following.

Relative to a given context c, we define the indexical content of the *L*-part of an LDRS φ as follows:

(31) $\mathbb{I}_{L,c}(\varphi) = \|\varphi\|^{\iota}_{L}$, with ι being the unique element of $\| \varphi \|^{f_0}_{\{k\},c}$, if such exists; otherwise undefined.

If $\|\varphi\|^{f_0}_{\{k\},c}$ is not a singleton set, c fails to determine unique values for all discourse referents in the k-layer, and $\mathbb{I}_{L,c}(\varphi)$ is undefined. Otherwise $\mathbb{I}_{L,c}(\varphi) = \|\varphi\|^{\iota}_{L}$, where ι is the unique embedding function determined by c. For example, if White is the

speaker in context c, $\mathbb{I}_{\{k,a\},c}(30b) = \| (30b) \|^{\{\langle x, White\rangle\}}_{\{k,a\}}$ = the set of worlds in which White is the speaker.

This type of LDRT implementation of direct reference is further developed and defended by Maier (2009), where it is combined with a theory of layered presupposition resolution. Although it is an attractive way of thinking about rigidity in DRT, we would like to explore here also a slightly different way of implementing a Kaplan-style notion of content in LDRT, which incorporates context parameters in the definition of $\| \cdot \|$. The main difference is that we can then interpret embedded k-layers, which is a feature that we need not so much for Kaplanian rigidity, but for our treatment of formal content in Sect. 6. We have to clear several choice points if we take this line, but the simplest solution we can think of mainly affects the interpretation of atomic conditions, which now comes out as follows:

a. $\|P_K(u_1,\ldots,u_n)\|^{f,c}_{L,w}$ is defined iff $\{u_1,\ldots,u_n\} \subseteq dom(f)$ and if $k \in K \cap L$, then $|c(P)| = 1$.
b. $\|P_K(u_1,\ldots,u_n)\|^{f,c}_{L,w} = \top$ iff $\|P_K(u_1,\ldots,u_n)\|^{f,c}_{L,w}$ is defined and one of the following holds:
   – $k \notin K \cap L$ and $\langle f(u_1),\ldots,f(u_n)\rangle \in w(P)$;
   – $k \in K \cap L$ and $\langle f(u_1),\ldots,f(u_n)\rangle \in c(P)$.

In other words: If an atomic condition is not on the k-layer, its interpretation is as in (26). If it is on the k-layer, its semantic value is undefined if the current context c fails to assign a unique value to its predicate. Suppose again that c(speaker) = White; then $\| (30b) \|^{f_0,c}_{\{k,a\},w} = \emptyset$ if White is not the speaker at w, or else $\| (30b) \|^{f_0,c}_{\{k,a\},w} = \{g\}$, where $dom(g) = \{x\}$ and $g(x)$ = White. Hence, $\| (30b) \|^{c}_{\{k,a\}} = \mathbb{I}_{\{k,a\},c}(30b)$ = the set of worlds in which White is the speaker.

So the two methods produce the same result in this case, as they do in many others, but they are not fully equivalent. First, and most importantly, with the second method, all discourse referents and conditions labeled 'k' are interpreted at the contextual index. The first method by contrast presupposes, in effect, that all k-material resides in the principal LDRS, or otherwise it will not be interpreted at c. Secondly, whereas the second method requires that each k-marked predicate be unique, the first method is less stringent in this respect, since it requires merely that, between them, the conditions in the k-layer determine unique values for all k-marked discourse refererents. The choice between these methods depends on considerations that go beyond the scope of this chapter. We should like to note, however, that on the whole the second method is more versatile, and brings out more clearly the relation between indexical content and what we call 'formal content', which is the topic of the next section.

6 Formal Content

We argued in Sect. 2.4 that we should aim for an analysis of expressions like 'the former' and 'the latter' which treats them as regular definite descriptions whose content is somewhat special. In this section we shall see that LDRT can provide us with such an analysis. As it turns out, 'the former', 'the latter', and their kin are context dependent in a way that resembles the context dependence of indexical expressions.

The LDRS in (33b) is a first stab at capturing the intuition that the expressions 'the former' and 'the latter' as used in (33a) are presuppositional devices whose content refers to formal properties of the previous discourse:

(33) a. As the beggar approached the bishop, the latter blessed the former.
b. [x_p, y_p: $beggar_p(x)$ $bishop_p(y)$ $approach_a(x,y) x \prec_{pf} y$ $bless_a$ (y,x)]

The intended interpretation of condition '$x \prec_{pf} y$' is something like: 'The last mention of x precedes the last mention of y.' (There may be other, and perhaps better, ways of rendering the meanings of 'the former' and 'the latter', but this one will do for our purposes.) This condition is part of the presupposition triggered by 'the former/latter', so it is on the p-layer, and it is also on the f-layer, because it refers to the form of the preceding discourse.

There is one problem with this proposal, for which we shall present a tentative solution. As it stands, our treatment of the interpretation of 'the former/latter' and related devices, such as grammatical gender, presupposes that the properties denoted by these expressions are properties of regular individuals: $'X \prec_{pf} Y'$ is true iff the last mention of the individual associated with x preceded the last mention of the individual associated with y. Example (10), repeated here as (34a), with its LDRS in (34b), demonstrates that this is not correct in general:

(34) a. If a beggar meets a bishop, then the latter will bless the former.
b. [: [x_a, y_a : $beggar_a(x)$ $bishop_a(y)$ $x \prec_{pf} y$ $meet_a(x,y)$] ⇒[: $bless_a(y,x)$]]

That is, for all beggar–bishop pairs in which the beggar is mentioned before the bishop, the second blesses the first. In order to heighten the dramatic impact of the example, we might suppose that the conditional is given a modal interpretation, and is construed as quantifying over worlds. But whatever the sentence quantifies over, its domain is not confined to states of affairs containing pairs of persons one of whom was mentioned before the other. The only mentioning event that is relevant is the actual utterance, which is part of the actual context.

If the predicate '$\prec$', as used in (34b), is not about beggars and bishops, be they possible or real, then what *is* it about? The answer, we would like to suggest, is 'discourse referents': the condition '$x \prec y$' states that the most recent use of x preceded the most recent use of y. Hence, a speaker who employs the 'former/latter' idiom in effect instructs the hearer to retrieve from the context a pair of

recently-used discourse referents. Note that this requires a notion of context that goes beyond the original Kaplanian context of utterance in containing not only individuals but also a structured representation of the actual utterance, as well as the discourse referents that that utterance has given rise to.[4]

In order to implement this idea, we propose to expand the interpretation of atomic conditions given in 4 along the following lines:

a. $\|P_K(u_1,\ldots,u_n)\|^{f,c}_{L,w}$ is defined iff one of the following holds:
   - k $\in K$ and $\{u_1,\ldots,u_n\} \subseteq dom(f)$ and $|c(P)| = 1$;
   - f $\in K$ and $|c(P)| = 1$;
   - $\{f,k\} \cap K = \varnothing$ and $\{u_1,\ldots,u_n\} \subseteq dom(f)$;

b. $\|P_K(u_1,\ldots,u_n)\|^{f,c}_{L,w} = \top$ iff $\|P_K(u_1,\ldots,u_n)\|^{f,c}_{L,w}$ is defined and one of the following holds:
   - k $\in K \cap L$ and $\langle f(u_1),\ldots,f(u_n)\rangle \in c(P)$;
   - f $\in K \cap L$ and $\langle u_1,\ldots,u_n\rangle \in c(P)$;
   - $\{f, k\} \cap K \cap L = \varnothing$ and $\langle f(u_1),\ldots,f(u_n)\rangle \in w(P)$.

According to this analysis, the interpretation of f-marked conditions depends on the context, and in this respect f-conditions and k-conditions are alike. But on the other hand, f-conditions are special in that they are about discourse referents, whilst all other conditions, k-conditions included, are about 'real' things in the world.

7 Conclusion

Our main objective in this chapter was to provide a general framework for representing and integrating all sorts of information that may be conveyed by linguistic means. Our proposal is LDRT. From a syntactical point of view, LDRT is perfectly straightforward. Label sets allow us to separate between different types of information, without severing binding relations. This device is uniform in the sense that, syntactically speaking, the *only* difference between one type of content and another consists in the labels they bear. But of course different types of content will differ in other ways as well—if they didn't, they wouldn't have to be distinguished in the first place. Such differences may be procedural; for example, certain types of content are cancellable whilst others are not. Other differences may be semantic; for example, indexical and formal content are context dependent in distinctive ways. But at the representational level, all kinds of content are equal.

[4] Incidentally, such a notion of context seems extremely useful for the study of signed languages, where discourse referents correspond to actual, visible points in the signing space, which signers point to and keep track of in a discourse. Cf. Schlenker (2010) for a discussion of DRT discourse referents in the analysis of sign language.

Acknowledgments We would like to thank Rob van der Sandt for discussion. Emar Maier is supported by the EU under FP7 (ERC Starting Grant 263890-BLENDS). Bart Geurts's research is supported by a grant from the Netherlands Organisation for Scientific Research (NWO).

References

Beaver, D. (2001). *Presupposition and assertion in dynamic semantics*. Stanford: CSLI.

Gazdar, G. (1979). *Pragmatics: Implicature, presupposition, and logical form*. New York: Academic Press.

Geurts, B. (1997). Good news about the description theory of names. *Journal of semantics, 14*, 319–348.

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

Geurts, B. (2006). Implicature as a discourse phenomenon. In E. Puig-Waldmüller (Ed.), *Proceedings of Sinn und Bedeutung 11*, (pp. 261–275). Barcelona: Universitat Pompeu Fabra.

Geurts, B. (2010). *Quantity implicatures*. Cambridge University Press.

Kamp, H. (1981). A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, & M. Stokhof (Eds.), *Formal methods in the study of language*, (pp. 277–322). Amsterdam: Mathematical Centre Tracts 135.

Kamp, H., & Reyle, U. (1993). *From discourse to logic*. Dordrecht: Kluwer.

Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford: Oxford University Press.

Karttunen, L., & Peters, S. (1979). Conventional implicature. In C. Oh & D. A. Dinneen (Eds.), *Syntax and semantics 11: Presupposition* (pp. 1–56). New York: Academic Press.

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

Kripke, S. (1972). Naming and necessity. In G. Harman & D. Davidson (Eds.), *Semantics of natural language* (pp. 253–355). Dordrecht: Reidel.

Maier, E. (2006). *Belief in context: Towards a unified semantics of de re and de se attitude reports*. Ph.D. thesis, Radboud University Nijmegen.

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

Maier, E. & van der Sandt, R. (2003). Denial and correction in layered DRT. In *Proceedings of DiaBruck 2003*. http://www.coli.uni-sb.de/diabruck/.

Schlenker, P. (2010). Donkey anaphora in sign language I: E-type vs. dynamic accounts. In M. Aloni (Ed.), *Logic, language and meaning: Proceedings of the Amsterdam Colloquium 2009*, (pp. 405–415).

Tasmowski-De Ryck, L., & Verluyten, P. (1982). Linguistic control of pronouns. *Journal of semantics, 1*, 323–346.

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

von Fintel, K. (2004). Would you believe it? The king of France is back! Presuppositions and truth-value intuitions. In M. Reimer, & A. Bezuidenhout (Eds.), *Descriptions and beyond*, (pp. 269–296). Oxford University Press.