

Pragmatic Discourse Representation Theory

Yafa Al-Raheb

National Centre for Language Technology
Dublin City University
Ireland
yafa.alraheb@gmail.com

Abstract

This paper presents a pragmatic approach to Discourse Representation Theory (DRT) in an attempt to address the pragmatic limitations of DRT (Werth 1999; Simons 2003). To achieve a more pragmatic DRT model, this paper extends standard DRT framework to incorporate more pragmatic elements such as representing agents' cognitive states and the complex process through which agents recognize utterances employing the linguistic content in forming mental representations of other agent's cognitive states. The paper gives focus to the usually ignored link in DRT literature between speaker beliefs and the linguistic content, and between the linguistic content and hearer's beliefs.

1 Introduction

Developments in dynamic semantics, resulting in DRT, have led to a framework suitable for the representation of linguistic phenomena (van Eijck and Kamp 1997). This is specifically due to the fact that, recognizing the importance of context, DRT concentrates on updating the context with the processing of each utterance. In addition, DRT can also be viewed as an agent's mental model of the world and not just a representation of the discourse. It is for these reasons that DRT holds great potential for incorporating more pragmatic phenomena.

However, despite the suitability of DRT for representing linguistic phenomena, some pragmatic limitations have been noted in the literature. Simons (2003) remarks that DRT is a theory of semantics and not pragmatics. Werth remarks that 'there

is no place in [DRT] for participant roles, setting, background knowledge, purposes, even inferences' (Werth 1999: 65). In general terms, we can say that the pragmatic dimension supplements semantic content by using context and cognitive states of agents in dialogue. The discipline of pragmatics is, therefore, concerned with the process by which agents infer information about elements of another agents' cognitive state such as their beliefs and intentions. Thus, this paper focuses on extending standard DRT pragmatically to model agents' cognitive states in the pragmatic context of dialogue.

2 A More Pragmatic DRT

This section presents a more pragmatic DRT focusing on the relationship between speaker generation and the linguistic content, and between the linguistic content and hearer recognition. Figure 1 represents the link between our representation of the speaker's cognitive state, the speaker's linguistic content and the hearer's cognitive state or DRS (Discourse Representation Structure). This relationship has not to our knowledge been explored in the literature and deserves investigation.

Generally speaking, to generate an utterance, there would be some discrepancy between the speaker's beliefs and the speaker's beliefs about the hearer's beliefs. The discrepancy leads to an utterance, i.e. linguistic content. The linguistic content is the window the hearer has onto the speaker's state of mind. It is what influences hearer *recognition*. By analysis of the linguistic content provided by the speaker, the hearer can propose a hypothesis regarding the speaker's state of mind.




Figure 1: Speaker DRS, Linguistic Content and Hearer DRS

2.1 New DR-Structures

The DRT representation introduced here extends standard DRT language and structure resulting in a suitable pragmatic-based framework for representing this pragmatic link. Separate DRSs are created to represent each agent. DRSs get updated with each new utterance. Each DRS representing an agent's cognitive state includes the two personal reference markers 'i' and 'you'. When 'i' is used in a DRS, it refers to the agent's self within that DRS; i.e. if the agent is the speaker, then 'i' refers to the speaker in the entire DRS. To refer to the other agent, 'you' is used. To follow from the speaker's example, 'you' in this case refers to the hearer. To account for agents' cognitive states and their meta-beliefs, a sub-DRS representing the agent's cognitive state called the *belief DRS* is created to include the speaker's beliefs about the hearer's beliefs. Additionally, a new DRS for representing weaker beliefs called *acceptance* is introduced. The same level of embedding offered to belief DRSs is introduced in acceptance DRSs. Acceptance DRS includes the speaker's acceptance DRS as well as what the speaker takes the hearer to accept. Provided the speaker has sufficient information, the speaker can also have the embedded DRS within the acceptance DRS that represents what the hearer takes the speaker to accept.

In addition to expanding the belief DRS, each agent's cognitive state contains an *intention DRS*. Intention in the sense used here refers to the agent's goals in making an utterance, which are represented by the corresponding dialogue act marked in the intention DRS. The hearer's intention DRS represents the recognized utterance and contains elements of utterance-making generally associated with pragmatics such as the function of an utterance, its dialogue act. This pragmatic enriching strengthens the

link between an agent's intentions and the linguistic form uttered. What is proposed is that the intention DRS be designed to include the linguistic content provided within utterances.

To further enhance the link between agents' cognitive states and the linguistic content of their utterances, the intention DRS contains the rich pragmatic information offered by explicitly marking the presupposition (given information) and the assertion (new information) of the current utterance. The intention DRS is a separate DRS from the belief DRS. The beliefs of an agent give the motivation for making an utterance, and the intention DRS represents the speaker's intended message. The recognition of an utterance gives the hearer an insight into the agent's beliefs. Depending upon the particular dialogue represented, the intention DRS could have the speaker's intention, the hearer's intentions or both. The intention DRS functions as the *immediate context*, the one containing the utterance being generated or recognized. The belief and acceptance DRSs function as *background context* containing information pertaining to the dialogue and not just the current utterance. This division of labour context-wise is useful in that the information represented in the intention DRS directly feeds into the speaker's utterance, and is then inferred by the hearer through the linguistic content. The hearer's intention DRS includes the inferred speaker intentions in uttering the current utterance. This gives the flexibility of being able to model information that the hearer has inferred but has not yet decided to accept or believe and is, therefore, not yet included in either the belief or acceptance DRS. For instance, while the hearer in example (1) has recognized S1's utterance, he has not yet accepted S1's utterance. This motivates separating the representation of beliefs from intentions.

- (1) S1: Bob’s trophy wife is cheating on him.
H1: When did Bob get married?

2.2 Extending DRT Language

In addition to the three DRSs introduced above, in order to make the link between speaker generation, linguistic content, and hearer recognition more explicit, *labels*, ‘label_{*n*}’, *n* an integer, are introduced. The labels mark the distinction between presupposition and assertion, and the distinction between weak and strong beliefs. Furthermore, the labels can be used to refer to a particular predicate by another complex predicate. The labels increase the expressive power from an essentially first-order formalism to a higher-order formalism. Presuppositions are marked by a presupposition label ‘p_{*n*}’. Similarly, DRSs inside the main speaker or hearer DRS are labeled ‘drs_{*n*}’. Assertions are marked by ‘a_{*n*}’ to strengthen the connections between the linguistic form (in the separation between presupposition and assertion) and the representation of beliefs. Believed information labeled ‘b_{*n*}’ inside a belief DRS or accepted information labeled ‘c_{*n*}’ inside an acceptance DRS can be either presupposed or asserted inside the intention DRS. Thus, the labels in the intention DRS can only be ‘p’ or ‘a’.

Conditions referring to attitudes (acceptance, beliefs, and intentions) have been added to the extended semantics of DRT. Figure 2 shows three embedded DRSs, acceptance DRS, drs2, belief DRS, drs4, and intention DRS, drs6 representing:

- (2) A: Tom is buying Mary a puppy.
B: That’s sweet.

DRSs are referred to by the attitude describing them. For example, attitude(i, ‘BEL’, drs4) refers to the DRS containing the speaker’s beliefs, using the label for the belief DRS, drs4. Other conditions are allowed to employ ‘i’ as an argument. Attitude(i, ‘accept’, drs2) refers to the DRS containing the speaker’s acceptance DRS, using the label for the acceptance DRS, drs2. Attitude(i, ‘INT’, drs6) refers to the DRS containing the speaker’s intention in uttering example (2), using the label for the intention DRS, drs6. The speaker’s acceptance DRS contains an embedded DRS for the hearer’s acceptance DRS, drs2. In this case, it is empty, as no weakly believed propositions have been introduced yet. Simi-

larly, the belief DRS contains space for the speaker’s beliefs about the hearer’s beliefs, drs5. The intention DRS contains the linguistic content of the utterance that the speaker is about to make, drs6, as well as the relevant dialogue acts.

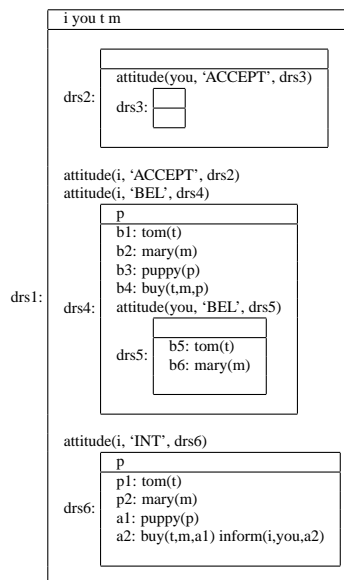


Figure 2: A’s initial Cognitive State

In Figure 2, there are essentially three levels of *embedding* in a main DRS. If we look at the belief DRS, the first embedded DRS is the agent’s own belief DRS. Level two is the agent’s beliefs about the other agent’s beliefs DRS. Level three is inserted when necessary and represents the agent’s beliefs about the other agent’s beliefs about the agent’s beliefs DRS. DRSs of the same level of embedding have similar status. For example, the agent’s acceptance and belief DRSs have equal status. However, the only discourse referents in common are the ones in the main DRS’s universe. Each equal-level embedding has its own set of discourse referents, as well as its own conditions.

Discourse referents of same and higher levels of embedding are accessible to lower levels of embedding and are therefore not represented in the lower level embedding universe. This does not entail that when a lower level embedding makes use of a discourse referent introduced in a higher level embedding the agent and other agent share the same internal or external anchors. For example, when talking about a rabbit, the speaker’s representation of rabbit

will be: $b1:rabbit(x)$, whereas the speaker's representation of the hearer's beliefs will be $b2:rabbit(x)$. This is to replace Kamp and Reyle's (1993) use of different discourse referents, where a new discourse referent is used every time the same object or individual is referred to in a new sentence (e.g. $rabbit(x)$, then $rabbit(y)$). The aim is to avoid having to overtly use the $x=y$ rule every time the same rabbit is referred to. The principles behind the equation predicate are still in place; i.e. every time rabbit is referred to, it is bound to the rabbit already in the context. However, we bind it to the previous properties of rabbit already in context through attaching it to the same discourse referent, $rabbit(x)$.

Both Kamp and Reyle's and our representation face revision when it transpires that the agents in dialogue have different referents in mind. For example, both the speaker and hearer might be talking about 'rabbit'. However, they might have a different 'rabbit' in mind, and assume the other participant is thinking of the rabbit they have in mind. The speaker might have a grey rabbit in mind, whereas the hearer has a white rabbit in mind. In this case, Kamp and Reyle's revision would consist of deleting $x=y$ predicate, and any previous equation predicate that may have been introduced each time rabbit was referred to. In our representation, the revision takes place by changing the other agent's discourse referent, $b2:rabbit(x)$ becomes $label2:rabbit(y)$.

Furthermore, the previous pragmatic extensions to standard DRT have been implemented computationally to approximate a computational model of communication and to enable us to see whether the extended DRT works logically. The implementation relates the linguistic content of utterances to the beliefs and intentions of the agents. The implementation operates with a specific dialogue, which can be modified, within a restricted domain. It seems reasonable to conclude on the basis of the implementation that the conceptual and formal proposals made provide a basis for further development.

3 Conclusion and Further Extensions

This paper pushes the treatment of linguistic phenomena in DRT more towards pragmatics, by bringing more pragmatic elements to the semantic/pragmatic interface which is DRT. It has

been the aim of this paper to achieve this by (a) expanding DRT structure to incorporate the pragmatic extensions introduced in this paper, (b) representing the complex process of speakers recognizing utterances and using the linguistic information in forming mental representations of hearers' mental representations, (c) enhancing the link between speaker beliefs, and between the linguistic content and the linguistic content and hearer's beliefs and (d) putting all these extensions and enhancements to the pragmatic side of DRT in a computational model.

While the work presented in this paper offers a more pragmatic approach to DRT, there is still more work to be done on making DRT more pragmatic. The possibility of extending the present treatment to include more agents remains for future work. In addition, future work can employ the intention DRS introduced in this paper, in order to enhance the complexity of the pragmatic representation of speaker/hearer intentions. For instance, embedding turn-taking acts within the intention DRS and relating them to agents' beliefs and intentions should be straightforward. It is also hoped that future work will address more aspects of context than the two detailed and implemented in this paper, namely, the immediate and background context. Furthermore, the sample implementation of the extensions suggested in this paper serves as an example of how the extensions to DRT can be implemented. One way of developing this implementation is to incorporate it into a dialogue system which aims to achieve a more balanced approach to the semantic/pragmatic interface in representing linguistic phenomena.

References

- Kamp, H. and Reyle, U. 1993. *From Discourse to Logic: Introduction to Model Theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Boston, Dordrecht: Kluwer.
- van Eijck, J. and Kamp, H. 1997. 'Representing Discourse in Context'. In: J. van Benthem and A. Ter Meulen (Eds.). *Handbook of Logic and Language*. pp. 179–237. Amsterdam: Elsevier.
- Simons, M. 2003. 'Presupposition and Accommodation: Understanding the Stalnakerian Picture'. *Philosophical Studies* 112, pp. 251–278.
- Werth, P. 1999. *Text Worlds: Representing Conceptual Space in Discourse*. New York: Longman.