

DORIS 2001: Underspecification, Resolution and Inference for Discourse Representation Structures

Johan Bos

Institute for Communicating and Collaborative Systems

Division of Informatics, University of Edinburgh

2 Buccleuch Place, Edinburgh EH8 9LW, Scotland, United Kingdom

`Johan.Bos@ed.ac.uk`

1 Introduction

DORIS (Discourse Oriented Representation and Inference System) translates English texts into Discourse Representation Structures (DRSs, from Discourse Representation Theory [Kam81]), thereby dealing with a range of linguistic phenomena including scope ambiguities, pronouns, presupposition triggers, plural noun phrases, and modal operators. DORIS is not only a framework for implementing semantic analysis—it is also a platform to investigate the use of first-order theorem proving technology within its semantic analysis.

DORIS was first launched in 1998, and since then has seen several technical improvements and increase of coverage of linguistic phenomena. The overall concept of DORIS has been reported in earlier work [BBKdN99]. The current paper not only describes DORIS in more detail, but also reports on recent changes in the system’s design, information flow and linguistic coverage. The latest version, DORIS 2001, uses compositional underspecified semantic representations, and an extended translation function from DRSs to first-order logic to cover a wider range of linguistic phenomena. DORIS 2001 is available via <http://www.coli.uni-sb.de/~bos/doris/>.

2 System Architecture and Functionality

Let’s start with describing the system architecture and functionality of DORIS. This will give us a rough idea of what the system is supposed to do and how it works. Three central themes motivate the design of DORIS: **underspecification**, **resolution**, and **inference**.

Underspecification. The input of DORIS is a (typed) English text (for a restricted domain). This text is clustered into sentences (using punctuation for segmentation clues) and fed to the parser which provides a syntactic analysis in the form of a tree structure. On the basis of the syntactic tree structure, semantic composition proceeds, producing underspecified discourse representations. Underspecification within DORIS focuses on anaphoric expressions (such as pronouns and presupposition triggers) and scope ambiguities (introduced by

quantifiers, negation, or modal operators). Section 3 gives a detailed account of underspecification in the composition process.

Resolution. The output of the parser are underspecified discourse representations. The next step in DORIS is to show how these can be resolved with respect to the DRS obtained from the input text so far (recall that the DRS gets incrementally updated, on the level of sentences clustered by the parser). Anaphoric constructs are resolved with respect to this DRS, following Van der Sandt’s algorithm for presupposition resolution [VdS92]. Scope ambiguities are encoded and resolved along the lines of Hole Semantics [Bos96]. Note that several DRSs could result from the resolution process, but the different solutions are ranked with a score obtained in our implementation of the presupposition resolution algorithm (accommodation is expensive, and yields a lower score than binding, for example). As DORIS implements its resolution tasks using known techniques, we won’t describe them here in detail and instead refer to [BB00] for an overview.

Inference. Finally, given this ordered set of DRSs, DORIS applies consistency and informativity checks to them, the so-called acceptability constraints (DRSs that fail these tests are rejected from the analysis). This requires serious reasoning, which is performed by translating the DRSs into first-order representations (Section 4), making use of theorem provers and model builders for first-order logic (Section 5).

Figure 1 wraps up the information flow within DORIS, and shows examples of the different kinds of representation associated with each level of analysis.

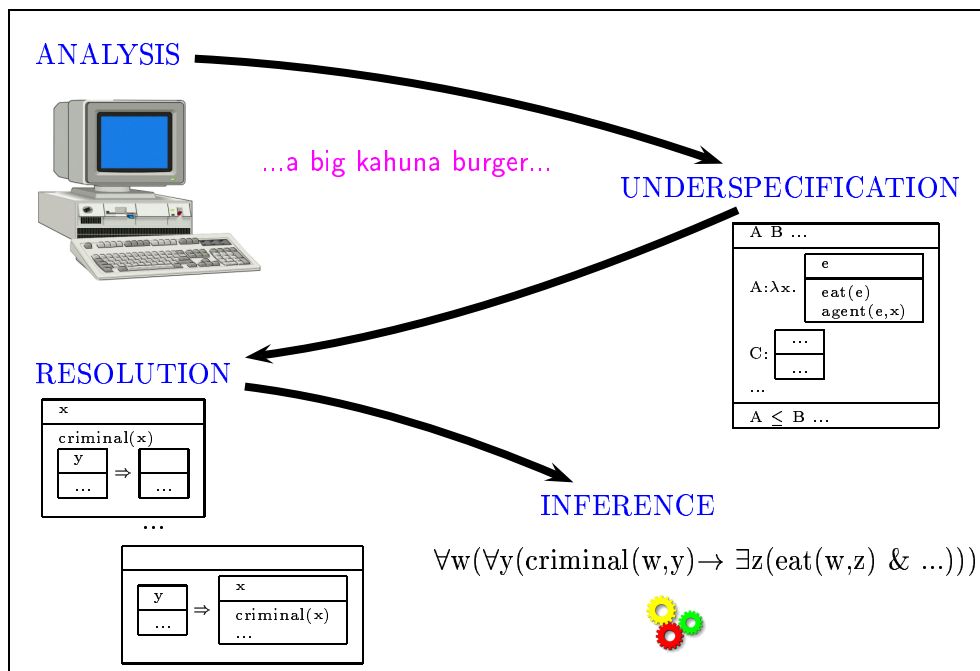


Figure 1: Information flow and representations used in DORIS 2001.

3 Compositional Semantics and Underspecification

DORIS combines the λ -calculus with known techniques for semantic underspecification, resulting in a compositional algorithm for constructing meaning representations. Two levels of representation are distinguished: the object level (which are DRSs or to be more precise λ -DRSs) and the meta level (describing the object level). In the composition process, λ -abstraction, functional application, and β -conversion are applied on both levels of representation. But let's first explain how underspecification is dealt with.

The underspecified discourse representations used in DORIS are based on Hole Semantics [Bos96]. Basically, the structures in Hole Semantics are three-place tuples consisting of a set of holes and labels, a set of labelled λ -DRSs, and subordination constraints. Holes and labels are meta-variables, where holes underspecify scope and labels point to specific λ -DRSs. However, note that holes and labels can occur in substructures of the labelled λ -DRSs as well (for lack of space we refrain from giving a definition of the syntax of underspecified discourse representations here). The subordination constraints are relations over labels and holes and express dominance within tree structures (after all, a DRS is a tree structure). For instance, $L \leq H$ states that H dominates L , or in other words, label L is in the scope of hole H . Resolving an underspecified representation amounts to plugging the holes with λ -DRS in such a way that none of the constraints are violated.

Compositional DRT systems marry type-theory with the basic DRS-language and use λ -abstraction, functional application, merge reduction (reducing a merged DRS (a DRS constructed out of two other DRSs) into one (logically) equivalent DRS), and β -conversion in the process of semantic construction [Mus96, KKP96]. The implementation of DORIS shifts these ideas to the level of underspecification, following three principles: (1) the object level representations are λ -DRSs, (2) the syntax of the meta level representations are extended to allow λ -abstraction and application, (3) merge reduction, and β -conversion is performed on the meta level in the composition process. Figure 2 illustrates this by showing the underspecified semantic representations of some phrases.

To conclude, consider the output of the parser in DORIS: this is an expression that yields an ordinary underspecified semantic representation (as in Hole Semantics) after applying β -conversion and merge reduction. Resolving this underspecified representation yields a set of expressions of λ -DRT. The result of applying β -conversion and merge reduction to each of these expressions is a set of 'ordinary' DRSs (see next section). So functional application and β -conversion are applied at both levels of representation.

4 Discourse Representation Structures

The syntax of the DRSs used by DORIS is as in standard DRT extended with modal operators. That is, a basic DRS is an ordered pair of a set of discourse referents and a set of DRS-conditions. The DRS-conditions are confined to: $R(x_1, \dots, x_n)$ where R is an predicate symbol with arity n and x_1, \dots, x_n are dis-

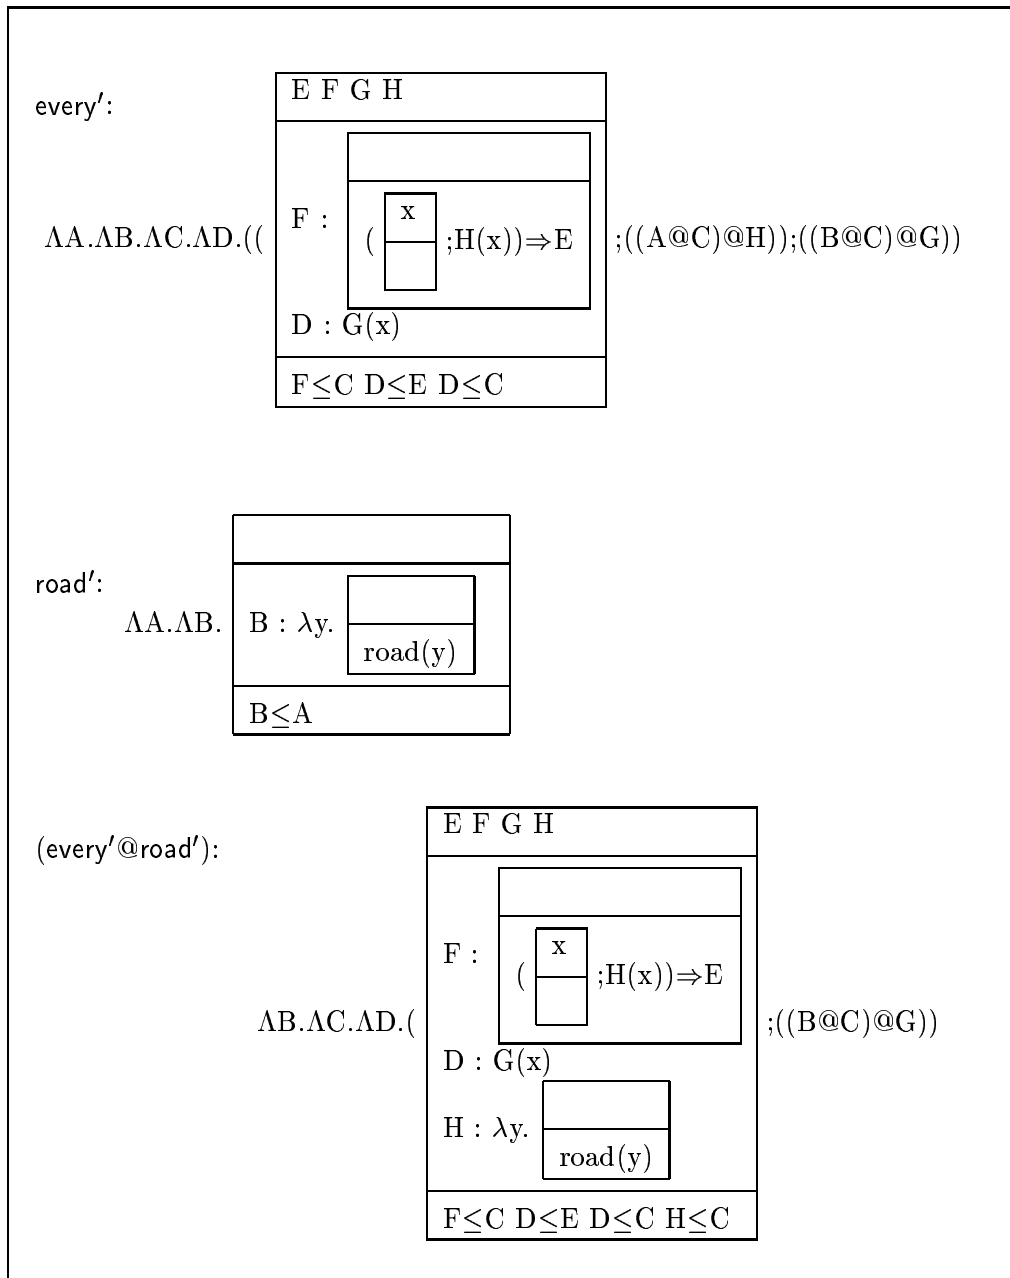


Figure 2: Translations (indicated by ') for the determiner every, the noun road, and the noun phrase every road (the result of combining the translations of every and road; for convenience, shown after β -conversion and merge reduction). Here @ and Λ stand for functional application and λ -abstraction at the meta level respectively, and the uppercase symbols are used to denote meta variables (i.e. holes and labels).

course referents (basic condition); $x_1=x_2$ for two discourse referents x_1 and x_2 (equality); $\neg B$, $\Box B$, $\Diamond B$, where B is a DRS (negation and modal operators); $B_1 \Rightarrow B_2$, $B_1 \vee B_2$, where B_1 and B_2 are DRSs (implication and disjunction); and finally $v:B$ where v is a discourse referent and B is a DRS (modal contexts).

Because the latter construct, $v:B$, is non-standard, it deserves some explanation. The domain D of interpretation is sorted, consisting of D_i (the individuals) and D_w (the possible worlds). So depending on their sort, discourse referents either denote individuals or possible worlds. What a DRS-condition $v:B$ in fact does is explicitly picking out a possible world, stating that the information expressed by the DRS B holds in the world denoted by v . This indeed shows similarities with hybrid breeds of modal logics [Bla00].

The interpretation of DRSs is done in an indirect manner, namely with the help of a translation function that maps our DRS language to first-order formulas. This translation is implemented as the function $(\cdot, \cdot)^{fo}$, from discourse referents (of the sort possible world) and DRSs to ordinary first-order formula syntax. The complete translation is shown in Figure 3.

$$\begin{aligned}
 (w, \frac{x_1 \dots x_n}{\gamma_1 \dots \gamma_m})^{fo} &\stackrel{\text{def}}{=} \exists x_1 \dots \exists x_n ((w, \gamma_1)^{fo} \wedge \dots \wedge (w, \gamma_m)^{fo}) \\
 (w, R(x_1, \dots, x_n))^{fo} &\stackrel{\text{def}}{=} R(w, x_1, \dots, x_n) \\
 (w, x_1 = x_2)^{fo} &\stackrel{\text{def}}{=} x_1 = x_2 \\
 (w, \neg B)^{fo} &\stackrel{\text{def}}{=} \neg (w, B)^{fo} \\
 (w, B_1 \vee B_2)^{fo} &\stackrel{\text{def}}{=} (w, B_1)^{fo} \vee (w, B_2)^{fo} \\
 (w, \frac{x_1 \dots x_n}{\gamma_1 \dots \gamma_m} \Rightarrow B)^{fo} &\stackrel{\text{def}}{=} \forall x_1 \dots \forall x_n (((w, \gamma_1)^{fo} \wedge \dots \wedge (w, \gamma_m)^{fo}) \rightarrow (w, B)^{fo}) \\
 (w, \Diamond B)^{fo} &\stackrel{\text{def}}{=} \exists v (R(w, v) \wedge (v, B)^{fo}) \\
 (w, \Box B)^{fo} &\stackrel{\text{def}}{=} \forall v (R(w, v) \rightarrow (v, B)^{fo}) \\
 (w, v : B)^{fo} &\stackrel{\text{def}}{=} (R(w, v) \wedge (v, B)^{fo})
 \end{aligned}$$

Figure 3: Translation from DRSs to First-Order Logic.

In essence, this translation is a combination of the standard translation for DRSs (e.g. [KR93]), the apparatus of Moore for dealing with the modal operators [Moo80], and the translations to first-order used in hybrid logics [Bla00]. Having first-order representations at our disposal opens the way for automated reasoning in DORIS, as nowadays many high speed theorem provers specialised for first-order problems are widely available.

5 Inference

DORIS implements Van der Sandt’s algorithm for presupposition projection which imposes checking the acceptability constraints: consistency and informativity. To illustrate *consistency*, suppose we have a DRS B . If we can prove that $\neg\exists w(w, B)^{fo}$ is valid, then we say that B is inconsistent. If, on the other hand, we find that $\exists w(w, B)^{fo}$ is satisfiable, we say that B is consistent. Now consider *informativity*, with respect to a DRS B_1 and a new DRS B_2 . If we find a proof for $\forall w((w, B_1)^{fo} \rightarrow (w, B_2)^{fo})$ we say that B_2 is not informative wrt B_1 . On the other hand, if we are able to show that both $\exists w((w, B_1)^{fo} \wedge (w, B_2)^{fo})$ and $\exists w((w, B_1)^{fo} \wedge \neg(w, B_2)^{fo})$ are satisfiable formulas, we say that B_2 is informative wrt B_1 .

This suggests that we not only need theorem provers (to check for validities) but also model builders (to check for satisfiable clauses), and this is indeed the case, although some theorem provers are able to detect satisfiability of clause sets [BBKdN99]. Moreover, although we are faced with the obvious limitations for reasoning with first-order logic (its undecidability and finite model satisfaction), DORIS tries to maximise its reasoning performance by using different inference engines concurrently. This is done with the help of MathWeb.

The MathWeb system distributes inference tasks among competing agents [FK99]. Using MathWeb as middleware, DORIS farms out consistency and informativity checking among off-the-shelf theorem provers and model builders for first-order logic, including BLIKSEM, FDPLL, MACE, OTTER, and SPASS. Using a distributed agent system such as MathWeb is motivated by two considerations: (1) DORIS generates a high number of independent problems at the same time, and (2) speed and coverage of the different provers available differ significantly. DORIS allows various options to use MathWeb: either by returning the first result available, or by showing all results of all selected inference engines.

6 Semantic Coverage (from a linguistic point of view)

With semantic coverage we not only refer to the process of deriving meaningful looking representations, but also performing the expected inferences with respect to consistency and informativity. This might involve additional background knowledge in the form of axioms, as is for instance the case for plural noun phrases, cardinals, and propositional attitude verbs.

With respect to noun phrases DORIS covers determiners (including quantifiers, the definite and indefinite article, cardinals, possessives), nouns, proper names, anaphoric and deictic pronouns. Most of these phenomena are implemented as DRT prescribes [KR93, VdS92], except for plural noun phrases. Plural objects are represented as discourse referents with a distinct sort (group), and membership relations between singular discourse referents with groups stipulate the constraints for plural noun phrases.

As for verb phrases, coverage include basic verbs, modal verbs, propositional attitude verbs, and *to be*. Basic verbs introduce discourse referents for

events, whereas identity statements or predications introduce states. Modal and propositional attitude verbs use the $v:B$ condition with additional axioms on accessibility of worlds to pin down their lexical semantics. Modifiers extend to modal adverbs (using the operators \Box and \Diamond), intersective adverbs (binding the event or state discourse referent), negation (as in DRT), presuppositional adjectives (such as *other*), prepositions, complementizers, relative clauses and certain cases of constituent coordination.

7 Prospects and Limitations

Disregarding some of the problem areas sketched below, DORIS performs deep semantic analyses in a reasonable amount of time. Take for instance a discourse introducing several discourse referents, some of them describing plural entities: An old dirty white Chevy barrels down a lonely street in Hollywood. Two young fellas are in the front seat. They are wearing cheap black coats. Jules is behind the wheel. The computing times for consistency verifications vary between less than a second up till four seconds.

An obvious question to ask is how far this enterprise will get us, or putting it differently: how far can we push first-order inference in natural language understanding. There is no straightforward positive answer to this question, as it partly depends on the amount of background knowledge supporting the inference problems. Background knowledge used in DORIS 2001 covers the axioms mentioned for plurals and modal relations, and ontological information (inheritance relations for nouns and verbs, and disjointness relations for adjectives and adverbs).

Some questions related to performance are easier to answer. The simplest examples that DORIS is able to generate and cause serious trouble for theorem provers (or model builders) are triggered by the use of cardinal expressions. To prove consistency for the sentence *two gangsters smoke* takes less than a second, stating that there is a group of individuals that solely has gangsters as members, and that there are actually two different individuals, that belong to this group. Increasing the group of smokers to three requires around two seconds. Four smoking gangsters take up more than 4 seconds, whereas five smokers already take 20 seconds. And six... well you've guessed rightly! Counting is not a particularly strong point of DORIS.

Future work includes extending the coverage with a treatment of tense (which is only handled in a trivial way in DORIS 2001) as well as dealing with actions. In fact, the machinery of DORIS is used as parser in a prototype dialogue system for instructing a mobile robot, as part of the IBL (Instruction-based Learning for Mobile Robots, GR/M90160) project [BLK⁺01].

References

- [BB00] Patrick Blackburn and Johan Bos. Representation and Inference for Natural Language. A First Course in Computational Semantics. Draft available at www.comsem.org, July 2000.

- [BBKdN99] Patrick Blackburn, Johan Bos, Michael Kohlhase, and Hans de Nivelle. Inference and Computational Semantics. In H.C. Bunt and E.G.C. Thijsse, editors, *Third International Workshop on Computational Semantics (IWCS-3)*, pages 5–19, 1999. Computational Linguistics, Tilburg University.
- [Bla00] Patrick Blackburn. Representation, reasoning, and relational structures: a hybrid logic manifesto. *Logic Journal of the IGPL*, 8(3):339–365, 2000.
- [BLK⁺01] Guido Bugmann, Stanislaw Lauria, Theodoros Kyriacou, Ewan Klein, Johan Bos, and Kenny Coventry. Using verbal instruction for route learning: Instruction analysis. In U. Nehmzow and C. Melhuish, editors, *TIMR 01 - Towards Intelligent Mobile Robots (TIMR 01)*. Department of Computer Science, Manchester University, 2001.
- [Bos96] Johan Bos. Predicate Logic Unplugged. In Paul Dekker and Martin Stokhof, editors, *Proceedings of the Tenth Amsterdam Colloquium*, pages 133–143, ILLC/Department of Philosophy, University of Amsterdam, 1996.
- [FK99] Andreas Franke and Michael Kohlhase. System description: Mathweb, an agent-based communication layer for distributed automated theorem proving. In *16th International Conference on Automated Deduction CADE-16*, 1999.
- [Kam81] Hans Kamp. A Theory of Truth and Semantic Representation. In Jeroen Groenendijk, Theo M.V. Janssen, and Martin Stokhof, editors, *Formal Methods in the Study of Language*, pages 277–322. Mathematical Centre Tract 135, Amsterdam, 1981.
- [KKP96] Michael Kohlhase, Susanna Kuschert, and Manfred Pinkal. A Type-Theoretic Semantics for λ -DRT. In Paul Dekker and Martin Stokhof, editors, *Proceedings of the Tenth Amsterdam Colloquium*, pages 479–498, ILLC/Department of Philosophy, University of Amsterdam, 1996.
- [KR93] Hans Kamp and Uwe Reyle. *From Discourse to Logic; An Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and DRT*. Kluwer, Dordrecht, 1993.
- [Moo80] Robert C Moore. Reasoning about knowledge and action. Technical Report 181, SRI International, California, 1980.
- [Mus96] Reinhard Muskens. Combining montague semantics and discourse representation. *Linguistics and Philosophy*, 19:143–186, 1996.
- [VdS92] Rob A. Van der Sandt. Presupposition Projection as Anaphora Resolution. *Journal of Semantics*, 9:333–377, 1992.