9

Labeling representations for effective disambiguation

Nicholas Asher and Tim Fernando

ABSTRACT. The effects of discourse on disambiguating (sub)expres-sions and the problem of computing those effects are treated formally by imposing a labeling structure over discourse representation structures, leading to barebones forms of Segmented and Underspecified Discourse Representation Theory.

Keywords: Ambiguity, underspecified semantic representations, information packaging, the semantics-pragmatics interface, dynamic interpretation in text

1. Introduction

From a formal semantic perspective, interpreting natural language utterances can be broken down into two tasks:

(I) Constructing logical forms or semantic representations of the utterances and

(II) Manipulating well-formed semantic representations.

Assuming semantic representations are drawn from a formal language, the problem (I) of *formalization* must be addressed before reasoning (II) *within* a formal system can commence. In *Discourse Representation Theory* (DRT, Kamp and Reyle [**10**]), the point is simply that a *Discourse Representation Structure* (DRS) must be associated with a natural language statement before one can speak of model-theoretic interpretations that lie behind the manipulations in (II), and which, in turn, provide the reason for bothering with (I).

Experience has shown that picking out a formal semantic representation faithful to a natural language statement is extremely difficult, resulting in pressure to blur the distinction between (I) and (II). Hence, the interest in *ambiguous* and *underspecified* semantic representations, which reduce the distance between natural language utterances and semantic representations. The price for shifting some of the burden of formalization (I) over to (II) is a gap within the formal system between on, the one hand, ambiguous and underspecified representations, and, on the other hand, unambiguous and fully specified representations (required to make sense of ambiguity and underspecification). The present paper is an attempt to combine two extensions of DRT, *Underspecified Discourse Representation Theory* (UDRT, Reyle [**13**]) and *Segmented Discourse Representation Theory* (SDRT, Asher [**1**], Lascarides and

Asher [**11**]), that concentrate on two complementary aspects of this gap: the first, on representing ambiguity; the second, on doing away with it. An example of an ambiguous expression excluded from Kamp and Reyle [**10**] is $x =?$, described in Asher [**1**] (following DRT folklore) as "an incomplete condition" that "anaphora resolution will make complete" (p. 73).[1] It is to complete such conditions (by supporting anaphora resolution and various other processes involved in discourse interpretation) that a system of *discourse relations* and *segmented discourse representation structures* (SDRSs) is introduced in SDRT. As shown in Asher [**2**], this extension can be presented perspicuously through a labeling system, similar to that used in UDRT for scope ambiguities. The present paper extends the use of labels to different levels (viz., word, sentence, discourse), with the intuition that the colon : between the label l and the expression e in $l : e$ serves as a "fence" for keeping the uncomputable content of e from spoiling the computational demands on disambiguation. Exactly what we mean by this will occupy much of the paper. At the end, we sketch some ideas on how to deal with ambiguity with a mix of "wait and see" and non-monotonic strategies.

2. From truth conditions to information packaging

To motivate our approach, we begin by recalling a formal distinction drawn in DRT between *conditions* and *discourse representation structures* (DRSs). A condition φ can be thought of as an ordinary "static" (truth-conditional) formula, interpreted (semantically), relative to a fixed set *MEP* of *model-embedding pairs*, as the set $[\![\varphi]\!] \subseteq MEP$ of model-embedding pairs at which it is *true* (à la Tarski, where an embedding assigns variables objects in the model paired with it):

$$[\![\varphi]\!] \quad = \quad \{s \in MEP \mid \varphi \text{ is true at } s\} \ .$$

By contrast, a DRS is a pair (U, C) consisting of finite sets U and C of discourse referents (or variables) and of conditions respectively, which is "dynamic" insofar as its meaning is taken to be a binary relation $[\![U, C]\!] \subseteq MEP \times MEP$ on model-embedding pairs, specifying its *input/output* behavior:

$$\begin{aligned}[\![U, C]\!] \quad = \quad \{(s, s') \in MEP \times MEP \mid & \text{ on input } s, \\ & (U, C) \text{ can output } s'\} \ .\end{aligned}$$

As the set C of conditions in a DRS (U, C) can be conjoined[2] into a single condition φ_C so that $[\![U, C]\!] = [\![U, \{\varphi_C\}]\!]$, and as $[\![\emptyset, \{\varphi_C\}]\!] = \{(s, s) \mid s \in [\![\varphi_C]\!]\}$, the difference between conditions and DRSs is that a DRS may have a non-empty set of discourse referents (i.e., U need not be $\emptyset$).

[1]The string $x =?$ should *not* be confused with the perfectly unambiguous random assignment $x :=?$, employed in Groenendijk and Stokhof [**7**] to assign an indeterminate value to a variable introduced by an indefinite, or with the corresponding DRS $(\{x\}, \emptyset)$.

[2]If necessary (following the strict conventions of DRT, according to which logical connectives are defined on DRSs) by first turning a condition φ into a DRS $(\emptyset, \{\varphi\})$.

The set U of discourse referents in a DRS (U, C) makes a genuine difference in discourse interpretation inasmuch as it provides some information concerning what the discourse (represented by (U, C)) is *about*. In particular, U supplies possible referents for anaphora resolution. Clearly, however, U need not encode sufficient information either to make a choice, in cases such as (1) where more than one candidate presents itself, or to construct the required referent, in say bridging examples like (2).

(1) Ernest has a brother. He feeds him.

(2) Ernest has a bike. The seat is bent out of shape.

One might turn to the set C of conditions for help — but again without any guarantee of finding it. At least two problems plague a truth-conditional approach to discourse interpretation:

(P1) the prohibitive costs of extracting truth-conditional information

and

(P2) the coarseness of grain of what can be extracted truth-conditionally.

By (P1), we have in mind, for instance, the undecidability of first-order logic, rendering consistency checks hopeless (non-r.e.). As for (P2), the difficulty is that the information typically encoded into C is, if abstracted from all matters but truth, not quite fine enough, even when supplemented with U; an example is anaphora resolution in (3), for which let us agree that `having a brother` is reduced to the condition of being in the domain of a symmetric binary relation.

(3) Ernest has a brother. He also has a pigeon.

To overcome obstacles (P1) and (P2), we will magnify, rather than minimize, the distinction between DRSs and conditions, by enriching the set U of discourse referents in a DRS (U, C) with more structure. Insofar as truth occupies a privileged place in semantics, it makes sense to dedicate C solely to truth conditions, consigning to our enrichment of U other matters, which may nonetheless bear on C, through a subsequent process of discourse interpretation. In particular, the aforementioned expressions $x =?$ (which are not among the conditions or DRSs assigned an unambiguous semantics via $[\![\cdot]\!]$) can be collected to pick out a set $U_? \subseteq U$ of discourse referents x requiring anaphora resolution; and the set U can then be expanded into the structure $\langle U, U_? \rangle$. Of course, in practice, it may be convenient to write $x =?$ among the other conditions in C (in the lower portion of the box). But even then, we ought not to forget that quite unlike conditions in C, the expression $x =?$ is merely a signal to some anaphora resolution module for attention, to be discarded as soon as it has spawned a genuine (if not true) condition. But as we have already mentioned, anaphora resolution may require much more than just U and C, suggesting that the structure $\langle U, U_? \rangle$ be enriched further. Without spelling out just now what that is, let us write P for that enrichment, the idea being that P describes information *p*ackaging, to complement the information *c*ontent kept in C.

Having passed from the DRS (U, C) to the pair (P, C), let us next relate the information packaging/content dichotomy to ambiguity. If any consensus has developed at all about formalizing ambiguity, it is the necessity of locating the uncertainty characteristic of ambiguity at a meta-level, removed from the object level where unambiguous expressions lie. Inasmuch as the conditions in C are unambiguous, it is natural to equate that meta-level with P, and to think of $x =?$ as an ambiguous expression subject to a form of disambiguation called anaphora resolution. The meta/object-level division between P and C is complex and involves rules not only for

(a) discriminating between information packaging (P) and information content (C),

but also for

(b) communication between P and C.

The motivation for item (a) comes from the problem (P2) of coarseness of grain in C, and also from the costs, (P1), of retrieving information from C, which is, in addition, crucial for item (b). The idea, put in operational terms, is that the process of discourse interpretation should have immediate unrestricted two-way access to P, but not to C, for which it must, in view of (P1), submit itself to a rigorous border control regime. By "two-way access to P", we mean that the discourse interpretation process should not only be able to feed information into P (thereby updating it, as well as C), but also request information from P. But then if the discourse interpretation process is to be, in some sense, "effectively computable," then so too must the logic for P. Hence, another reason for separating out P from C is that it would be wishful thinking to expect the logic for C to be effectively computable. By contrast, insofar as information packaging must necessarily be, in some sense, finite (if only because of the infeasibility of unpacking something infinite, let alone the problem of storing an infinite amount of packaging material), it is more plausible to assume that the logic for P would be more tractable. What makes the entire business rather delicate, however, is that it is not at all clear that discourse interpretation can get by without occasionally requesting information from C. Consider, for instance, the connection argued in van der Sandt [**14**] between anaphora resolution and presupposition projection. In effect, van der Sandt [**14**] supplements quasi-conditions like $x =?$ with restrictions of the kind in Barwise and Cooper [**4**] that must somehow be accomodated within (the representation that defines) C. This poses considerable computational challenges through its unrestricted use of a logic for C (including consistency checks). As an initial step to addressing this problem (something one of us has done elsewhere), our notational system exposes it. More specifically, our proposal for item (b) amounts to a labeling system, similar to that used in Reyle [**13**], with the crucial additional idea of employing the colon : in $l : e$ as a fence separating the label l (belonging to P) from the expression e (belonging to C) which (in contrast to l) is potentially arbitrarily complex. As the expressions e in Reyle [**13**] are treated only as pieces of syntactic material, our use of : is essentially different. On the other hand, it does share with Reyle [**13**] the aim of dealing with ambiguity in a computationally effective manner — an aim that requires integrating different levels of structure.

3. Between discourse macro-structure and micro-structure

It is one thing to isolate an instance of ambiguity within a word (as in the case of lexical ambiguity) or say, a sentence (as in the case of quantifier scope ambiguity); it is quite another matter to resolve that ambiguity. The former suggests that the problem is a local one; whereas experience with the latter teaches us that the solution often involves a more global consideration of the discourse. Consider the discourse (4) below, from Asher and Lascarides [**3**], where the intervening sentence (4b) makes all the difference in disambiguating *bar*.

(4) a. The judge asked where the defendant was.
b. The barrister apologised, and said he was at the pub across the street.
c. The court bailiff found him slumped underneath the bar.

To account for (4), we must recognize that a discourse is, as argued in Grosz and Sidner [**8**], composed of *segments*, which means we must determine

(T1) what those boundaries are

and

(T2) how the segments are joined.

In SDRT, the tasks (T1) and (T2) are intertwined within a hierarchical structure of *segmented DRSs*, defined inductively from DRSs through *discourse relations* modeled after Hobbs [**9**], Polanyi [**12**] and Thompson and Mann [**15**]. Under the perspective within which we have been proceeding, taking DRSs as atomic forms of discourse segments is perhaps natural enough — assuming that these atomic forms can, in fact, be decomposed at a lower (e.g., intra-sentential) level. What may require more motivation is the choice of discourse relations that join segments so as to form larger segments. Here we will steer clear of this controversy, and concentrate (in a neutral vein) on more abstract matters.

The proposal that there is some structure to a multi-sentence discourse that must, if the discourse is to be interpreted, be uncovered (along the lines suggested by (T1) and (T2)) is essentially no different from what is by now the widely held view that some form of syntactic analysis of a sentence is a prerequisite to interpreting that sentence. The only difference is the step up in scale from a sentence to a piece of discourse, plus the possibility that the larger discourse structure may not only presuppose structure at the lower (intra-sentential) level, but, in fact, determine what that structure is. Given this two-way interaction between macro- and micro-levels, one might be tempted to do away altogether with the separation into levels. But just as we have (in the previous section) been at pains to maintain a distinction between P and

C (even though discourse interpretation may require both), let us insist on differentiating between macro- and micro-levels (noting parenthetically that the division between macro- and micro-levels is orthogonal to that between object- and meta-levels, underlying C and P, respectively). The next three sections flesh out the preceding ideas in somewhat precise terms.

4. Labeled representations and links (without holes)

To pick out a specific occurrence of an expression e, one might, as in Reyle [**13**], introduce a label l and, applying it to that occurrence, write $l : e$. In order to carry this out for an occurrence of e within a larger expression, it is useful to consider how expressions are generated. An expression, such as the formula $(\forall x)(x = y \ \vee \ f(x) < z)$ from predicate logic, is the product of *constructors*, such as $\forall$, x, $=$, y, $\vee$, f, $<$ and z, each of which is assigned an *arity* indicating the number of arguments it takes. For example, x, y and z each have arity 0 (taking, as they do, no argument), while $\forall$ has arity 2 (one for a variable, and another for a formula), as do $=$, $\vee$ and $<$ (though these require different *sorts* of arguments, which we will, for the sake of simplicity, ignore[3]). Now, to accomodate labels l within an expression e, we can simply increase the arity of each constructor by 1, using that additional argument place for a label (which can remain implicit in case that occurrence is of no interest). In UDRT and in Bos [**5**], the other ("genuine") arguments of a constructor are often left as "holes" representing scope ambiguity, the idea being that disambiguation consists, as Bos [**5**] puts it, of "plugging" these holes with labels. We will attend to holes in the next section, but keep things in this section whole in order to describe another application of labeling, for which holes need not enter (at least not immediately).

The application we have in mind is discourse macro-structure, of the kind described in the preceding section. Without spelling out what exactly discourse segments and discourse relations are, suffice it to say that a discourse relation is a constructor that takes n discourse segments as arguments. For instance, in Asher [**1**], $n = 2$ and

(i) a discourse segment is either a DRS or a *segmented DRS* (SDRS), which, like a DRS, is a pair $(\mathcal{U}, \mathcal{C})$ of sets consisting of a finite set $\mathcal{U}$ of discourse segments, and a finite set $\mathcal{C}$ of SDRS conditions on $\mathcal{U}$, where

(ii) an *SDRS condition on* $\mathcal{U}$ is an expression obtained by applying a discourse relation to discourse segments from $\mathcal{U}$.

We will exploit the parallel between an SDRS and discourse micro-structure given by a DRS (U, C) with a view towards effective disambiguation (unhindered, to the extent possible, by the problems (P1) and (P2) noted in §2). The step from (U, C) to $(\mathcal{U}, \mathcal{C})$ will be reformulated as one from (U, C) to (P, C), where P represents information packaging (that can come at either the macro-

[3]A sorted version would take an arity to be a sequence of sorts indicating the sort required at each argument place. (E.g., $\forall$ has arity (*variable,formula*).) This complicates the notation, but can otherwise be incorporated in a straightforward way to what follows.

or micro-level). Towards that end, observe that a discourse relation R holds between discourse segments $x_1, \ldots, x_n$ *relative to some context* y, which itself constitutes a discourse segment.[4] That is, a discourse relation R is more properly conceived as having arity not n, but rather $n+1$, as it would under a labeling system, where $R(x_1, \ldots, x_n)$ is transformed to $y : R(x_1, \ldots, x_n)$. To bring this observation out into the notation, it is natural to define discourse relations not on SDRSs directly, but rather on instances, or labels, of them as in Asher [**2**].[5] Hence, the elements in the set $\mathcal{U}$ forming an SDRS $(\mathcal{U}, \mathcal{C})$ can be regarded simply (and more generally) as labels.

To capture the labeling structure shared by both macro- and micro-levels of discourse structure, let us formulate the notion of labeled representation as a relational structure (in the sense of predicate logic). The idea is, given some set of constructors, to associate with a constructor f of arity n an $(n+1)$-ary relation symbol R_f, so that for a constructor e with arity 0,

$$l : e \quad \text{becomes} \quad R_e(l)$$

and a constructor f of arity 2,

$$l : f(e, e') \quad \text{becomes} \quad R_f(l_e, l_{e'}, l)$$

where l_e labels the occurrence of e in $l : f(e, e')$ and $l_{e'}$ does the same for e'. For the record,

Definition. Given a set Σ of constructors f with specified arities $n_f \in \{0, 1, \ldots\}$, a *labeled* Σ*-representation* is a pair $\langle U, I \rangle$ where

(L1) U is a non-empty set

(L2) I is an (interpretation) function with domain Σ, such that for every $f \in \Sigma$ of arity n, $I(f)$ is an $(n+1)$-ary relation on U (i.e., $I(f) \subseteq U^{n+1}$)

and (take a deep breath)

(L3) the binary relation *Succ* on U is well-founded, where (by definition) *Succ* consists of exactly the pairs $(l, l') \in U \times U$ such that for some $f \in \Sigma$ of positive arity $n > 0$, and some $(n-1)$-tuple $l_1, \ldots, l_{n-1}$ from

[4] More concretely, for instance, the sentence *John pushed him* can be said to "explain" *Max fell* within the discourse (5), analyzed in Lascarides and Asher [**11**], but (presumably) not within (6).

(5) Max fell. John pushed him.

(6) Max fell. He got up. John pushed him.

[5] In Asher [**1**] and other previous work, the contextual nature of the semantics of discourse relations was captured by the fact that each DRS generated as part of an SDRS had a distinct set of discourse referents. Accordingly, it was perfectly consistent to maintain that the relation *Explanation* held between the two DRSs in the SDRS for (5) in footnote 5 but not between the two (different) DRSs generated by the same sentences in (6). The label notation is an improvement in that it allows one to ignore details about the DRS construction procedure.

U,

$$(l, l_1, \ldots, l_{n-1}, l') \in I(f) \;\text{ or }\; (l_1, l, l_2 \ldots, l_{n-1}, l') \in I(f)$$
$$\text{or} \;\cdots\; \text{or} \;(l_1, \ldots, l_{n-1}, l, l') \in I(f)\;.$$

In item (L3), the binary relation *Succ* is what is sometimes referred to in the literature as a "successor" or immediate dominance relation, the transitive closure of which is subordination. To see why it should be well-founded, let us define the notion that an element $l \in U$ *labels* a Σ-expression e (built recursively from the constructors in Σ) by induction on e as follows.

> If e is an atomic expression — i.e., a constructor in Σ of arity 0 — then l labels e iff $l \in I(e)$. For the inductive step, l labels $f(e_1, \ldots, e_n)$ iff $(l_1, \ldots, l_n, l) \in I(f)$ for some $l_1, \ldots, l_n \in U$ such that l_1 labels e_1 and ... and l_n labels e_n.

Returning to (L3), observe that *Succ* traces the inductive generation of an expression, and must thus be well-founded, provided that

(a) for every $l \in U$, l cannot be re-used as a label within the same occurrence of an expression

and

(b) for some fixed number i, every $l \in U$ can label at most i occurrences of expressions.[6]

Part (a) is implied already by (L3), whereas (b) will follow from some finiteness assumptions that will be imposed in §6. For now, let us isolate a notion that will prove useful for relating labeled Σ-representations induced by discourse micro- and macro-structure, respectively. Given a labeled Σ-representation $\langle U, I\rangle$, and the binary relation *Succ* defined on U as in (L3) above, call an element $l \in U$ *atomic* if there is *no* $l' \in U$ for which $Succ(l', l)$. That is, $l \in U$ is atomic precisely if l is minimal relative to subordination.

Next, let us associate labeled Σ-representations with discourse micro- and macro-structure by choosing $\Sigma_\circ$ and Σ° respectively as follows.

> **Micro-structure**$_\circ$. Let $\Sigma_\circ$ be a set of constructors, yielding (as expressions) discourse referents and conditions. More precisely, the constructors in $\Sigma_\circ$ of arity 0 are discourse referents. Constructors in $\Sigma_\circ$ of arity > 0 return conditions, when applied to either conditions or discourse referents. That is, constructors of arity > 0 include not only atomic relation symbols, but also logical connectives, quantifiers, etc.[7]

[6] To the extent that a label merely signifies an occurrence, it is natural to require that a label apply to at most one occurrence (i.e., $i = 1$). Recall, however, that for applications to discourse macro-structure, a label also serves as a "context," relative to which a discourse relation joins certain discourse segments.

[7] Recall that we have been ignoring sorting of expressions to simplify the notation. There is also the complication that in DRT, the logical connectives and quantifiers operate not on conditions but on DRSs. It is easy enough, however, to get the constructors in the form mentioned above, by defining, for example, the operations on conditions φ according to the

Macro-structure$^\circ$. Constructors in Σ° of arity 0 are discourse segments; constructors of arity > 0 are discourse relations.

Beyond the fact that discourse segments have arity 0, and discourse relations have certain positive arities, the precise nature of discourse segments and discourse relations can, for the purpose of inducing a notion of labeled Σ°-representation, be left open. The force or "semantic content" of these (syntactic) constructors can be brought out elsewhere, through, for example, the following link with micro-structure. Given

a labeled $\Sigma_\circ$-representation $\langle U_\circ, I_\circ\rangle$ (for micro-structure), the atomic elements of which form the set $A_\circ$

and

a labeled Σ°-representation $\langle U^\circ, I^\circ\rangle$ (for macro-structure), the atomic elements of which form the set A°,

let $\in_U$ and $\in_C$ be two relations $\in_U \subseteq A_\circ \times A^\circ$ and $\in_C \subseteq (U_\circ - A_\circ) \times A^\circ$, supporting the intuition that for every $u \in A_\circ$ and $k \in A^\circ$,

$$u \in_U k \quad \text{iff} \quad u \text{ labels a discourse referent in the (universe of the) DRS labeled by } k$$

and for every $c \in U_\circ - A_\circ$,

$$c \in_U k \quad \text{iff} \quad c \text{ labels a condition in the DRS labeled by } k\ .$$

That is, $\in_U$ and $\in_C$ associate with every atomic element $k \in A^\circ$ (of macro-structure) the DRS (U_k, C_k), where

$$\begin{aligned} U_k &= \{f \in \Sigma_\circ \ :\ (\exists u \in A_\circ)\ u \in_U k \text{ and } u \in I_\circ(f)\} \\ C_k &= \{\varphi \ :\ (\exists c \in U_\circ)\ c \in_C k \text{ and } c \text{ labels } \varphi \\ &\qquad\qquad (\text{relative to } \langle U_\circ, I_\circ\rangle)\}\ . \end{aligned}$$

Clearly, a relation $\in_U$ can be hardwired as a relation between constructors in $\Sigma_\circ$ of arity 0 (i.e., discourse referents) and constructors in Σ° of arity 0 (i.e., discourse segments) by encoding a set of discourse referents into the notion of a discourse segment (as in the case of a DRS). On the other hand, the relation $\in_C$ can*not* be hardwired (in the same way) as a relation between constructors in $\Sigma_\circ$ of positive arity, and constructors in Σ° of arity 0, since a constructor in $\Sigma_\circ$ of positive arity does not constitute a (completed) condition, requiring, as it were, a labeled $\Sigma_\circ$-representation for completion.

Turning to the non-atomic elements of U°, an $l \in U^\circ - A^\circ$ has discourse macro-structure given by I°, and inherits the micro-structure of its atomic constituents (i.e., the atomic elements related to it by subordination). As described in §6, the discourse relations connecting discourse segments may impose certain constraints on the micro-structure of those discourse segments. Focussing for now on macro-structure, let us mention a simple "connectedness" constraint on $\langle U^\circ, I^\circ\rangle$ amounting to a very primitive notion of *coherence*;

effects of the DRT operations on the DRSs $(\emptyset, \{\varphi\})$; similarly, when discourse referents as well as conditions are involved.

namely, that there is a label $Top \in U^\circ$ such that for every $l \in U^\circ - \{Top\}$, (l, Top) is in the transitive closure of the binary relation *Succ* mentioned in (L3) above — i.e., every $l \in U^\circ - \{Top\}$ is subordinate to *Top*. Note that by (L3), there is at most one such label *Top*. The reader accustomed to drawing a DRS box, with one condition written on top of another, might flatten *Top* into such a box, where micro-structure is horizontal (in the sense that it describes a particular condition occupying a row in the box), and macro-structure is vertical (insofar as discourse relations connect conditions listed down the box).

Collecting together notions from this section, let us define a $\Sigma_\circ, \Sigma^\circ$-*packet* $\hat{P}$ to be a 5-tuple specifying a particular choice of

$$\langle U_\circ, I_\circ \rangle,\ \langle U^\circ, I^\circ \rangle,\ \in_U,\ \in_C,\ \mathit{Top} \quad .$$

We will often drop the parameters $\Sigma_\circ, \Sigma^\circ$ (when they can be ignored in the background), and speak simply of packets.

5. Discourse interpretation and underspecification

What, one might ask, have we gained by reformulating (in the previous section) the step from discourse micro-structure (U, C) to macro-structure $(\mathcal{U}, \mathcal{C})$ as one from (U, C) to $(\hat{P}, C)$? The process of discourse interpretation can, as in DRT, be conceived as an *updating of discourse contexts*, which, as every programmer knows, requires some coding or representation if the process is to be realized computationally. Now, the point behind spelling out a notion of a labeled representation is to lay down a solid basis for integrating SDRT and UDRT. But to actually carry out such an integration, we will need to bore holes into our labeled representations.

Let us explain. The notion of a *discourse context* given in Asher [**2**] specifies a particular packet $\hat{P}$, as well as (i) an element $Current \in U^\circ$ that labels the last discourse segment analyzed, and (ii) a certain function (we will get to in the next section) assigning discourse segments *main eventualities*. But confining our attention at present to packets $\hat{P}$, a crucial question to ponder is

(Q) should the initial discourse context in a discourse interpretation process, in fact, *single out a particular* packet $\hat{P}$, and, even if it did, should the output of a discourse interpretation process also specify exactly one packet?

Notice that we have smuggled into (Q) the presupposition of a definite output. Both the uniqueness and existence of an output might very well be challenged, under an approach to ambiguity based on non-determinism. Be that as it may, the widespread interest in underspecification suggests that the answer "no" to both parts of (Q) is worth exploring.

The idea behind underspecified representations is simply to leave things out. We have gone through great lengths to formulate labeled $\Sigma_\circ$-representations as relational structures so that, for instance, the first-order language of these relational structures can be used to describe families of labeled $\Sigma_\circ$-representations. Writing R_f for the $(n+1)$-ary relation symbol induced by a constructor f in $\Sigma_\circ$ of arity n, and adding a binary relation symbol $\leq$ for the reflexive-transitive closure of the relation *Succ* mentioned in (L3), the generalization of UDRT in Bos [**5**] is obtained by considering atomic formulas of the form $R_f(x_1, \ldots, x_n, l)$ and $x \leq y$, where holes are nothing more than variables (at this meta-level of packaging). Other forms of ambiguity (e.g., lexical) can be captured if the constructors can be replaced by variables, as in higher-order logic. SDRT invites similar investigations of labeled Σ°-representations, separate from or in conjunction with labeled $\Sigma_\circ$-representations via fixed or varying choices of links $\in_U$ and $\in_C$. Doubts (such as that expressed in Wilson [**16**]) that discourse relations of the sort used in SDRT must be identified in order to interpret a discourse suggest that perhaps R_f ought, at least in some cases, to be left underspecified (i.e., as a hole). An essential consideration that must be brought into the picture is, as Wilson [**16**] argues, "processing effort." It is precisely that issue which the next section tries to address, in formal terms.

6. Finitely packaging infinite content

Having been concerned, in §§4 and 5, almost exclusively with information packaging, let us now peek at the content inside. But as we will be interested in just how much effort that takes, let us define a labeled Σ-representation $\langle U, I\rangle$ to be a *finitely labeled* Σ*-representation* if, in addition to (L1), (L2), and (L3), it meets the following requirement (reflecting the finiteness of discourse)

(L4) U is finite, and $I(f) \neq \emptyset$ for at most finitely many f's in Σ.

Henceforth, all labeled Σ-representations (in particular, those in packets $\hat{P}$) will be assumed to be finitely labeled. It follows that a packet is a perfectly finite object that can readily be Gödel-numbered, and fed to a computer.

Writing Φ for the collection of conditions, let us assume also that we are given some notion $\vdash\ \subseteq \mathsf{Pow}(\Phi) \times \Phi$ of logical entailment on these conditions (by virtue of which the conditions can be said to have *content*[8]). Now, fix a packet $\hat{P}$. Recall from §4 that every atomic element k of the macro-structure $\langle U^\circ, I^\circ\rangle$ is assigned a set $C_k \subseteq \Phi$ of conditions. The key to the content encoded by a packet $\hat{P}$ is the relation $:_\vdash$ between such elements k and conditions φ, defined by

$$k :_\vdash \varphi \quad \text{iff} \quad C_k \vdash \varphi \ .$$

The relation $:_\vdash$ can be extended to arbitrary (non-atomic) elements in the macro-structure $\langle U^\circ, I^\circ\rangle$, including *Top*, for which

$$Top :_\vdash \varphi \quad \text{iff} \quad (\bigcup_{k\in A^\circ} C_k) \vdash \varphi \ .$$

[8] "Content" here can be understood in the model-theoretic terms mentioned in §2, assuming $\vdash$ is connected to that by a completeness theorem, as in first-order logic.

The computational costs of $:_\vdash$ can be traced to $\vdash$, for which the best we can hope for is that it be recursively enumerable (as in the case of first-order logic). It is for this reason that we have described the relation $:$ earlier as a "fence" for keeping discourse interpretation computationally effective.

As we have also mentioned, however, the pressure to jump the fence can be irresistible when theorizing about discourse without regard for computation. To determine a range of temporal relations, Lascarides and Asher [**11**] make do with the trivial "inspection" notion $:_I$ of inference relation, given by

$$k :_I \varphi \quad \text{iff} \quad \varphi \in C_k \ ,$$

together with a function Ev picking out the *main eventuality* $Ev(k)$ associated with the label k. The idea, in general, is to enrich a packet with some finite structure $\mathcal{E}$ around the image of Ev, approximating the worlds picked out by $\{\varphi \in \Phi \mid Top :_\vdash \varphi\}$ (according to the model-theoretic interpretation of conditions mentioned in §2 above). An additional feature of SDRT is the use of a non-monotonic logic to determine discourse relations between discourse segments, and (for instance) the temporal relations between the eventualities. The non-monotonicity here reflects a human tendency to disambiguate discourse without necessarily incontrovertible evidence. But the computational cost of nonmonotonically inferring the appropriate SDRS is high even in the best of cases, involving only $:_I$ and quantification over the aforementioned finite structure, and hopeless in the more realistic case, where the use of $:_\vdash$ is less restricted.

In addition to making nonmonotonic leaps in discourse interpretation, however, humans also live with ambiguity. Indeed, part of the interest in underspecification is the possibility of filling in holes monotonically as the discourse progresses. To what degreee can the non-monotonicity in SDRT be reduced through underspecification? Without attempting to settle this question now once and for all, let us try to be more precise about what we mean by a discourse context. An *idealized discourse context* Ω is a tuple consisting of a (finite) packet $\hat{P}$, a label *Current* $\in U^\circ$ (pointing to the last discourse segment analyzed), a function Ev with domain U°, and possibly some additional finite structure built around the image of Ev.

To turn an idealized discourse context Ω into a *discourse context*, two ingredients are needed:

(a) a finite set $U_?$ of *pieces* of Ω (to be treated as holes)

and

(b) a set Γ of *constraints* on Ω (specifying requirements on how to fill the holes in $U_?$).

By "pieces" of Ω, we mean not only

(I) labels $l \in U_\circ \cup U^\circ$ in the packet specified by Ω

but also

(II) constructors $f \in \Sigma_\circ \cup \Sigma^\circ$ (buried) in the domain of $I_\circ$ and of I°.

Confining our attention first to micro-structure, scope ambiguities represented in UDRT and in Bos [**5**] can be captured by

(a) including among $U_?$ labels from $U_\circ$ (as allowed under (I)), and
(b) encoding into Γ the requirement that such holes be filled (in a bijective manner, if we wish) from some set of pieces of Ω disjoint from $U_?$, in accordance with certain subordination constraints (expressed through an auxiliary symbol $\leq$).[9]

In addition, cases of lexical and pronominal ambiguity can be represented by

(a) throwing into $U_?$ constructors from $\Sigma_\circ$ (as allowed under (II)),

although for these cases,

(b) it is not at all clear that Γ should require that such holes be plugged with material from Ω.

That is to say, a crucial difference between the structural ambiguities represented in UDRT and the lexical and pronominal ambiguities disambiguated in SDRT is that in the latter, specifying the material with which to plug the holes can be highly non-trivial (as illustrated by bridging phenomena). In particular, it is doubtful whether Γ should require that those holes be filled by pieces in Ω.

At the discourse macro-level, pieces in item (I) specify *where* to attach a discourse segment, while pieces in item (II) specify *what* discourse relation to link discourse segments with. In SDRT, Γ would constrain all labels in $U^\circ - \{Top\}$ to be subordinate to Top, and relate the main eventualities of discourse segments linked by particular discourse relations according to the rules set out, for instance, in Lascarides and Asher [**11**]. Now, by including into $U_?$ discourse macro-items from (I) and (II), we allow for the possibility of underspecifying the structure of discourse, according to considerations centering around "processing effort" (Wilson [**16**]). That is, suppose we have a discourse segment K (possibly underspecified) and a label l that is to be incorporated into some (initial) discourse context. One option — a "lazy" and trivially monotonic one — is to simply add the label l to $U_?$ and any other underspecified elements in K to $U_?$. Another option is to resolve some of the underspecification by using various bits of the SDRS update procedure of Asher [**1, 2**]. Which bits we want to resolve will depend, as relevance theorists have pointed out, on how much computational cost we think is associated with the resolution and the amount of information the interpreter hopes to gain from the computation. For instance, the presence of a clue

[9] The reader familiar with the *underspecified representations* $\langle H, L, C\rangle$ of Bos [**5**] should be able to extract such triples as special instances of the notion of a discourse context above, noting the formal (but otherwise inessential) difference that holes are treated (for the sake of convenience) as labels above.

By introducing discourse structure into the representations, traces of syntax of the kind left, for example, in the notion of a plugging (present in the revised semantics of URs in Definition 8 of Bos [**5**]) can be analyzed semantically, yielding a fully compositional interpretation of ambiguous expressions, relative to some (possibly partial) notions of disambiguation (Fernando [**6**]).

word or conditions in the look-up may make it almost effortless to compute the discourse relation through which to attach the new label l. On the other hand, although it may be quite expensive to compute the discourse relation, the resolution of other underspecifications — both lexical and pronominal — may depend on this computation (see Asher and Lascarides [**3**]). In this case too, it might seem appropriate for the interpreter to calculate the discourse relation. If, however, the other underspecified elements do not rely on a choice of discourse relation, the interpreter might leave the choice of R unresolved (if the computation of R is itself not trivial). Similarly, computing an attachment site might, in some cases, be very easy (especially if there is only one possible site), but, in other cases, quite expensive.

7. Conclusion

There is a striking logical contrast in disambiguation between the adventurous non-monotonicity in SDRT and the cautious stillness of UDRT. This is, no doubt, due in no small measure to the very different directions in which SDRT and UDRT have been pushed — SDRT towards the semantics/pragmatics interface, and UDRT towards computational work. We have tried above to outline a level of representations at which such a connection might be established. Marrying SDRT with UDRT has dramatic consequences for discourse interpretation. By allowing underspecification at all levels, we can reduce the requirements on contextual update, disambiguating only when we think the gains are worth the computational effort — a point we hope to take up elsewhere.

References

[1] Nicholas Asher. *Reference to Abstract Objects in Discourse*. Kluwer, Dordrecht, 1993.

[2] Nicholas Asher. Mathematical treatments of discourse contexts. In P. Dekker and M. Stokhof, editors, *Proc. Tenth Amsterdam Colloquium*. ILLC, University of Amsterdam, 1996.

[3] N. Asher and A. Lascarides. Lexical disambiguation in a discourse context. *J. Semantics*, 12, 1995.

[4] J. Barwise and R. Cooper. Extended Kamp notation: a graphical notation for situation theory. In P. Aczel, et al., editors, *Situation Theory and Its Applications*, volume 3. CSLI, Stanford, 1993.

[5] Johan Bos. Predicate logic unplugged. In P. Dekker and M. Stokhof, editors, *Proc. Tenth Amsterdam Colloquium*. ILLC, University of Amsterdam, 1996.

[6] Tim Fernando. Ambiguity under changing contexts. Manuscript, 1996. Preliminary version presented in *Mathematics of Language 4* (Philadelphia, October 1995), and submitted for publication.

[7] J. Groenendijk and M. Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14, 1991.

[8] B. Grosz and C. Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3), 1986.

[9] Jerry R. Hobbs. On the coherence and structure of discourse. Technical Report 85-37, Center for the Study of Language and Information, Stanford, 1985.

[10] H. Kamp and U. Reyle. *From Discourse to Logic*. Kluwer, Dordrecht, 1993.

[11] A. Lascarides and N. Asher. Temporal interpretation, discourse relations and commonsense entailment. *Linguistics and Philosophy*, 16(5), 1993.

[12] Livia Polanyi. A theory of discourse structure and discourse coherence. In P.D. Kroeberger, et al., editors, *General Session, 21st Regional Meeting of the Chicago Linguistics Society*. Chicago, 1985.

[13] Uwe Reyle. Dealing with ambiguities by underspecification: construction, representation and deduction. *J. Semantics*, 10(2), 1993.

[14] Rob A. van der Sandt. Presupposition projection as anaphora resolution. *J. Semantics*, 9(4), 1992.

[15] S. Thompson and W. Mann. Rhetorical structure theory: a framework for the analysis of texts. In *IPRA Papers in Pragmatics*, volume 1. 1987.

[16] Deirdre Wilson. Truth, coherence, relevance: recent approaches to pragmatics. Invited talk, ITALLC 96, London, July 1996.