Reasoning Dynamically about What One Says

Nicholas Asher
IRIT
Université Paul Sabatier, Toulouse
asher@irit.fr

Alex Lascarides
School of Informatics,
University of Edinburgh
alex@inf.ed.ac.uk

Abstract

In this paper we make SDRT's glue logic for computing logical form dynamic. This allows a dialogue agent to anticipate what the update of the semantic representation of the dialogue would be after his next contribution, including the effects of the rhetorical moves that he is contemplating performing next. This is a pre-requisite for planning what to say next. We make the glue logic dynamic by extending a dynamic public announcement logic (PAL) with the capacity to perform default reasoning—an essential component of inferring the pragmatic effects of one's dialogue moves. We add to the PAL language a new type of announcement, known as *ceteris paribus* announcement, and this is used to model how an agent anticipates the default consequences of his next dialogue move. Our extended PAL validates more intuitively compelling patterns of default inference than existing PALs for practical reaosning, and we demonstrate via the proof of reduction axioms that the dynamic glue logic, like its static version, remains decidable.

1 Introduction

Speakers in dialogue anticipate their interlocutors' interpretations and adjust their utterances accordingly. Researchers adopt planning (e.g., Stone (1998)), decision theory (e.g., Williams and Young (2007)) or game theory (e.g., van Rooij (2001)) to model such decisions. But these approaches tend to use models of semantics that don't capture constraints on interpretation stemming from logical structure (Kamp and Reyle, 1993) or discourse coherence (Hobbs et al., 1993, Asher and Lascarides, 2003). On the other hand, semantic models that do capture such constraints either do not interpret the rules for constructing logical form (Poesio and Traum, 1998), or do so in a *static* logic (e.g., weighted abduction (Hobbs et al., 1993), or nonmonotonic deduction (Asher and Lascarides, 2003)). Since dynamic discourse update *uses* the static axiomatisation but is *not* a part of it, a speaker cannot compare his candidate next moves, inferring how they will be interpreted—but doing this is essential for axiomatising decisions about what to say.

This paper aims to fill this gap. We start with Segmented Discourse Representation Theory (SDRT, Asher and Lascarides (2003)), a model of discourse where interpretation depends on logical and rhetorical structure. We will make SDRT's existing, static *glue logic* for constructing logical forms dynamic, incorporating discourse update into the axiomatisation. We thus achieve a pre-requisite for making strategic decisions during conversation—the speaker can reason about an interlocutor's interpretation—but we leave to future work the task of interfacing these expected outcomes with the speaker's preferences or goals.

Section 2 describes the logical form of dialogue in SDRT and Section 3 presents its existing *glue logic* and the accompanying dynamic discourse update. Section 4 replaces the static glue logic with a dynamic one that incorporates discourse update into the axiomatisation: we extend a dynamic logic of public announcement with the capability to perform default reasoning, a necessary feature of dialogue interpretation. We prove that the dynamic glue logic has the same computational complexity as the static one; so constructing logical form remains computable.

2 Logical Forms for Dialogue

A fundamental decision that a speaker must make about his next move is its effects on agreement. Lascarides and Asher (2009) argue that *rhetorical relations* (e.g., *Narration*, *Explanation*) are crucial for capturing *implicit agreement*: representing the illocutionary contribution of an agent's utterance via rhetorical relations reflects his commitments to another agent's commitments, even when this is linguistically implicit. For example, Karen's utterance (1c), taken from (Sacks et al., 1974, p717), commits her to

(1b) thanks to the semantic consequences of the relational speech act *Explanation*$(1b, 1c)$ that she has performed:

(1) a. *Mark (to Karen and Sharon):*
Karen 'n' I're having a fight,
b. after she went out with Keith and not me.
c. *Karen (to Mark and Sharon):*
Wul Mark, you never asked me out.

Arguably, by committing to (1b) Karen also commits its *illocutionary effects*—(1b) explains (1a). These commitments are not *monotonically* entailed by (1c)'s compositional semantics nor by Karen's asserting it. Rather, Karen's implicit acceptance of Mark's contribution is logically dependent on the relational speech acts they perform and their semantics.

More generally, Lascarides and Asher (2009) propose that the commitments of each agent at a given dialogue turn (where a turn boundary occurs whenever the speaker changes) is a Segmented Discourse Representation Structure (SDRS, Asher and Lascarides (2003)) as shown in each cell of Table 1, the proposed logical form for (1). Roughly, an SDRS is a rooted hierarchical set of discourse segments or *labels*, with each label π associated with some content ϕ (written $\pi{:}\phi$). The contents ϕ are expressed in a language $\mathcal{L}$ of SDRS-formulae, and the hierarchical structure occurs because $R(\pi, \pi') \in \mathcal{L}$ where R is a rhetorical relation—in other words, the content of a segment can feature rhetorical connections among sub-segments. For simplicity, we have omitted from Table 1 the contents of the clauses (1a) to (1c), corresponding to labels π_1 to π_3, and adopted a convention that the root label of agent a's SDRS for turn j is π_{ja}. We may also refer to the content of label π as K_π, and to the SDRS that agent a commits to in turn j as $T^a(j)$.

The logical form of dialogue (e.g., Table 1) is called a Dialogue SDRS (DSDRS). The agents' SDRSs can share labels, and each label is always associated with the same content. Thus an agent can commit to the the content expressed by prior speech acts, even if they were performed by another agent. For example, *Explanation*(π_1, π_2) is a part of Karen's SDRS, making her committed to her and Mark having a fight because she went out with Keith.

Formally, Lascarides and Asher (2009) define a dynamic semantics $\models_d$ for DSDRSs in terms of that for SDRSs (Asher and Lascarides, 2003), where $\models_d$ captures shared commitments. With agreement being shared public commitment, the logical form of (1) makes the following agreed upon: Mark and Karen were having a fight because she went out with Keith and not Mark.

Lascarides and Asher (2009) propose a number of default axioms for constructing these logical forms: they predict the semantic scope of implicit and explicit endorsements and challenges, and provide the basis for adding *Explanation*(π_1, π_2) to Karen's SDRS in (1). In general, the principles are designed to maximise one's ongoing commitments from prior turns, subject to them being consistent with default inferences about the illocutionary contribution one intends to make in the current turn.

This principle therefore predicts that Karen makes different commitments if (1a) is replaced with (1a′) (and the word *after* is removed from (1b)):

(1) a′. Karen is a bitch.

In this case, as before, Karen is committed to *Explanation*(b, c) and therefore is committed to the content of (1b). However, this time, the inference that (1c) explains (1b) supports a further default inference about the speech act that Karen performed in uttering (1c): namely, the *Explanation* segment provides *Counterevidence* to the content of (1a′). This captures the intuition that Karen uses (1c) to *justify* her choices as conveyed in (1b), making those choices reasonable rather than vindictive and thereby undermining (1b) as an explanation of (1a′)—the content that Mark committed to. Thus, since default inferences about the illocutionary contribution of (1c) makes Karen committed, via *Counterevidence*, to the denial of (1a′), and this is inconsistent with a commitment to *Explanation*(a', b), the logic for constructing logical form does not add Mark's commitments from the first turn to Karen's commitments for the second turn, even though Karen is committed to a part of what Mark committed to: namely, (1b).

3 SDRT's Glue Logic

Asher and Lascarides (2003) argue that constructing logical form should be decidable, so as to provide a competence model of language users who largely agree on what was said if not its cognitive effects (Lewis, 1969). This *glue logic* must involve nonmonotonic reasoning (see, for instance, the above discussion of (1)), and hence consistency tests: agents have only partial information about the context, including the speech acts that they intended to perform or commit to. So to remain decidable, the glue logic must be separate from but related to the logic $\models_d$ for *interpreting* logical form. That is, computing what was said doesn't require evaluating whether what was

Turn	Mark's SDRS	Karen's SDRS
1	π_{1M} : $Explanation(\pi_1, \pi_2)$	$\emptyset$
2	π_{1M} : $Explanation(\pi_1, \pi_2)$	π_{2K} : $Explanation(\pi_1, \pi_2) \wedge Explanation(\pi_2, \pi_3)$

Table 1: A representation of dialogue (1).

said is true. SDRT achieves this separation through underspecified semantics (e.g., Egg et al. (2001)).

An underspecified logical form (ULF) is a *partial description* of the *form* of the intended logical form, here a DSDRS. SDRT's glue logic builds ULFs, and so its language $\mathcal{L}_{ulf}$ describes DSDRSs: each n-ary constructor in $\mathcal{L}$ corresponds in $\mathcal{L}_{ulf}$ to an $n{+}1$-ary predicate symbol that takes labels as arguments. Labels denote scopal positions in the DSDRS being described; so the first n arguments to the predicate P in $\mathcal{L}_{ulf}$—where P corresponds to the n-place constructor P in $\mathcal{L}$—denotes the scopal positions of the arguments to P in the described DSDRS(s), and the $(n+1)^{th}$ label is the scopal position of P itself. This label is written to the left of a colon with its predicate and other arguments to the right. So l:$\mathsf{dog}(l') \wedge l'$:$\mathsf{d}$ in $\mathcal{L}_{ulf}$ describes $dog(d)$ in $\mathcal{L}$. $\mathcal{L}_{ulf}$ also includes variables like ?, which indicate that the value of a constructor in the described DSDRS(s) is unknown. For instance, the compositional semantics of a pronoun introduces into the ULF the formula l_0:$=(l_1, l_2) \wedge l_1$:$\mathsf{x} \wedge l_2$:?—an equality between x and some unknown variable. The language $\mathcal{L}_{ulf}$ also features Boolean connectives like $\wedge$, and a weak conditional $>$ that's used to formalise defaults ($A > B$ means "If A then normally B").

The glue logic derives a logical form (or, in fact, a ULF) via schemata that predict pragmatically preferred values for underspecified semantic elements. They in particular predict rhetorical connections, and have the form given in (2a):

$$(\lambda{:}?(\alpha, \beta) \wedge \mathsf{Info}(\alpha, \beta, \lambda)) > \lambda{:}\mathsf{R}(\alpha, \beta, \lambda) \quad (2a)$$

$$(\lambda{:}?(\alpha, \beta) \wedge \alpha{:}\mathsf{int}) > \lambda{:}\mathsf{IQAP}(\alpha, \beta) \quad (2b)$$

In words, if segment β is rhetorically connected to α as part of a segment λ but the relation is unknown, and moreover $\mathsf{Info}(\alpha, \beta, \lambda)$ holds of the content labelled by λ, α and β, then normally the rhetorical relation is R. The conjunct $\mathsf{Info}(\alpha, \beta, \lambda)$ is proxy for particular ULF formulae and the axioms are justified on the basis of linguistic knowledge, world knowledge, or knowledge of cognitive states. Rule (2b) is an example from Asher and Lascarides (2003) where int stands for interrogative mood; so (2b) states that a response to a question is normally an indirect answer (*IQAP* stands for Indirect Question Answer Pair).

Definition 1 gives the model theory for the glue language $\mathcal{L}_{ulf}$: intuitively, $\mathcal{M}, s \models_g \phi$ means that ϕ is a (perhaps partial) description of the DSDRS s. As we've said, not all $\models_d$-consequences from the dynamic semantics of DSDRSs are transferred into the glue logic: those arising from substitution of equalities, $\wedge$- and $\vee$-elimination and $\exists$-introduction are validated by $\vdash_g$ but $\exists$-elimination is not and so $\vdash_g$ loses the logical equivalence between the SDRS-formulae $\neg\exists x \neg\phi$ and $\forall x \phi$. This model theory is static; in Section 4 we make it dynamic.

Definition 1 Static Glue Model Theory

A model $\mathcal{M} = \langle S, *, V \rangle$ for $\mathcal{L}_{ulf}$ consists of:

- *A a set of states S where each $s \in S$ is a unique DSDRS,*[1]
- *A function $*$ from a state and a set of states to a set of states (for interpreting $>$), and*
- *A function V for interpreting $\mathcal{L}_{ulf}$'s non-logical constants (so V is constrained by the partial transfer of $\models_d$-entailments from $\mathcal{L}$ described above).*

Then the truth definitions are:

- $\mathcal{M}, s \models_g \phi$ *iff* $s \in V(\phi)$ *for atomic* ϕ
  $\mathcal{M}, s \models_g \phi \wedge \psi$ *iff* $\mathcal{M}, s \models_g \phi$ *and* $\mathcal{M}, s \models_g \psi$
  $\mathcal{M}, s \models_g \neg\phi$ *iff* $\mathcal{M}, s \not\models_g \phi$
  $\mathcal{M}, s \models_g \phi > \psi$ *iff* $*^{\mathcal{M}}(s, [\![\phi]\!]^{\mathcal{M}}) \subseteq [\![\psi]\!]^{\mathcal{M}}$,
  where $[\![\phi]\!]^{\mathcal{M}} = \{s' : \mathcal{M}, s' \models_g \phi\}$

The glue logic also has monotonic and nonmonotonic relations $\vdash_g$ and $\mid\!\sim_g$ (Asher and Morreau, 1991, Asher, 1995), with $\mid\!\sim_g$ yielding default inferences, via axioms like (2b), about discourse interpretation, including particular resolutions of anaphora. $\vdash_g$ abides by axioms of classical logic plus axioms and rules on $>$-formulae such as those in (3) (corresponding con-

[1] In fact, it is a *finite constructor tree* (Egg et al., 2001), which is a function from a tree domain (i.e., a subset of $\mathbb{N}*$ which is closed under prefix and sibling) to constructors in $\mathcal{L}$. Each finite constructor tree thus corresponds to a unique DSDRS.

straints on $*$ are omitted here):

$$\frac{\vdash_g A \rightarrow B}{\vdash_g A > B} \tag{3a}$$

$$\frac{\vdash_g A \rightarrow B}{\vdash_g A > B} \tag{3b}$$

$$\frac{\vdash_g B \rightarrow C}{A > B \vdash_g A > C} \tag{3c}$$

$$\frac{\vdash_g A \rightarrow B}{A > C, B > \neg C \vdash_g B > \neg A} \tag{3d}$$

$\mid\!\sim_g$-consequences are computed via $\vdash_g$ by converting $>$-formulae into $\rightarrow$-ones proviso the result being $\vdash_g$-consistent. $\mid\!\sim_g$ validates many intuitively compelling inferences such as those below, and the logic is sound, complete and decidable.

Defeasible Modus Ponens: $\phi, \phi > \psi \mid\!\sim \psi$

Penguin Principle: If $\phi \vdash_g \psi$ then $\phi, \phi > \chi, \psi > \neg\chi \mid\!\sim_g \chi$

Nixon Diamond: If $\phi \not\vdash_g \psi$ and $\psi \not\vdash_g \phi$ then $\phi, \psi, \phi > \chi, \psi > \neg\chi \not\mid\!\sim_g \chi$ (and $\not\mid\!\sim_g \neg\chi$)

Weak Deduction: if (a) $\Gamma, \phi \mid\!\sim_g \psi$, (b) $\Gamma \not\mid\!\sim_g \psi$ and (c) $\Gamma \not\mid\!\sim_g \neg(\phi > \psi)$, then (d) $\Gamma \mid\!\sim_g (\phi > \psi)$

The Penguin Principle is valid because (3d) is.

Definition 2 makes the updated ULF include all the $\mid\!\sim_g$-consequences of the old and the new information (see **Simple Update**). So update always *adds* constraints to what the dialogue means. If there is more than one choice of labels that the new content attaches to, then update is conservative and generalises over all the possibilities (see **Discourse Update**).

Definition 2 Discourse Update for DSDRSs

Simple Update *of a context with new content* β, *given a particular attachment site* α.

Let $T(d, m, \lambda) \in \mathcal{L}_{ulf}$ *mean that the label* λ *is a part of the* SDRS $T^d(m)$ *in the* DSDRS *being described. So the* ULF*-formula* $\lambda{:}?(\alpha, \beta) \wedge T(d, m, \lambda)$ *specifies that the new information* β *attaches to the* DSDRS *as a part of the* SDRS $T^d(m)$. *Let* σ *be a set of (fully-specified)* DSDRS*s, and let* $Th(\sigma)$ *be the set of all* ULF*s that partially describe the* DSDRS*s in* σ. *Let* ψ *be either (a) a* ULF $\mathcal{K}_\beta$, *or (b) a formula* $\lambda{:}?(\alpha, \beta) \wedge T(d, m, \lambda)$, *where* $Th(\sigma) \vdash_g \mathcal{K}_\beta$. *Then:*

$$\sigma + \psi = \{\tau : \text{ if } Th(\sigma), \psi \mid\!\sim_g \phi \text{ then } \tau \vdash_g \phi\}, \text{ provided this is not } \emptyset;$$
$$\sigma + \psi = \sigma \text{ otherwise}$$

Discourse Update. *Suppose that* A *is the set of available attachment points in the old information* σ. $update_{\text{SDRT}}(\sigma, \mathcal{K}_\beta)$ *is the* union *of* DSDRS*s that results from a sequence of* $+$*-operations for each member of the power set* $\mathcal{P}(A)$ *together with a stipulation that the last element of the updated* DSDRS *is* β.

The power set $\mathcal{P}(A)$ represents all possible choices for what labels in σ the new label β is attached to, so $update_{\text{SDRT}}$ is neutral about which member of $\mathcal{P}(A)$ is the 'right' choice.

Discourse update typically doesn't yield a specific enough ULF to identify a *unique* logical form or DSDRS. But intuitively, some DSDRSs that satisfy the $\mid\!\sim_g$-consequences are 'preferred' because they are *more coherent*. SDRT makes degree of coherence influence interpretation by *ranking* the DSDRSs in the update into a partial order. This partial order adheres to some very conservative assumptions about what contributes to coherence, as stated in Definition 3 from Asher and Lascarides (2003).

Definition 3 Maximise Discourse Coherence (MDC) *Discourse is interpreted so as to maximise discourse coherence, where the (partial) ranking* $\succ$ *among interpretations adheres to the following:*

1. *All else being equal, if* DSDRS ϕ *has more rhetorical connections between two labels than* DSDRS ψ, *then* $\phi \succ \psi$.

2. *All else being equal,* $\phi \succ \psi$ *if* ϕ *features more semantic values that support* $\mid\!\sim_g$*-inferences for particular rhetorical relations.*

3. *Some rhetorical relations are inherently scalar. For example, the quality of a* Narration *is dependent on the specificity of its common topic. All else being equal,* $\phi \succ \psi$ *if* ϕ *features higher quality rhetorical relations.*

4. *All else being equal,* $\phi \succ \psi$ *if* ϕ *has fewer labels but no semantic anomalies: e.g.,* $\pi_0{:}Contrast(\pi_1, \pi_2) \wedge Condition(\pi_2, \pi_3)$ *is anomalous because the first speech act 'asserts'* K_{π_2} *and the second doesn't, but* $\pi_0{:}Contrast(\pi_1, \pi), \pi{:}Condition(\pi_2, \pi_3)$ *isn't anomalous.*

4 Dynamic Commitments in SDRT

Definition 2 uses the entailment relation $\mid\!\sim_g$ but it is external to it; so is MDC. It is impossible to reason about dynamic updates *within* a static glue logic and so we need to make it dynamic so as to support strategic decisions about what to say. We do this with a public announcement logic (PAL) (Baltag et al., 1999). A PAL features the action of announcing a formula, which changes the model by restricting the states in the output model to those in which the announced formula is true. SDRT's glue logic can thus be recast in terms of the effects of announcing a formula: the states of the model are still DSDRSs (see Definition 1) with announcements eliminating DSDRSs from the input model that fail to satisfy the announcement.

As Definition 2 suggests, we need to specify the effects of three sorts of announcements:

1. $\mathcal{K}_\beta$—the ULF of an utterance or segment.
2. $\lambda{:}?(\alpha, \beta) \wedge T(d, j, \lambda)$—a choice of where to attach a new segment.
3. $last = \beta$—β is the last entered element.

If all consequences of one's announcements were monotonic, then simple PAL would do. But Section 3 makes plain that *nonmonotonic* consequences of announcements determine the DSDRSs, since discourse interpretation is generally a product of commonsense reasoning.

Extensions to PAL that support nonmonotonic reasoning exist. For instance, van Benthem (2007) and Baltag and Smets (2006) propose dynamic PALs for modelling belief revision; they incorporate into a standard PAL conditional doxastic models, with logics equivalent to AGM belief revision theory (Alchourrón et al. , 1985). Like their PALs, ours is extended by introducing a weak conditional connective. However, our logic differs from theirs in that our extension to PAL, being based on the connective from Commonsense Entailment (Asher and Morreau, 1991, Asher, 1995), validates the monotonic axiom (3d) and hence also validates the Penguin Principle. The Penguin Principle incorporates an important and intuitively compelling principle of nonmonotonic inference that we have shown extensively elsewhere is vital for accurately predicting the logical form of coherent discourse (Lascarides and Asher , 1993, Asher and Lascarides, 2003). So it is essential that our dynamic PAL version of the glue logic continue to support this type of inference.

Our strategy, then, is to introduce another sort of announcement—not simple announcement but *announcement ceteris paribus* or ACP—and we will define ACPs in terms of the conditional $>$, so that ACPs support similar inference patterns to those supported in the static version of the Glue Logic.

We will convert the static model theory from Definition 1 into a dynamic one for interpreting ACPs. This involves (a) extending the language $\mathcal{L}_{ulf}$ to express announcements; and (b) defining how models are transformed by such announcements in interpretation. As is standard in PAL, we add a modality $[!\phi]$ to $\mathcal{L}_{ulf}$, to express the announcement that ϕ. The formula $[!\phi]\psi$ means that ψ follows from announcing ϕ. The above three values of ϕ are all $>$-free (although $\mathcal{K}_\beta$ may contain a predicate symbol corresponding to the distinct constructor $>$ in $\mathcal{L}$). So we make announcements $>$-free. We extend standard PAL by introducing a new modality $[!\phi]^{cp}$ for ACPs, where $[!\phi]^{cp}\psi$ means that ψ normally follows from announcing ϕ.

Definition 4 assigns this extended language $\mathcal{L}_{ulf}$ a dynamic model theory, with announcements transforming models. Observe that $[!\phi]^{cp}\psi$ is defined in terms of the conditional connective $>$.

Definition 4 Dynamic Glue Model Theory

Let $\mathcal{M} = \langle S, *, V \rangle$ be a model as in Definition 1. We define $\mathcal{M}^\phi$ in the standard way and $\mathcal{M}^{cp(\phi)}$ using the nonmonotonic closure of ϕ given the background truths of glue logic GL that characterise SDRT*'s constraints on attachment of new constituents in a given* SDRS *and on what relations can be inferred between two given points of attachment (e.g., see axiom (2b)). We assume that this background theory holds in all models and, crucially, is finite. For such finite theories, there exists a* prime implicate *or strongest finite formula that is a nonmonotonic consequence of the theory and from which all other nonmonotonic consequences follow (Asher, 1995). The prime implicate of the background theory together with ϕ, which we shall write as $\Im_\phi$, characterises the nonmonotonic closure of ϕ. In other words:*

$$\forall\psi \text{ such that } \phi\mid\!\sim\psi, \Im_\phi \rightarrow \psi$$

Because $\mid\!\sim$ is supra-classical (see (3a)), we have as an axiom $\Im_\phi \rightarrow \phi$. Furthermore, $\Im_\phi$ incorporates the consequences of axiom (3d) and it abides by the Penguin Principle: in other words, it entails the nonmonotonic consequences of more specific information that follows from $\Im_\phi$ when it conflicts with the nonmonotonic consequences of less specific information in $\Im_\phi$.

We now define:

$$\begin{aligned} \mathcal{M}^\phi = &\ \langle S^\phi,\ *^{\mathcal{M}}|S^\phi, V\rangle, \text{ where} \\ & S^\phi = S^{\mathcal{M}} \cap [\![\phi]\!]^{\mathcal{M}} \\ \mathcal{M}^{cp(\phi)} = &\ \langle S^{cp(\phi)}, *^{\mathcal{M}}|S^{cp(\phi)}, V\rangle, \text{ where} \\ & S^{cp(\phi)} = S^{\mathcal{M}} \cap [\![\Im_\phi]\!]^{\mathcal{M}} \\ \Im_\phi \rightarrow \psi, &\quad \forall\psi \text{ such that } \phi \mid\!\sim \psi \end{aligned}$$

The dynamic interpretations of $[!\phi]\psi$ and $[!\phi]^{cp}\psi$ are:

$$\begin{aligned} \mathcal{M}, s \models [!\phi]\psi &\quad \text{iff} \quad \mathcal{M}, s \models \phi \rightarrow \mathcal{M}^\phi, s \models \psi \\ \mathcal{M}, s \models [!\phi]^{cp}\psi &\quad \text{iff} \quad \mathcal{M}, s \models \Im_\phi \rightarrow \mathcal{M}^{cp(\phi)}, s \models \psi \end{aligned}$$

In words, model $\mathcal{M}^{\phi}$ is formed from $\mathcal{M}$ by eliminating all states that don't satisfy the monotonic consequences of announcing ϕ; and $\mathcal{M}^{cp(\phi)}$ is formed by eliminating all states that don't satisfy the nonmonotonic consequences of announcing ϕ. Note that because $>$ is supra-classical (see (3a)), ACPs, like 'simple' announcements, presuppose that the announcement is true; i.e. $S^{cp(\phi)} \subseteq [\![\phi]\!]$. This means that the ULF of an utterance is always a ULF for the entire dialogue; it does *not* mean that the the utterance is true or even that the speaker is committed to it.

Glue logic axioms like (2b) make the consequences of ACPs express information about rhetorical connections or specific values for other underspecified elements introduced by linguistic syntax. The axioms from Lascarides and Asher (2009) (but omitted here) also ensure that ACPs predict which commitments from prior turns are current commitments. For instance, for dialogue (1), the axioms ensure that $\mathcal{M}, s \models [!\mathcal{K}_{\pi_3}]^{cp}([!(\pi_{2K}{:}?(\pi_2, \pi_3) \wedge T(K, 2, \pi_{2K}))]^{cp}\pi_{2K}{:}\mathsf{Explanation}(\pi_1, \pi_2))$, where $\mathcal{M}$ is the model constructed by updating with utterances π_1 and π_2 in that order and $s \in \mathcal{M}$.

It is now simple to define discourse update within the logic. We imagine that the set of DSDRSs σ is simply the set of states of a model $\mathcal{M}_\sigma$:

Definition 5 Dynamic Simple Update:
$\sigma + \phi \models \psi$ *iff* $M_\sigma \models [!\phi]^{cp}\psi$

To define full DSDRS update, we take Boolean combinations of ACP updates so as to match the Discourse Update process from Definition 2 (see its second paragraph).

Definition 6 Dynamic Discourse Update:
Let $\mathcal{M}_\sigma$ be 'old' information and the ULF *$\mathcal{K}_\beta$ be new information. Let $\Sigma_1, \ldots \Sigma_n$ be all the jointly composible attachment sites of β, chosen from the set A of all possible attachment sites for each* DSDRS *in σ. Let k_i be an enumeration of the compossible attachment sites in Σ_i, $1 \leq i \leq n$. And let k_i be the sequence of assumptions about attachment provided by the enumeration k_i of sites in Σ_i:*

$$\begin{aligned}\mathsf{k}_i = \;& \lambda_1^i{:}?(\alpha_1^i, \beta) \wedge T(d, j, \lambda_1^i) \wedge \ldots \\ & \wedge \lambda_{k_i}^i{:}?(\alpha_{k_i}^i, \beta) \wedge T(d, j, \lambda_{k_i}^i)\end{aligned}$$

Then
$$\begin{aligned}& Update(\mathcal{M}_\sigma, \mathcal{K}_\beta) \models \psi \textit{ iff } \forall s \in S^{\mathcal{M}_\sigma}, \\ & \mathcal{M}_\sigma, s \models [!(\mathcal{K}_\beta \wedge last = \beta)](\textstyle\bigwedge_{i=1}^n [!\mathsf{k}_i]^{cp}\psi)\end{aligned}$$

We now also axiomtise MDC from Definition 3 within the glue logic. We will take MDC as imposing a coherence order on what is announced—$\phi \preceq \psi$ means in words that *ceteris paribus* announcing ϕ (in this particular context) is less coherent than *ceteris paribus* announcing ψ. We assume that the ϕ and ψ that get ordered are essentially states that verify formulae in the background description of the discourse: in other words, if the (partial) description of the logical form of the discourse context is $Th(\sigma)$ and $\mathcal{K}_\beta$ is the ULF for the new information, then ϕ and ψ resolve some underspecified semantic elements in $Th(\sigma) \wedge \mathcal{K}_\beta$. Without loss of generality, we can assume that ϕ and ψ are conjunctive formulae, where at least one of the conjuncts is an assumption $\lambda{:}?(\alpha, \beta)$ about attachment of the new information. This is because if ϕ and ψ differ only in how some other aspect of underspecified content is resolved—such as the antecedent to a pronoun, for instance—then ϕ and ψ can include the *same* conjunct $\lambda{:}?(\alpha, \beta)$ and differ only in the conjunct that fixes the antecedent. Since discourse must be coherent, we know that β must attach to something, and hence there is no loss of generality in including that attachment in all ACPs. We also assume a partial ordering $\leq$ on rhetorical relations: $R \leq R'$ means that R is a less coherent relation than R' (see clause 3. from Definition 3). For example, $\mathsf{Background} \leq \mathsf{Explanation}$ would make MDC prefer interpreting new information as an *Explanation* rather than as a *Background*, all else being equal.

The principles that govern degree of coherence are then stipulated in Definition 7. Clause 1. says that an ACP ϕ that yields a less coherent rhetorical connection compared with the ACP ψ is normally less coherent. Clause 2. says that an ACP ϕ that resolves fewer underspecified elements than ACP ψ is normally less coherent. Finally, clause 3. says that an ACP ϕ that results in a logical form with more segments than ACP ψ is normally less coherent. Stipulating clauses 2. and 3. calls for an extension to our object language $\mathcal{L}_{ulf}$, where all variables—including higher order ones—are typed as variables and distinguished from types corresponding to constant symbols. We also need existential quantification over labels and variables (see clauses 3. and 2. respectively).

Definition 7 Coding up MDC

1. $(Th(\sigma) \wedge \mathcal{K}_\beta) > \phi \preceq \psi$ *if* *there's a permutation f on $\Pi_{Th(\sigma)} \cup \Pi_{K_\beta}$ st* $R' \leq R \wedge \forall \pi_1 \pi_2([!\psi]^{cp} R(\pi_1, \pi_2) \rightarrow [!\psi]^{cp} R'(f(\pi_1), f(\pi_2)))$

2. $(Th(\sigma) \wedge K_\beta) > \phi \preceq \psi$ *if* $[!\phi]^{cp} \exists_{\geq n} ?_1, \ldots ?_n \rightarrow [!\psi]^{cp} \exists_{\geq n} ?_1, \ldots ?_n$

3. $(Th(\sigma) \wedge K_\beta) > \phi \preceq \psi$ *if* $[!\psi] \exists_{\geq n} \pi_1, \ldots \pi_n \rightarrow [!\phi] \exists_{\geq n} \pi_1, \ldots \pi_n$

This is obviously only an approximation of the principles described in Definition 3. We have not, for instance, encoded the principle that interpretations with more rhetorical connections are more coherent than those without (see clause 1. from Definition 3). But given that the number of rhetorical relation symbols in $\mathcal{L}$ is finite, it would be very straightforward to express this coherence factor via the ordering relation $\leq$ on predicate symbols in $\mathcal{L}_{ulf}$.

The axioms in Definition 7 can be used to influence interpretation. Our existing update function *Update*$(\mathcal{M}, \phi_\beta)$ abstracts away entirely from how interpretation is influenced by degree of coherence. But we add to it a new update function *Best-update*$(\mathcal{M}, \phi_\beta)$, whose definition is exactly like that of *Update*, save that the consequences of the announcement $[!\mathcal{K}_\beta \wedge \mathit{last} = \beta]$ are restricted to those conjuncts about attachment that are maximal on the partial ordering $\preceq$ given by Definition 7. Thus a speaker can anticipate what the most coherent interpretation of his announcement will be, as well as the range of possible coherent interpretations, as given by *Update*.

The logic will support an inference that $\phi \preceq \psi$ only if one or all the axioms in Definition 7 are verified—so just like Definition 3 of MDC, ψ must be at least as coherent as ϕ in all three respects, and more coherent in at least one of them. If, for example, ϕ is a logical form with more segments than ψ, but ψ features lower quality rhetorical connections, then the default axioms whose antecedents are satisfied will conflict, with one having consequent $\phi \preceq \psi$ and the the other having conflicting consequent $\psi \preceq \phi$. This results in a Nixon Diamond and no inferences about the relative coherence of the announcements ϕ and ψ.

We have extended the language to include existential quantification over a *finite* number of labels and variables, and permutations over a finite number of labels. Because the quantification and permutations are over finite domains we can do without quantification. So $\mid\!\sim$ remains decidable even if the premises and conclusions are Σ_1-formulae. But we must now extend the models to include a fixed set of labels that remains constant as we move from M to M^ϕ.

The computational complexity of PAL is demonstrated by proving **reduction axioms and rules** (Balbiani et al., 2007). The reduction axioms for the $[!\phi]$ operator are quite standard—axiom IV is equivalent to one given in van Benthem (2007), for instance, though we offer a proof for this reduction axiom here.

I $[!\phi]p \leftrightarrow (\phi \rightarrow p)$

II $[!\phi](\psi \wedge \chi) \leftrightarrow ([!\phi]\psi \wedge [!\phi]\chi)$

III $[!\phi]\neg\psi \leftrightarrow (\phi \rightarrow \neg[!\phi]\psi)$

IV $[!\phi](\psi > \chi) \leftrightarrow (\phi \rightarrow ((\phi \wedge [!\phi]\psi) > [!\phi]\chi))$

As regards axiom IV, note that $[\![\phi \wedge [!\phi]\psi]\!]^M = [\![\psi]\!]^{M^\phi}$, since $s \in [\![\phi \wedge [!\phi]\psi]\!]^M$ iff $s \in [\![\phi]\!]^M \cap [\![[!\phi]\psi]\!]^M$ iff $s \in [\![\psi]\!]^{M^\phi}$. To prove axiom IV, we observe $*^{M^\phi}([\![\phi \wedge [!\phi]\psi]\!]^M, s) \subseteq *^M([\![\phi \wedge [!\phi]\psi]\!]^M, s)$ by definition. Since $*$ is reflexive in the model theory of GL and common sense entailment (i.e., $*(p, s) \subseteq p$), $\forall s' \in *^M([\![\phi \wedge [!\phi]\psi]\!]^M, s)$, $s' \in [\![\phi]\!]$, and this means $*^{M^\phi}([\![\phi \wedge [!\phi]\psi]\!]^M, s) = *^M([\![\phi \wedge [!\phi]\psi]\!]^M, s)$. This suffices to prove axiom IV, since $M^\phi, s \models \psi > \chi$ iff $*^{M^\phi}([\![\psi]\!]^{M^\phi}, s) \subseteq [\![\chi]\!]^{M^\phi}$ iff $*^M([\![\phi \wedge [!\phi]\psi]\!]^M, s) \subseteq [\![\chi]\!]^{M^\phi}$ iff $*^M([\![\phi \wedge [!\phi]\psi]\!]^M, s) \subseteq [\![\phi \wedge [!\phi]\chi]\!]^M$ iff $*^M([\![\phi \wedge [!\phi]\psi]\!]^M, s) \subseteq [\![[!\phi]\chi]\!]^M$ (the last step follows again because of the reflexivity of $*$ in GL). $\square$.

The more interesting question concerns ACPs. Simply using the definition of ACPs, we have the following additional reduction axioms:

V $$\frac{\Gamma \vdash [!\phi]^{cp}\psi}{\Gamma, \phi \mid\!\sim \psi}$$

VI $$\frac{\Gamma, \phi \mid\!\sim \psi}{\Gamma \vdash [!\phi]^{cp}\psi}$$

Rules V and VI follow directly from Definition 4. A strengthened reduction rule like (4) using *defeasible* inference with ACPs is *not* valid, however, because like many nonmonotonic logics ours suffers from the Drowning Problem (Benferhat et al., 1993)—defaults from ϕ 'drown out' those from Γ when they are mixed together.

(4) $$\frac{\Gamma \mid\!\sim [!\phi]^{cp}\psi}{\Gamma, \phi \mid\!\sim \psi}$$

The problem comes from nested conditionals, such as those in (5).

(5) a. $\Gamma\colon \{A, A > D, A > ((B \wedge E) > C)\}$
b. $\phi\colon B \wedge E \wedge ((A \wedge E) > \neg D)$
c. $\psi\colon C$

In (5), $\Gamma \mid\!\sim (B \wedge E) > C$ and therefore $\Gamma \mid\!\sim [!\phi]^{cp}\psi$. But $\Gamma, \phi \not\mid\!\sim \psi$ since $\Gamma, \phi \not\mid\!\sim (B \wedge E) > C$—the defaults from ϕ 'drown out' those from Γ when they are mixed together.

Using the prime implicate $\mathfrak{I}_\phi$, we can get proper reduction axioms for $[!\phi]^{cp}\psi$. This is because the prime implicate encapsulates the nonmonotonic reasoning inherent in $\mid\!\sim$. The connection between $\mathcal{M}$ and $\mathcal{M}^{cp(\phi)}$ is this: $[\![\psi]\!]^{\mathcal{M}^{cp(\phi)}} = [\![\mathfrak{I}_\phi \wedge [!\phi]^{cp}\psi]\!]^{\mathcal{M}}$.

VII $[!\phi]^{cp}p \leftrightarrow (\Im_\phi \rightarrow p)$

IX $[!\phi]^{cp}\neg\psi \leftrightarrow (\Im_\phi \rightarrow \neg[!\phi]^{cp}\psi)$

X $[!\phi]^{cp}(\psi \wedge \chi) \leftrightarrow ([!\phi]^{cp}\psi \wedge [!\phi]^{cp}\chi)$

XI $[!\phi]^{cp}(\psi > \chi) \leftrightarrow$
$(\Im_\phi \rightarrow ((\Im_\phi \wedge [!\phi]^{cp}\psi) > [!\phi]^{cp}\chi))$

Given the restrictions that hold of GL's background theory, prime implicates exist and can be decidably computed (Asher, 1995). Furthermore, the base logic of GL is decidable (Asher and Lascarides, 2003). These reduction axioms thus ensure that our extension of PAL is decidable as well.

Fact 1 *Dynamic GL is decidable*

As to the proofs of (VII-XI), they pattern closely with the proofs for (I-IV). We sketch here proofs for (VII) and (XI). To prove the base case, assume for the left to right direction that $M, s \models [!\phi]^{cp}p$. If $s \notin [\![\Im_\phi]\!]^M$, then we are done. So assume $s \in [\![\Im_\phi]\!]^M$. So $s \in S^{cp(\phi)}$, and by the satisfaction definition $M^{cp(\phi)}, s \models p$. The right to left direction follows straightforwardly from the definitions.

To prove axiom XI, start by observing that $*^{M^\phi}([\![(\Im_\phi \wedge [!\phi]^{cp}\psi]\!]^M, s) \subseteq *^M([\![\Im_\phi \wedge [!\phi]^{cp}\psi]\!]^M, s)$ by definition. Since $*$ is reflexive in the model theory of GL and common sense entailment, $\forall s' \in *^M([\![\Im_\phi \wedge [!\phi]^{cp}\psi]\!]^M, s)$, $s' \in [\![\phi]\!]$, and this means $*^{M^{cp(\phi)}}([\![\Im_\phi \wedge [!\phi]^{cp}\psi]\!]^M, s) = *^M([\![\Im_\phi \wedge [!\phi]^{cp}\psi]\!]^M, s)$. This suffices to prove (XI), since $M^{cp(\phi)}, s \models \psi > \chi$ iff $*^{M^{cp(\phi)}}([\![\psi]\!]^{M^{cp(\phi)}}, s) \subseteq [\![\chi]\!]^{M^{cp(\phi)}}$ iff $*^M([\![\Im_\phi \wedge [!\phi]^{cp}\psi]\!]^M, s) \subseteq [\![\chi]\!]^{M^{cp(\phi)}}$ iff $*^M([\![\Im_\phi \wedge [!\phi]^{cp}\psi]\!]^M, s) \subseteq [\![\Im_\phi \wedge [!\phi]^{cp}\chi]\!]^M$ iff $*^M([\![\Im_\phi \wedge [!\phi]^{cp}\psi]\!]^M, s) \subseteq [\![[!\phi]^{cp}\chi]\!]^M$ (the last step follows again because of the reflexivity of $*$ in GL).

The reduction axioms VII to XI go beyond those given in Baltag and Smets (2006) and van Benthem (2007) since they don't include in their logics the distinct *ceteris paribus* type of announcement. None of the reduction axioms VII to XI are particular to the $>$-axiom of (3d) or to the Penguin Principle, however. These latter properties of our logic are reflected in the nature of the prime implicates and by the particular inferences our PAL logic will license. We note that it is due to the particular characteristics of SDRT's glue logic that such prime implicates exist.

5 Conclusion

In this paper we have made SDRT's glue logic for computing logical form dynamic. This allows a dialogue agent to reason about what the update of the DSDRS will be after his contribution, including the effects of his candidate rhetorical moves. This is a prerequisite for *planning* one's next move, but so is reasoning about attitudes like preferences. The next step is to examine SDRT's other shallow logic, the logic of cognitive modeling, so as to optimise the trade offs between expected interpretations and speaker preferences.

References

CE Alchourrón and P. Gärdenfors and D. Makinson. On the Logic of Theory Change: Partial Meeting Contraction and Revision Functions, *Journal of Symbolic Logic*, 50:510–530, 1985.

N. Asher. Commonsense entailment: A conditional logic for some generics. In G. Crocco *et al* (eds) *Conditionals: From Philosophy to Computer Science*, pages 103–145. OUP, 1995.

N. Asher and A. Lascarides. *Logics of Conversation*. Cambridge University Press, 2003.

N. Asher and M. Morreau. Commonsense Entailment. In *Proceedings of the 12th International Joing Conference on Artificial Intelligence*, pages 387–392.

P. Balbiani, H. van Ditmarsch, A. Herzig, and T. de Lima. A Tableau Method for Public Announcement Logics. In *Proceedings of the International Conference on Automated Reasoning with Analytic Tableaux and Related Methods (TABLEAUX)*, 2007.

A. Baltag and S. Smets. Conditional Doxastic Models: A Qualitative Approach to Dynamic Belief Revision. in G. Mints and R. de Queiroz (eds.), *Electronic Notes in Theoretical Computer Science*, 165, pages 5–21, 2006.

A. Baltag, L.S. Moss, and S. Solecki. The Logic of Public Announcements, Common Knowledge and Private Suspicions. Technical Report SEN-R9922, Centrum voor Wiskunde en Informatica, 1999.

S. Benferhat, D. Dubois, and H. Prade. Argumentative Inference in Uncertain and Inconsistent Knowledge Bases. In *Proceedings of Uncertainty in Artificial Intelligence*, pages 411–419, 1993.

J. van Benthem. Dynamic Logic for Belief Revision. *Journal of Applied Non-Classical Logics* 17(2):129-155, 2007.

M. Egg, A. Koller, and J. Niehren. The Constraint Language for Lambda Structures. *Journal of Logic, Language, and Information*, 10:457–485, 2001.

J. R. Hobbs, M. Stickel, D. Appelt, and P. Martin. Interpretation as Abduction. *Artificial Intelligence*, 63(1–2):69–142, 1993.

H. Kamp and U. Reyle. *From Discourse to the Lexicon*. Kluwer Academic Publishers, 1993.

A. Lascarides and N. Asher. Temporal Interpretation, Discourse Relations and Commonsense Entailment. *Linguistic and Philosophy*, 16(5):437–493, Kluwer Academic Publishers, 1993.

A. Lascarides and N. Asher. Agreement, Disputes and Commitment In Dialogue. *Journal of Semantics*, 26(2):109–158, 2009.

D. Lewis. *Convention: A Philosophical Study*. Harvard University Press, 1969.

M. Poesio and D. Traum. Towards an Axiomatisation of Dialogue Acts. In *Proceedings of SEMDIAL*. 1998.

H. Sacks, E. A. Schegloff, and G. Jefferson. Systematics for the Organization of Turn-Taking in Conversation. *Language*, 50(4):696–735, 1974.

M. Stone. *Modality in Dialogue: Planning, Pragmatics and Computation*. PhD thesis, University of Pennsylvania, 1998.

R. van Rooij. Relevance of Communicative Acts. In *Proceedings of TARK*, pages 83–97, 2001.

JD Williams and SJ Young. Partially Observable Markov Decision Processes for Spoken Dialog Systems. *Computer Speech and Language*, 21(2): 231–422, 2007.