

Reporting, telling, and showing dreams

Emar Maier

University of Groningen

Abstract Dreams are not real, so when we recount them we prefix an intensional operator like ‘I dreamed that . . .’. Linguists will analyze this construction in terms of clausal complementation syntax and possible worlds semantics. But talking about a dream is often more like telling a story, with a potentially complex discourse structure (involving propositional discourse units connected by coherence relations like NARRATION, BACKGROUND, and EXPLANATION) that is hard to fit inside a single syntactically embedded *that*-clause (or a sequence of independently embedded clauses). I show how we can analyze actual, complex dream report stories using a formal discourse semantics framework. I then explore how to extend this discourse framework to visual dream reporting, like in movies and comics, where it’s not immediately clear that we even have any intensional operators or embeddings to begin with.

keywords: dream reports, semantics, narrative, discourse structure, comics, film

Contents

1	Dreaming that p	2
1.1	Dream reports as intensional operators	2
1.2	Beyond the syntax/semantics interface	3
2	Dream reporting as storytelling	7
2.1	Talk about dreaming in context	7
2.2	Discourse and coherence	8
2.3	Attribution	11
2.4	Reporting with embedding and without	12
2.5	Dream discourse revisited	14
3	Showing dreams	18
3.1	Discourse across media	18
3.2	PicDRT and comics panels	19
3.3	SDRT and comics	20
3.4	Visual dream attribution	22

3.5	Dream bubbles	22
3.6	Dream sequences	24
4	Conclusion	27

1 Dreaming that *p*

1.1 Dream reports as intensional operators

Semanticists’ interest in dream reporting grew out of a general interest in reporting mental states with attitude verbs and *that*-clause, first in philosophy (Frege 1892, Quine 1956) and later also in linguistic semantics (Schlenker 2003).

A classic and much discussed example of a linguistic peculiarity of dream reporting is (1) (ascribed to McCawley by Lakoff 1968):

- (1) I dreamed that I was Brigitte Bardot and that I kissed me.

The observation is that while ‘I kissed me’ as an independent main clause seems ungrammatical (cf. ‘I kissed myself’, which is grammatical but pragmatically odd), it is grammatical and meaningful in the context of this dream report. In fact, it means something quite different from the alternative with the reflexive.

Inspired by (1), semanticists have been investigating the different readings that arise for pronouns in dream reports. von Stechow (1982) observes that pronouns in dream reports are ambiguous between (i) a *de se* reading, where the pronoun refers to the dream-self, i.e., the first-person protagonist of the dream, and (ii) a *de re* reading where the pronoun refers to the actual dreamer. Since the actual dreamer may feature in the dream as a third person, distinct from the dream-self, these readings can come apart. In (1) we can thus get a reading where the *de se* dream-self (Bardot) kisses the *de re* subject (McCawley). With two pronouns this would predict up to four distinct readings. Interestingly, Percus & Sauerland (2003) observe that one of these is systematically ruled out: (2a) cannot report a dream where John (*de re*) is marrying the dream-self’s (*de se*) granddaughter (i.e., a dream in which the dream-self might exclaim “Oh no. that guy John is marrying my granddaughter!”). Similarly, Anand (2006) observes that (2b) doesn’t allow the most likely intended reading where the actual dreamer (*de re*) is chopping up the dream-self-carrot (*de se*).

- (2) a. John dreamed that he was marrying his granddaughter.
[Percus & Sauerland 2003]
 b. #I dreamed that I was a carrot and that I chopped me up for dinner.
[Anand 2006]

The empirical generalization is that in dream complements, a *de re* pronoun may not c-command a *de se* pronoun.¹ Ninan (2008) further complicates the picture by noting that orthogonal to the *de re/de se* distinction we also have a so-called inside/outside distinction, which will increase the number of potential readings (you can dream that you (*de se*) are chopped up for dinner “from the inside”, i.e., imagining being chopped up, or from the outside, i.e., witnessing your carrot-self being chopped up as seen from a third-person perspective while still screaming “Oh no I’m being chopped up!”).

In sum, philosophers and linguists have uncovered some puzzling properties of dream report constructions, with much room for further research. Note that all these issues so far arise squarely within the syntax–semantics interface, involving intensional operators, intra-sentential pronoun binding and co-reference, and *de re/de se* ambiguities. In the current paper I want to go beyond the syntax and semantics of ‘*x* dreamed that *p*’ constructions and look at dream reporting in the semantics–discourse interface. The phenomena I’m interested in here have gone unnoticed by the linguists and philosophers cited above, but in fact the classic Brigitte Bardot example already illustrates one key aspect thus far overlooked: real dream reports do not fit inside a single embedded *that*-clause.

1.2 Beyond the syntax/semantics interface

The first conjunct of (1) is indeed just an attitude verb with a clausal complement (‘I dreamed that I was Brigitte Bardot’), but the example continues with what appears to be a second *that*-clause embedded under a second, and therefore presumably phonologically elided, dream verb.

(3) I dreamed that I was Brigitte Bardot and <I dreamed> that I kissed me.

Intuitively, the second complement further specifies the content of the dream introduced in the first. But on the reconstruction in (3), a standard semantic analysis of attitude verbs as quantifying over possible worlds or contexts (Hintikka 1969, Kaplan 1989, Schlenker 2003) would not obviously derive that these complements combine to describe a single dream. Alternative syntactic reconstructions, more sophisticated dynamic or event-based attitude semantics, and/or hidden anaphoric elements might help us out with this particular made-up example. But there’s a deeper problem lurking in the background: when we look at naturally occurring dream reports, or introspect how we recount our own dreams, we find that most dream reports extend beyond the single *that*-clause in various ways, typically involving not just multiple conjoined *that*-clauses, but many separate, independent

¹ See Pearson & Dery (2013) for empirical validation of the effect.

sentences. Here is how a chatbot quite naturally continues a story that starts with “I dreamed that I was Brigitte Bardot”.²

- (4) I dreamed that I was Brigitte Bardot. The sun-kissed beaches of the French Riviera stretched out before me, their golden sands inviting and warm. My hair, a cascade of honey-blond waves, framed my face as I strolled along the shoreline. [...] Paparazzi trailed me, their cameras clicking incessantly. But I paid them no mind; after all, I was Brigitte – the enchantress of cinema, the embodiment of sensuality. [...] But as dawn approached, reality seeped in. I was not Brigitte Bardot; I was just a dreamer lost in reverie. I woke up, my heart heavy with longing, and wondered if dreams could ever be more than fleeting illusions.

The whole story is several paragraphs long, and the majority of sentences, apart from the opening line and the concluding paragraph, all clearly describe events and states that occur only in the dream, involving a dream-self who is the famous actress. Only the opening line involves an overt embedding, the rest are just independent main clauses. So, do we really want to maintain that all these sentences are prefixed with hidden operators that moreover maintain main clause syntax?

One attempt to motivate such a hidden operator proposal is to gesture at the broader phenomenon of ‘modal subordination’ (Roberts 1987):

- (5) a. If Edna forgets to fill the bird feeder, she will feel bad. The birds will get hungry.
b. A wolf might come in. It would eat you first.(Roberts 1989)

In both of the above, the second sentence, a seemingly independent main clause, appears to be interpreted relative to some kind of hidden modal operator or conditional antecedent that is reconstructed on the basis the previous sentence: <in all those worlds where Edna forgets> the birds will get hungry, or: <if the wolf comes in> it will eat you first. Indeed, in (4) we likewise have one overt operator, and a sequence of main clauses interpreted as if under a modal that we can reconstruct from that first overt dream operator (by simply copying that first operator, as we did in (3)).³

Of course, pointing out some parallels to a known linguistic phenomenon doesn’t really help us much unless we have an independently motivated, explanatory (and preferably elegant and commonly accepted) analysis of that second phenomenon

² Bing Copilot, April 2024, prompt: “write a short story that starts with ‘I dreamed that I was Brigitte Bardot’”.

³ Note that (4) does not actually fall under Roberts’ definition of modal subordination since she stresses the importance of a modal element in the second sentence (i.e., the ‘will’ and ‘would’ in (5)).

– which is unclear in this case.⁴ More importantly, the (already tenuous) analogy with modal subordination breaks down in cases like (6) where a dream is introduced without any overt operators at all.

(6) Last night I had a terrible nightmare. I was chased by a giant octopus.

In section 3 we'll go a couple of steps further and look at dream reporting conventions in non-verbal discourse, like comics and film, where likewise there's often no evidence of any overt intensional operators and hence no question of modal subordination.

Next, one might point to the well-known phenomenon of free indirect discourse, where sequences of independent main clauses are similarly interpreted as descriptions of an individual's mental state, arguably involving hidden quotation marks (Maier 2015) and/or (monstrous) attitude operators (Schlenker 2004, Sharvit 2008). Charnavel (2024) indeed suggests analyzing complex dream report discourses like our (4) as involving free indirect discourse. She discusses the following concrete example (in the West-African language Ewe):⁵

(7) Kofi koudrin be yè bidzi. Marie zu yè.
Kofi dream COMPL LOG angry Mary insult LOG
'Kofi dreamed that he was angry. Mary insulted him.' [Ewe, Pearson 2015]

However, none of the supposedly covertly prefixed report sentences in (4) show any of the characteristics of free indirect discourse: no shifted temporal indexicals or epithets, no linguistic/idiolectal peculiarities (ascribed to the dreamer), no exclamatives, vocatives, questions, or other non-assertive speech acts, etc. (Banfield 1982, Fludernik 1995, Maier 2015, Eckardt 2014). So, following the argumentation laid out in Bary & Maier (2014) (against similar claims of finding free indirect discourse, in Ancient Greek), what we have here are instances of what they call Unembedded Indirect Discourse (UID).

In any case, if we go down the hidden operators road, whether by invoking modal subordination, FID, UID, or some other mechanism, one empirical problem we run into is that we might be unable to account for some of the apparent coherence

⁴ See Roberts (2020) for an overview of the current state of the debate that shows the complexity of the “puzzle” of modal subordination.

⁵ In the context of her manuscript Charnavel is primarily interested in the occurrence of the logophoric pronoun in (7), and is somewhat cautious in classifying this as free indirect discourse (FID): “The second sentence in [(i)] is not syntactically embedded under the attitude verb *dream* and the complementizer *be*. But as noted by Pearson, it is semantically embedded as it implies that the insult took place in the dream. This seems to correspond to FID”. Pearson (2015) uses the example to confirm an observation from Clements (1975) about the use of logophoric pronouns in such dream report continuations, but neither Pearson nor Clements uses the term FID in this connection.

inferences that we read between the lines in a discourse. For instance, adding the hypothesized hidden operators in the middle of (4), we'd get:

- (8) ... <I dreamed that> Paparazzi trailed me, their cameras clicking. <I dreamed that> (But) I paid them no mind; ...

Ignoring the awkward clash between the main clause syntax and the reconstructed operator, the use of 'but' in the second sentence marks a perceived contrast between being trailed by paparazzi (stressful), and paying the paparazzi no mind (relaxed). Both of these contrasting events, and hence the contrast between them,⁶ clearly exist only in the dream-world (as presented by the dreamer's re-telling). In intensional operator terms, the intended interpretation corresponds to the dream operator outscoping the contrast marker:

- (9) I dreamed that ((paparazzi trailed me) but (I paid them no mind)).

Adding two separate operators at the level of sentential syntax/semantics, as in (8), would allow us at best to interpret 'but' as marking a contrast between the two dream reports:

- (10) (I dreamed that (paparazzi trailed me)) but (I dreamed that (I paid them no mind)).

This second reading makes less sense, as there's no obvious contrast between having a dream where you're trailed by paparazzi, and having a dream where you pay the paparazzi no mind. Even if we manage to force these copied operators in different independent sentences to report on the same dream, we know that dreaming frequently involves non-sequiturs so the contrast between dreaming about a paparazzi chase and dreaming about staying relaxed is at least weaker than that between a paparazzi chase and staying relaxed. In other words, the most likely reading of (8) is one where the teller is signaling a contrast between two dreamed events, not a contrast between two dreaming events.

In the next section I'll show how we can derive the intended interpretations of examples like (4) and (6) without any hidden operators, relying solely on independently motivated mechanisms of inferring discourse coherence.

⁶ Stokke (p.c.) observes that in similar cases it's not so clear that the dreamer feels this contrast in the dream itself, but rather constructs it in the real-world telling of the dream. What we're analyzing here however is precisely the semantics of the dreamer's recounting of the dream. Regardless of what actually went on in their mind/brain when they were asleep, the dreamer, in their telling, presents the dream-world as involving a contrast between two events, which means that the contrast should be semantically represented at the same level as the dreamed events, not at the level of the telling events.

2 Dream reporting as storytelling

2.1 Talk about dreaming in context

Let's consider some actual human dream reports, drawn from the dreambank.net corpus (Schneider & Domhoff 2024).

- (11) I dreamed I was at a party. There were many people crowded around an extremely long tea table. As I started to talk to the people who were about my age (M & F 20-25) my teeth began to drop out. One by one each tooth fell out and soon my face started to fold. [...]

We have here the now familiar pattern: an overtly embedded dream report construction, followed by a number of independent main clauses that are clearly intended as further descriptions of the subject's dream. The corpus contains many examples of this general form, but also many where there is no overt report embedding. Many of those are just the story of the dream, without any 'frame' mentioning the dreamer having a dream, others are more like (6) above, i.e., there's an opening sentence that explicitly mentions a dream but without an intensional operator or clausal embedding:

- (12) This dream took place in P.. My mother and I were in the public library together. I left her for a few moments and met an older man (who is about 40 years old, was the chaperon on a couple of geology trips I was on and who apparently took a great interest in me). He asked me to go to a bar and have a couple of drinks with him. We started walking toward the city hall and all at once the street and sidewalks were crowded with 'madly' rushing people shouting and screaming.

Although the second sentence of (12) is in principle ambiguous between a description of the dreamer in the real world (we went to the library and there I fell asleep and dreamed ...) and a description of the dream world, it quickly becomes apparent that the latter interpretation is more likely for this and the following sentences – with the exception of the bracketed remark, where the teller appears to temporarily 'pop out' of their dream description.

In dream reports like these, as in any type of storytelling, the interpreter is presented with a sequence of sentences, describing states and events, and has to make sense of them as telling a coherent story, by inferring logical connections between these described states and events (henceforth: eventualities). Determining which sentences (or rather 'discourse units') describe eventualities in the actual world and which describe eventualities in a dream is part of this general inferential process of maximizing discourse coherence. The process can be aided by overt

embeddings under a dream operator, but does not require such operators. Below I present a well-established discourse semantic framework designed to study these processes in more detail.

2.2 Discourse and coherence

Before analyzing the discourse structure of complex dream reports, consider a simple narrative:

- (13) It was late. Susan got up from the couch and looked in the fridge. She was hungry.

Each of these sentences individually describes one or more eventualities (e.g. the event of getting up from the couch and the state of being hungry). When we read them in sequence, as part of a coherent discourse, we reliably and effortlessly infer certain connections, known as discourse relations, between these eventualities: *while* it was getting late, Susan got up, and *then* she looked in the fridge, *because* she was hungry. Discourse semantics studies these and other inferential processes in discourse interpretation.

The most comprehensive and formally precise theory of discourse semantics to date is Segmented Discourse Representation Theory (SDRT, Asher & Lascarides 2003). It is often presented as an extension of Discourse Representation Theory (DRT, Kamp 1981), adding a formal semantic theory of discourse relations and discourse structure. I'll introduce the basics of SDRT by applying it to (13), i.e., I will show how the framework generates a logical, model-theoretically interpretable, representation that captures both the sentential and discourse semantic content of the story in (13).

The first step in discourse interpretation is segmentation: we split the discourse into minimal proposition-expressing units, typically clauses:

- (14) π_1 :It was late.
 π_2 :Susan got up from the couch
 π_3 : (and) looked in the fridge.
 π_4 :She was hungry.

Each elementary discourse unit, π_i , asserts the existence of eventualities and/or other entities and describes certain properties and relations between them. We'll represent these elementary semantic contents in basic DRT box notation: the top compartment lists the novel eventualities and other entities introduced into the discourse record as discourse referents, and the bottom compartment specifies some descriptive properties and relations linguistically encoded in the clause. Since we're

interested here in the discourse structural semantics (rather than the intra-clausal semantic composition), we'll drastically abbreviate and simplify our DRS boxes. We adopt a neo-Davidsonian approach to clausal semantics so that (most) clauses introduce at least a discourse referent for an eventuality (e.g., π_2 introduces an event discourse referent e_2 of getting up along with an entity discourse referent x_2 for the agent Susan).⁷

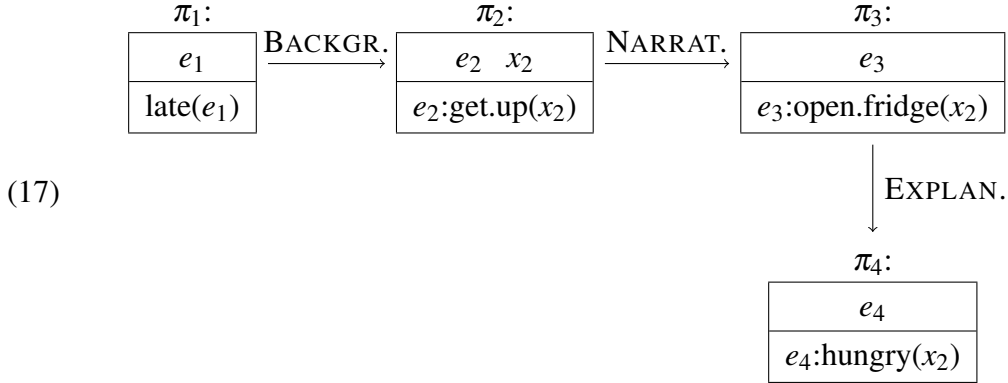
$$(15) \quad \pi_2 \rightsquigarrow \boxed{\begin{array}{c} e_2 \ x_2 \\ e_2:\text{get.up}(x_2) \end{array}} \dots$$

Next, we infer discourse relations between the elementary units, like BACKGROUND, NARRATION, EXPLANATION. These may be explicitly encoded with overt connectives like 'while', 'and then', and 'because', respectively, but they can also be left implicit, as in (13), in which case they have to be recovered by the hearer based on a variety of contextual and linguistic cues. For instance, a state description followed by an event description may indicate a BACKGROUND relation, while two event descriptions instead may indicate NARRATION. SDRT's so-called Glue Logic axioms specify which configurations defeasibly entail which relations. Since multiple discourse relation entailments may well be possible, we measure the global coherence of the resulting potential discourse representation options (by considering numbers of anaphoric connections, numbers of discourse relations, logical consistency etc.). The end result is a maximally coherent directed graph structure with discourse units as nodes and discourse relations as edges. When such graphs are visualized, we typically draw coordinating discourse relations, that move the story forward, as horizontal, labeled arrows, and subordinating relations, where a subordinate unit elaborates on a "nuclear" unit, as vertical arrows.

$$(16) \quad \begin{array}{ccccc} \pi_1 & \xrightarrow{\text{BACKGROUND}} & \pi_2 & \xrightarrow{\text{NARRATION}} & \pi_3 \\ & & & & \downarrow \text{EXPLANATION} \\ & & & & \pi_4 \end{array}$$

I will combine the discourse structure graph and the (simplified) DRS box representations of the elementary units into a single, comprehensive structured discourse representation:

⁷ We ignore the fact that definites, including 'Susan' and 'the couch' are better analyzed as presupposition triggers.



SDRT comes with a model-theoretic interpretation that extends that of DRT. Assuming full-fledged representations like (17), we interpret the graph by interpreting each pair of nodes connected by an arrow, following the directions of the arrows:

$$(18) \quad \llbracket \alpha_1 \longrightarrow \alpha_2 \longrightarrow \dots \rrbracket = \llbracket \alpha_1 \longrightarrow \alpha_2 \rrbracket \circ \llbracket \alpha_2 \longrightarrow \dots \rrbracket$$

We're assuming a dynamic semantic interpretation of the elementary discourse unit DRSs, i.e., $\llbracket \alpha \rrbracket$ in (18) denotes a function that maps an input information state (input context) to an updated output information state (output context). \circ denotes function composition, a kind of dynamic conjunction, i.e., when p and q are context change potentials, $p \circ q$ is a context change potential that tells us to first update the input context with p and then with q (Kamp et al. 2003, Asher & Lascarides 2003).

Next, we define for each of the finite set of discourse relations a dynamic semantic interpretation:

$$(19) \quad \begin{array}{l} \text{a.} \quad \left\| \begin{array}{c} \alpha \\ \downarrow \text{BACKGROUND} \\ \beta \end{array} \right\| = \llbracket \alpha \rrbracket \circ \llbracket \beta \rrbracket \circ \llbracket e_\alpha \circ e_\beta \rrbracket \\ \text{b.} \quad \left\| \begin{array}{c} \alpha \xrightarrow{\text{NARRATION}} \beta \end{array} \right\| = \llbracket \alpha \rrbracket \circ \llbracket \beta \rrbracket \circ \llbracket e_\alpha \prec e_\beta \rrbracket \\ \text{c.} \quad \left\| \begin{array}{c} \alpha \\ \downarrow \text{EXPLANATION} \\ \beta \end{array} \right\| = \llbracket \alpha \rrbracket \circ \llbracket \beta \rrbracket \circ \llbracket \text{cause}(e_\beta, e_\alpha) \rrbracket \end{array}$$

With interpretation rules like these we can extend the standard dynamic semantics of basic DRSs to interpret SDRSs like (17) as context change potentials. In words, the meaning of (17) is that context change potential that tells us to update an input

context with the dynamic content of the first DRS (adding a discourse referent for an eventuality of it getting late), then that of the second DRS (adding two new discourse referents, for an event and an individual), then adding the BACKGROUND requirement (that the first eventuality we just added overlaps with the second), then the content of the third DRS, then the NARRATION requirement (that the event introduced by the second DRS immediately precedes the event introduced by the third), etc..


2.3 Attribution

In the case of (13) the discourse relations were all implicit, the reader had to infer them based on general Glue Logic axioms and world-knowledge. But I also noted above that natural languages provide various ways to make discourse relations explicit. For instance, we treat a connective like ‘because’ in (20) not as a sentential operator (at the level of compositional semantics), but rather as a lexical realization of an instruction to the discourse structure building algorithm (i.e., as a lexically encoded Glue Logic axiom) to label the arrow connecting the two conjuncts of the connective as EXPLANATION.

(20) I got mad because Mia insulted me. She said that I’m a phony.

Following Hunter (2016), Maier (2023) I treat overt reporting constructions like ‘*x* said/thinks/dreamed that *p*’ the same way, viz. as encoding an instruction to connect the two discourse units (‘She said’ and ‘I’m a phony’ in (20)) with a discourse relation of ATTRIBUTION (i.e., clausal embedding report constructions are treated as grammaticalized Glue Logic axioms).⁸

(21) a. π_1 : I got mad π_2 :(because) Mia insulted me.
 π_3 :She said π_4 :(that) I’m a phony.



⁸ While ‘she said’ may not be a grammatical clause by itself in English, our Davidsonian event semantics does readily assign it an interpretable DRS, viz. one that introduces an event and describes it as an event of saying, so nothing prevents us from analyzing it as a separate discourse unit.

ATTRIBUTION is a subordinating relation, drawn vertically. Importantly, the information that is attributed (viz., that the speaker is a phony), should not get added to the discourse context (the speaker is not committing herself to actually being a phony). In SDRT terms this means that ATTRIBUTION is a so-called (right-)non-veridical discourse relation. I propose the following simple semantics, in which $\wedge\beta$ denotes the possible worlds proposition expressed by β (see [Maier 2023](#) for a more comprehensive semantics of ATTRIBUTION that covers also direct and mixed quotation as well as free indirect discourse):⁹

$$(22) \quad \left[\left[\begin{array}{c} \alpha \\ \downarrow \text{ATTRIBUTION} \\ \beta \end{array} \right] \right] = \llbracket \alpha \rrbracket \circ \llbracket \text{content}(e_\alpha) = \wedge\beta \rrbracket$$

This discourse semantic implementation of ATTRIBUTION mirrors Kratzer’s compositional analysis of attitude and speech reports in which a matrix verb introduces contentful eventuality (of speaking, believing, imagining etc), whose propositional contents are then specified by the embedded clause ([Kratzer 2006](#), [Hacquard 2010](#), [Ozyildiz 2021](#)). In fact, for this simple, single sentence example, it really doesn’t matter, semantically speaking, whether we analyze it as consisting of two discourse units connected by ATTRIBUTION (with the semantics in (22)), or more traditionally, as a single discourse unit with a compositional logical form containing either an intensional operator (with Hintikka/Kaplan-style semantics) or a contentful-event modifier (with Kratzer-style semantics). The predicted truth conditions (and/or context change potentials) are exactly the same. The difference, at this point, is architectural, i.e., it concerns where, when, and how the testable truth conditions are derived, viz., at the level of the syntax/semantics interface, or at the level of discourse processing. As we’ll see, pushing the semantics of reporting from compositional semantics into discourse semantics, as proposed here, will have actual, empirical benefits when we look at more complex forms of reporting in discourse, especially in the kind of complex dream report stories we encountered above, where, we argued in section 1, intensional operator approaches falter.

2.4 Reporting with embedding and without

A major selling point of an ATTRIBUTION-based discourse approach to reporting concerns the phenomenon of reports extending beyond the complement, i.e., the cluster of phenomena we tentatively discussed in section 1.2 under the header of

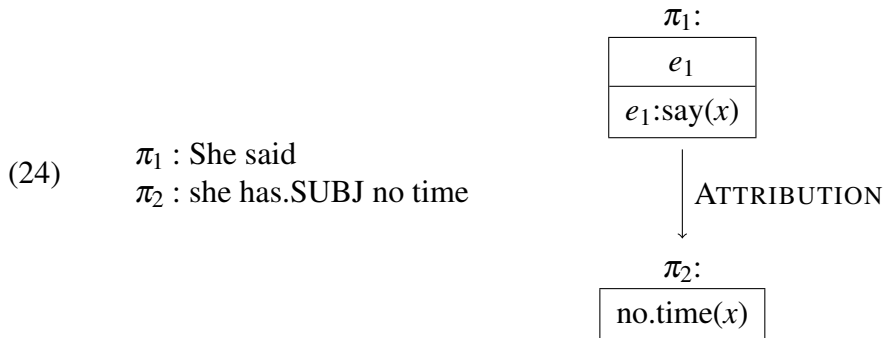
⁹ Above we assumed a dynamic semantics for DRSs, but the possible worlds proposition expressed by a DRS is a static notion, which we can derive from the classical truth-conditional semantics of DRT, or derive from the dynamic interpretations, see [Kamp et al. \(2003\)](#).

unembedded indirect discourse or modal subordination. Here’s a typical example of an ‘extended speech report’:

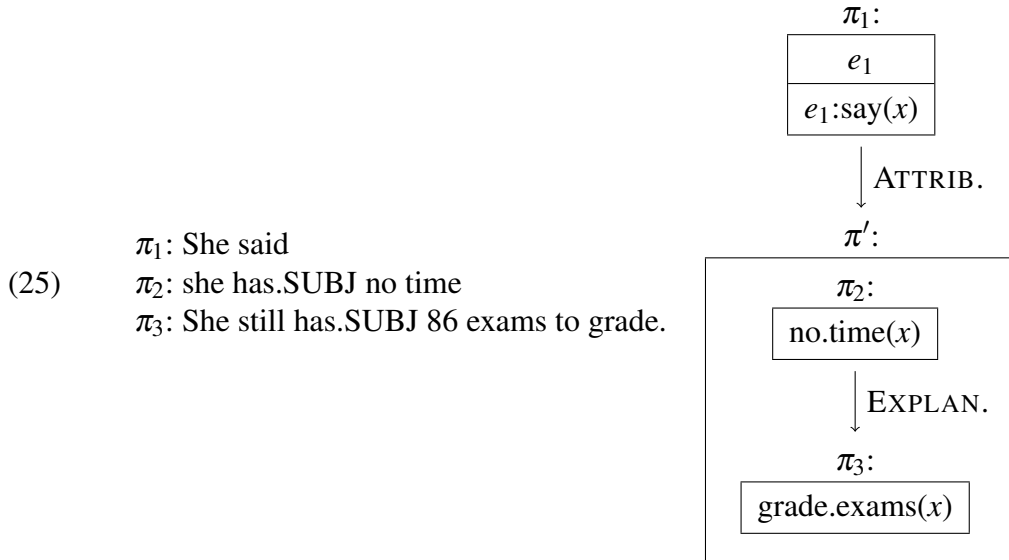
- (23) Sie sagte sie habe keine Zeit. Sie müsse noch 86 Prüfungen bewerten.
 ‘She said she has.SUBJ no time. She still has.SUBJ 86 exams to grade’
 (Bary & Maier 2021)

The verbal inflection glossed SUBJ is known as the reportive subjunctive (Fabricius-Hansen & Sæbø 2004). It occurs in speech report complements but also in free-standing main clauses, in which case it serves to mark the proposition as part of a report. Thus, the second sentence of (23) is interpreted as the subject saying that she still has 86 exams to grade. In SDRT terms, this mood is just another lexicalized marker that tells us (via a lexical Glue Logic axiom, presumably) that the marked verb is part of an attribution (i.e., that the interpretation of that verb should be represented in a box under a downward arrow labeled as ATTRIBUTION).

Let’s look at the processing of this specific example in a bit more detail before returning to dream reports. Given the discourse structure constraints imposed both grammatically (by the indirect speech construction) and lexically (by the subjunctive morpheme), the first sentence gets segmented and unambiguously represented as follows:



Now the second sentence comes in, expressing the proposition that the subject has 86 exams to grade. If we would ignore the mood morpheme, this discourse unit, π_3 , could in principle be connected to the existing discourse graph in (24) in two ways: it might attach to π_1 (via BACKGROUND or perhaps EXPLANATION), or it might connect to π_2 , via EXPLANATION. In the actual German example we have the subjunctive morpheme to disambiguate, enforcing the second discourse construal, shown in (25). A final technical note: connecting multiple units in this way under a non-veridical relation like ATTRIBUTION requires the formation of a so-called complex discourse unit π' that groups these ‘embedded’ units together.



Subsequent discourse units can now be connected either to π_3 (continuing the reported explanation), or to π_2 (ending the reported explanation and continuing the description of the reported sequence of events), or to π_1 (ending the report and continuing the description of eventualities occurring in the actual world).

By merely following the independently motivated mechanisms of coherence maximization and discourse relation inference, we’ve effectively managed to correctly represent the truth conditions of the discourse in (23) as involving a complex report. Anticipating our answer to the challenge raised in section 1.2, note that in so doing we’ve inferred a reported explanation, i.e. we’ve inferred not that the speaker believes grading exams explains the subject’s lack of time, but that according to (the reasoning ascribed to) the subject the grading explains the lack of time.

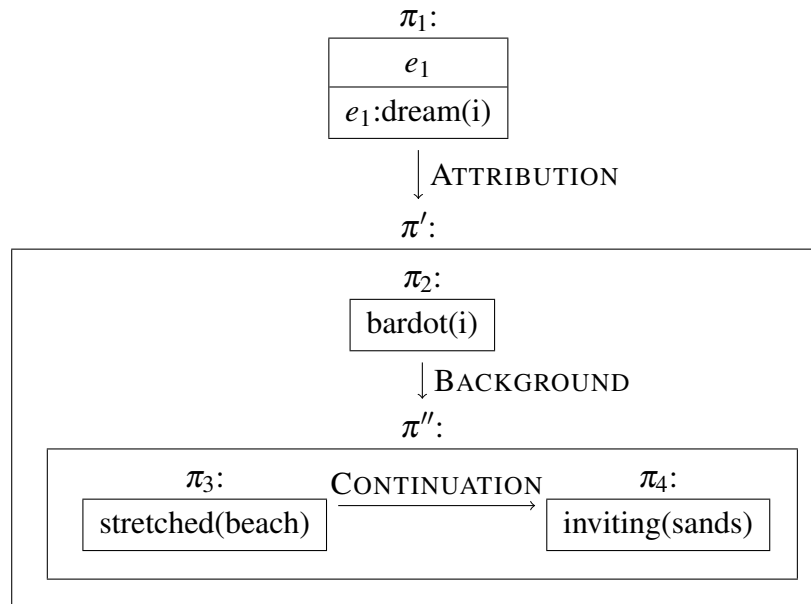
This concludes our exposition of the basic SDRT framework, including a non-veridical event-based semantics of ATTRIBUTION, lexical/grammatical Glue Logic axioms (for causal connectives, subjunctive morphemes, and clausal embedding report constructions), and complex discourse units. This gives us all the tools we need to analyze, and interpret, the types of actual dream reports you find in the dreambank corpus.

2.5 Dream discourse revisited

Let’s go over some examples to see how the general theory of discourse and reporting applies to dream reports, especially those that are arguably difficult to account for with covert operators.

First, the extended Bardot example. The opening lines show the pattern of the extended report discussed above:

- (26) π_1 :I dreamed π_2 :that I was Brigitte Bardot. π_3 :The sun-kissed beaches of the French Riviera stretched out before me, π_4 :their golden sands inviting and warm.

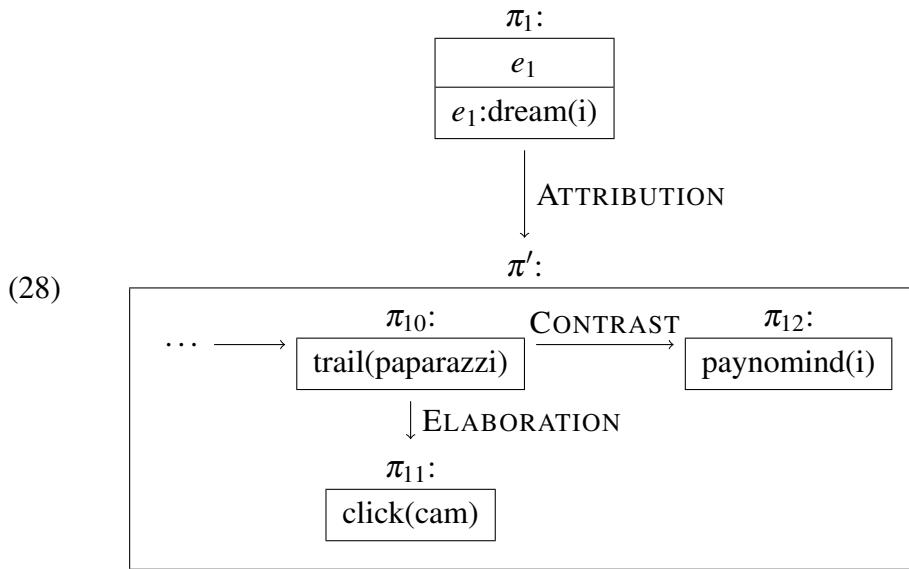


The example illustrates, again, how the description of a dream need not fit into an embedded clause, nor be distributed over a sequence of clauses embedded under (hidden) operators. The second sentence is, compositionally speaking, treated here as just an independent main clause, a separate discourse unit. This unit may, discourse-theoretically speaking, attach to either π_1 , continuing the speaker’s description of what’s happening in the real world (or the ‘base story world’, if the whole thing is a fiction), or π_2 , continuing the description of the dream world. The latter makes much more sense from a global coherence perspective (it’s Bardot who is associated with the sun-kissed beaches of the French Riviera), so for pragmatic reasons we opt for that graph construal, resulting in the structure in (26). The model-theoretic interpretation associated with this graph structure is one in which there is an event of the speaker dreaming, and the content of that dreaming-event is the proposition expressed by the complex SDRS box π' (roughly the proposition that the speaker is Bardot, located at the French Riviera etc.).

Scrolling a bit further through the story we get to the contrastive marker discussed in section 1.2.

- (27) [...] π_{10} :Paparazzi trailed me, π_{11} :their cameras clicking incessantly.
 π_{12} :But I paid them no mind;

As discussed there, it is hard for a compositional semantic analysis in terms of elided dream operators to derive the intended interpretation, where *but* marks a contrast between the paparazzi trailing and the paying no mind (i.e., a contrast within the scope of a dream operator). Here, we get that interpretation for free:



Finally, consider the dreambank example without dream operators, (12):

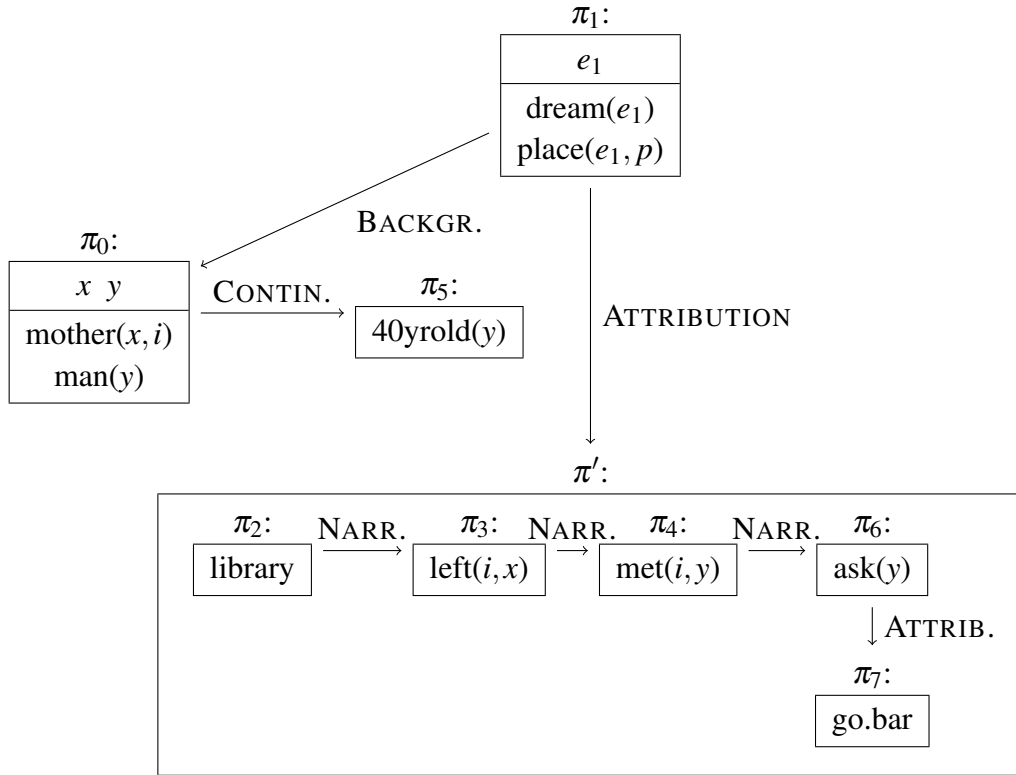
- (29) π_1 :This dream took place in P. π_2 :My mother and I were in the public library together. π_3 :I left her for a few moments π_4 :and met an older man π_5 :(who is about 40 years old [...]). π_6 :He asked me π_7 :to go to a bar π_8 and have a couple of drinks with him

In the first unit there's a nominal rather than verbal description of a dream event, and no sign of clausal embedding at all, which is problematic for a 'modal subordination' approach (see section 1.2). On the current account there's no issue. The first unit introduces a dream, and since dreams are the types of things that can have propositional contents, the following unit can be connected via *ATTRIBUTION*, which gives the intended interpretation where the subject is dreaming that they were in the public library. The other interpretation, where the second sentence connects via some other, veridical discourse relation and hence describes what the subject actually did, is initially available as well, though in the context of the next couple of sentences the first reading is clearly the most coherent.

The next thing to note in (29) is the bracketed remark, π_5 , where the subject seems to give some real-world background facts about the man in her dream. We can model this straightforwardly by just connecting the bracketed remarks to π_1 or perhaps some other discourse unit in the context, outside the dream box. Such a global attachment may be driven solely by global coherence considerations, but presumably the switch from past to present tense and the use of brackets can be modeled as clues that inform this decision.¹⁰

A final complication is the interpretation of the indefinite ‘a man’ that the bracketed relative clause unit is attached to. To get a consistent graph we must interpret this as a specific indefinite, i.e., represented with wide scope, outside the dream (Sæbø 2013). In other words, the subject is recounting a *de re* dream about a certain man whom she was in fact already acquainted with in the real world. Of course, definites like ‘my mother’ are likewise to be construed *de re*. I can’t offer here a detailed analysis of presupposition and anaphora resolution, *de re* attitudes, names, indexicals, and specific indefinites within the framework of SDRT, but the idea is that a specific indefinite, like a name, doesn’t introduce a new discourse referent in its local DRS but rather attaches to a discourse referent already in the context, or else creates one via global accommodation. Concretely, the mother (x) and the man (y) or represented in an accommodated unit π_0 , serving as background to π_1 , so we can connect the bracketed information in π_5 to that unit.

¹⁰ A comprehensive inventory and analysis of such discourse structure clues, as grammatical/prosodic Glue Logic axioms or otherwise, is beyond the scope of this paper.



My proposal, in sum, is that *dream that* is not an intensional operator but an (optionally) overt realization of a non-veridical discourse relation of **ATTRIBUTION**. This helps us make sense of the intuition that dream reporting discourse is a form of storytelling, not about the real-world (or fictional base world), but about a dream-world. Since storytelling is not restricted to verbal, linguistic discourse, as witness comics, picture books, and movies (and perhaps also music and dance), we should expect our discourse structural analysis of dream reporting to carry over to dream reporting in other modalities. Below I show how some common non-linguistic (or at least, non-verbal) conventions for reporting dreams in film and comics can be analyzed by the exact same discourse structural analysis proposed for the linguistic report structures above.

3 Showing dreams

3.1 Discourse across media

Building on recent work in super linguistics and pictorial discourse semantics (Schlöder & Altshuler 2023, Cumming et al. 2017, Maier 2024, Bateman &

Wildfeuer 2014), I assume that any sequence of propositional units can be used to tell stories. Regardless of the medium of expression, storytelling exploits the human tendency for coherence maximization by the inference of discourse relations between eventualities described, depicted, or otherwise represented. As long as we can (i) segment the representational medium into a discrete sequence of units and (ii) associate propositional expressions of eventualities with those individual units, we can use our SDRT framework to study the discourse interpretation process.

A first salient example of a segmentable/propositional representation system evidently suited for storytelling is comics, where the units correspond to panels and propositions are expressed via depiction. Film too uses depiction to express propositions, but here the units correspond to so-called shots, which we can think of as ‘moving pictures’ corresponding to a single run of a camera. The table below summarizes the important differences and similarities between some different representational media.

	speech	writing	comics	film
primary modality	auditive	visual	visual	audiovisual
discourse unit	utterance	sentence	panel	shot
sequencing	temporal	linear	2D grid	temporal
unit semantics	compositional	compositional	pictorial	pictorial


In all these representation systems, interpreters are tasked with extracting coherent stories out of sequences of propositional units, and it is this process that frameworks like SDRT are supposed to model. Since we know that comics and film are expressive enough to convey stories featuring dreams, the question arises if indeed our *ATTRIBUTE*-based semantics of dream reporting can be used to investigate dream reporting strategies and conventions that have emerged in these media.

3.2 PicDRT and comics panels

Our presentation of SDRT used DRSs to represent the dynamic contents of elementary discourse units. We’ll follow Maier & Bimpikou’s (2019) (in turn inspired by Abusch 2012 and Greenberg 2013, among others) extension of DRT to represent the dynamic contents of pictorial units, i.e., the panels of comics.

In PicDRT a picture is semantically represented as a picture condition in a DRS box. In the interpretation process, first, a number of salient regions are identified and labeled with fresh discourse referents. In the case of the drawing in (30) we have two such regions, which we’ll label x_1 and y_1 . These represent the two entities that the picture is about. In addition, we create a discourse referent to represent the inferred spatio-temporal viewpoint, v_1 from which the picture is, supposedly, taken.

The content of the picture as a whole is thus represented as a DRS box introducing three new discourse referents:¹¹



The PicDRS in (30) is supposed to mean, roughly, that there are two individuals and a viewpoint, and the world looks like that picture when viewed from that viewpoint, and moreover, that the two individuals look like the two labeled regions (when viewed from that same viewpoint). We can integrate this into a proper model-theoretic DRT semantics by defining ‘looks like’ in terms of a formal picture semantics via geometric projection (Greenberg 2013, Abusch 2020). A projection is a certain type of function mapping 3D scenes into 2D depictions. The most commonly used type of projection is linear perspective, which corresponds to the familiar drawing techniques in which nearby objects appear larger and parallel lines converge on the horizon. Given a projection function Π , a world (or object in a world) ‘looks like’ a given picture (from some viewpoint) iff applying the projection function to that world (or object) and viewpoint as inputs gives us precisely that picture as output.

Summing up: in a dynamic semantic setting a single panel of a comic strip can be semantically represented in the form of a PicDRS (i.e., with a picture condition that contains a copy of the picture itself). A PicDRS contributes discourse referents for the objects that are saliently depicted in (or contextually inferred from) the picture, and says that the world (as seen from that viewpoint) looks like that picture (or ore precisely, projects onto that picture).

3.3 SDRT and comics

Maier & Bimpikou (2019) use standard (Pic)DRT to model the dynamics of interpretation for sequences of multiple panels. Here we use SDRT instead, because we’re interested in modeling the structure of (dream reporting) discourse.

Consider a two-panel sequence:




We combine basic SDRT with the idea that panels correspond to elementary discourse

¹¹ Police and squirrel images taken from Maier & Bimpikou 2019.

units, semantically represented as PicDRSs. That means we have here two segments, π_1 and π_2 each depicting what the world looks like from some viewpoint. Coherence maximization tells us to connect these world depictions with a discourse relation (now without the help of lexicalized or grammaticalized Glue Logic axioms for tense, mood, aspect, causal connectives etc).

The default relation for the kind of storytelling commonly associated with comics is NARRATION: first the world looks like this *and then* it looks like that, but we'll see examples of other relations, specifically *ATTRIBUTION*, below. Recall from section 2.2 that the semantics of NARRATION and many other discourse relations requires that we identify the main eventuality introduced by a discourse unit. There's an ongoing debate about whether and what kind of eventualities can be depicted in (or reliably inferred from) a picture, but following Maier (2024) I hold that panels like the above introduce discourse referents for events, i.e., the first depicts an event of a policeman running, rather than a state of a policeman floating in mid-air with stretched out legs.

Since NARRATION is often assumed to demand spatio-temporal proximity between the linked eventualities, maximal coherence entails that the policeman depicted in the first is (probably) the same policemen as that depicted in the second, i.e., we equate discourse referents introduced by consecutive pictures whenever we coherently and plausibly can. Discourse referent equation is a form of pragmatic strengthening, on a par with further inferences we might represent in the DRSs, like that x_1 is a policemen and e_1 is a event of x_1 chasing y_1 . For readability I will not explicitly add such inferred conditions (like 'policeman(x_1)') to the PicDRSs.



Film, we said, consists of shots that are stitched together in deliberate temporal sequence to tell coherent stories. We may think of shots as, literally, moving pictures, so the DRSs representing the elementary discourse units will contain such moving pictures, along with discourse referents for salient (moving) picture regions, (moving) viewpoints, and events. We won't go into the details of formalizing the mathematical and conceptual complications posed by addition movement, just like we'll ignore audio and music. Instead we'll simply represent film shots as if they are single pictures in a wordless comics strip.


3.4 Visual dream attribution

Many genres of film and comics rely heavily on NARRATION, adjacent panels/shots show consecutive events. Occasionally we find ‘wandering eye’, stative scenery shots (BACKGROUND), or a FLASHBACK, or PARALLEL actions montage (McCloud 1993). More to the point, sometimes we are shown not what happens in the base story-world, but what some character subjectively experiences. Modeling such perspective shifts in the visual domain has led some researchers to invoke hidden intensional operators again (Abusch & Rooth 2022), but in the current approach we can do without; we just infer our non-veridical discourse relation of ATTRIBUTION.

Focusing on dream representations in visual media we can identify two main types, dream bubbles and dream sequences. We’ll examine both visual dream report strategies and see that they mirror the explicit operator embedding and the extended unembedded reporting strategies that we’ve encountered in linguistic reporting discourse in section 2. Following our uniform analysis of linguistic dream reporting I propose to model both in terms of ATTRIBUTION again.

3.5 Dream bubbles

Bubbles are mostly used in comics, and then mostly for linguistic thought reports, but occasionally for dreams as well. It involves showing both the actual dreamer and their dream in the same panel, typically with a bubble shape in (33a) to mark the perspective shift.¹²



Predating the modern comics convention of bubbles and balloons we find paintings like (33b) exemplifying the same strategy: we see St. Joseph, asleep in the base world, and we see the angel that visits him in his dream. In both examples we have a single image that combines depictions of two distinct worlds.¹³

¹² Garfield comics 9-9-1987 by Jim Davis. Painting *Il sogno di San Giuseppe* by Antonio Ciseri, 1850.

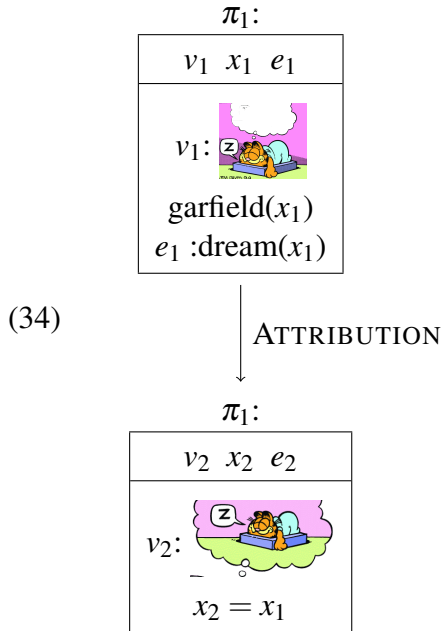
¹³ This kind of representation is used in theater as well, and as in the painting of St. Joseph above, lighting can be a way of separating the sleeping dreamer’s world from the content of their dream. Although such representations blend real world and dream world they are not what Maier & Bimpikou

Maier (2019) analyzes speech and thought bubbles containing words as a form of quotation. The account sits firmly in the syntax/semantics interface, decomposing the picture into a picture proper (interpreted iconically, via geometric projection), a bubble (interpreted symbolically, as a quotation operator), and a string of letters (constituting the argument of the quotation operator). The bubble is quite literally a morphologically realized operator in the logical form of the picture.¹⁴ In the dream report in (33a), the content of the bubble is a picture rather than text, but we could, in principle, try to extend the compositional operator-based analysis. Here, instead, I will take an alternative, discourse-level approach. The reason is entirely parallel to the reason why I opted for a discourse approach for linguistic reports: complex dream representations that extend over various discourse units and involve internal discourse structure, as we'll see in section 3.6 below.

The discourse approach for reports straightforwardly extends to dream bubbles. We simply treat the bubble not as an intensional dream operator at LF but via a grammaticalized Glue Logic axiom forcing the content of the bubble picture to be represented under an *ATTRIBUTION*. Applied to the Garfield dream this means we segment it as consisting of two separate pictorial discourse units whose contents are linked via *ATTRIBUTION*:

(2019) call 'blended perspective shots'. That term is reserved for the combination of an objective, third person viewpoint with a depiction of a subjective (dreamed or imagined) world. Some of the examples in section 3.6 do involve such blended perspective.

¹⁴ Maier (2023) argues that we can, and should, extend our *ATTRIBUTION* analysis of indirect discourse to quotation in natural language, so then even an analysis of thought bubbles as quotation will ultimately lead us to a discourse-based analysis involving *ATTRIBUTION*.



The model-theoretic interpretation rule for ATTRIBUTION demands (i) that the top unit introduces a potentially contentful event, and (ii) that the bottom unit can be assigned a possible worlds proposition. Both demands are met in (34). First, π_1 evidently depicts a sleeping Garfield so we can indeed easily accommodate that he is dreaming, which is contentful. Then, π_2 is a picture, represented as a PicDRS, which has truth conditions and hence can be assigned a propositional content: $\hat{K}\pi_2 \approx$ the set of worlds w for which we can find a viewpoint from where w looks like that.


As in the linguistic case, for simple reports with a single dream proposition represented inside another discourse unit, not much hinges on the architectural choice between a compositional analysis ‘a la Maier (2019) or an ATTRIBUTION approach as outlined above. But for multi-panel ‘dream sequences’ we really need a discourse approach.

3.6 Dream sequences

In both film and comics, we often find dream sequences that extend beyond a single bubble inside a single panel. As a direct comics parallel to our extended Bardot example, we find explicit dream bubbles, followed by multiple separate discourse units (i.e., panels) interpreted as representations of the same dream.¹⁵

¹⁵ Donald Duck, *Perchance to dream*, by William Van Horn, 2000.

(35)



In fact, in (35) we see traces of the bubble morphology also in the extended dream panels, suggesting an analysis of bubbly edges as the visual analogue of a reportative mood (as in the German extended speech report example in (23)).

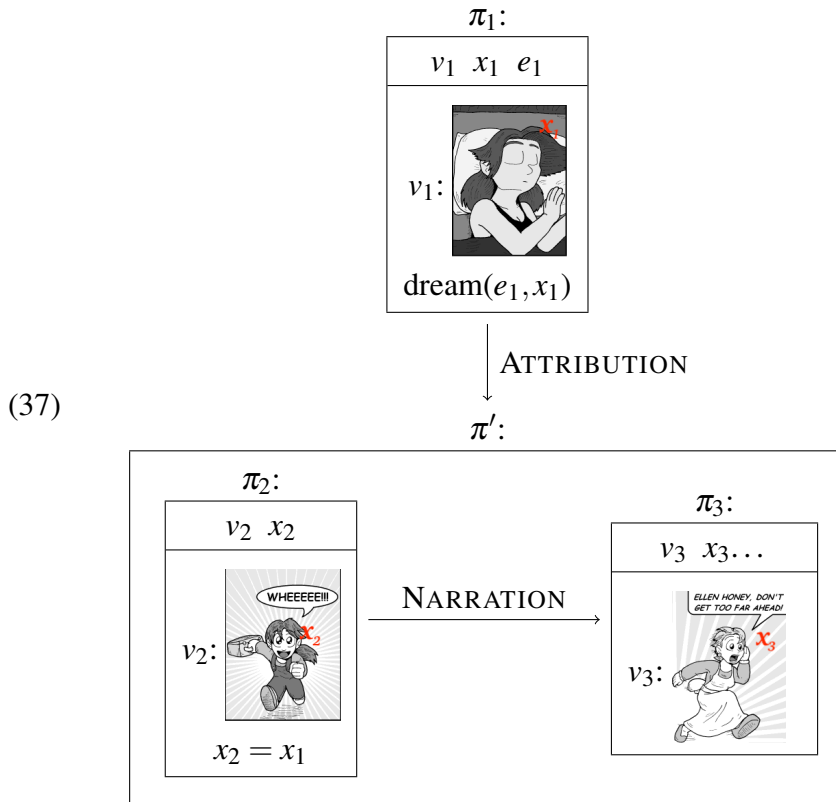
In section 2.5 we also saw that complex dream reporting need not involve any embedding or mood markers, and can be driven wholly pragmatically, as long as we are told, or can plausibly accommodate, that there's an event of dreaming (see example (12) about the library). The same happens in comics. In (36) we see the protagonist asleep in her bed, followed by a sequence of panels depicting scenes with a much younger version of that protagonist and her mother.¹⁶

(36)



There's no bubbles or other morphological marking that we might analyze as the realization of an operator at LF taking scope over the dream images. We gather from context that we're viewing the protagonist's dream here. The resulting discourse structure for the first three panels comes out something like this (with the contextually inferred information that the event depicted in π_1 is an event of dreaming, and that x_2 is the same individual as x_1 represented as pragmatic enrichments in the corresponding DRSs):


¹⁶ 'Second life: First day of school', *El Goonish Shive*, by Dan Shive, 2004.



In film, where we don't (usually) have bubbles, dreams are likewise presented via sequences of shots that form a dream story. In (38), we see a shot zooming in on the sleeping dreamer, Sheldon, followed by a sequence of shots depicting the dream world (featuring multiple version of Sheldon, sitting in sci-fi chairs in a strange white world). The cut between the base story shots and the dream story shots is marked by a blue swirl transition effect.¹⁷



¹⁷ 'The relaxation integration', *The Big Bang Theory*, season 11, episode 3, 2017.



Comics artists and filmmakers may leave various cues to help the viewer infer that there's a perspective shift, including a visual or musical transition effect between shots (like the blue swirl in (38)), or distinct music, lighting, backgrounds, or filters during the dream sequence. We may think of some of these markers as filmic parallels of linguistic cues like reportative mood, or the bubble morphology of comics. A complete inventory and classification of audiovisual morphemes and other multi-modal discourse structure markers, stylistic cues, and genre conventions (and their linguistic analogues) is beyond the scope of this paper.

In some cases, the artist or filmmaker chooses to leave the dream event and/or the scope of the *ATTRIBUTION* entirely unmarked so that the viewer has to arrive at the dream interpretation solely via contextual reasoning (in the service of maximizing global coherence). Playing with a dream/reality ambiguity, leaving audiences in the dark until a final reveal, is a tried and tested storytelling technique, in visual as well as verbal storytelling, and the current discourse approach is well-suited to modeling the interpretation and re-interpretation processes involved (see [Altshuler & Kim 2023](#) on 're-analysis' in an SDRT framework).

4 Conclusion

There's interesting unresolved syntactic, semantic, and empirical issues about pronoun interpretation in '*x* dreamed that *p*' sentences. But dreams are not usually reported with a single embedded clause. When we look at what's going on in actual dream reporting discourse, the traditional analysis in terms of intensional operators and clausal embedding falters.

I have argued for an alternative analysis where 'dream that' is an overt realization of a non-veridical discourse relation. While this move makes little difference for the prediction of truth-conditions of single sentence dream reports, it gives us a handle on a range of complex and implicit dream reporting strategies that we find in the dream reports from the DreamBank corpus. Conceptually, my proposal follows the intuition that dream reporting is a form of storytelling, where the reported dream can be represented by multiple discourse units that form a coherent but potentially complex discourse structure.

Since telling complex but coherent stories is not restricted to verbal linguistic discourse, I used my formal discourse semantics framework to look for analogous

dream reporting strategies in other media. We saw how dream bubbles and dream sequences in comics and film exactly mirror ‘dream that’ complements and the extended main clause reporting strategies from verbal dream discourse, and can hence be given a uniform analysis in terms of ATTRIBUTION.

References

- Abusch, Dorit. 2012. Applying Discourse Semantics and Pragmatics to Co-reference in Picture Sequences. *Proceedings of Sinn und Bedeutung* 17. <http://ecommons.cornell.edu/handle/1813/30598>.
- Abusch, Dorit. 2020. Possible-Worlds Semantics for Pictures. In Daniel Gutzmann, Lisa Matthewson, Cécile Meier, Hotze Rullmann & Thomas Zimmermann (eds.), *The Wiley Blackwell Companion to Semantics*, 1–31. Wiley. <https://doi.org/10.1002/9781118788516.sem003>.
- Abusch, Dorit & Mats Rooth. 2022. Pictorial free perception. *Linguistics and Philosophy* <http://dx.doi.org/10.1007/s10988-022-09370-1>.
- Altshuler, Daniel & Christina S. Kim. 2023. Towards a Typology of Narrative Frustration. *Topoi* <http://dx.doi.org/10.1007/s11245-023-09942-1>.
- Anand, Pranav. 2006. *De De Se*: MIT dissertation. <http://hdl.handle.net/1721.1/37418>.
- Asher, Nicholas & Alex Lascarides. 2003. *Logics of Conversation*. Cambridge: Cambridge University Press.
- Banfield, Ann. 1982. *Unspeakable Sentences: Narration and Representation in the Language of Fiction*. London: Routledge & Kegan Paul.
- Bary, Corien & Emar Maier. 2014. Unembedded Indirect Discourse. *Proceedings of Sinn und Bedeutung* 18. 77–94. <http://semanticsarchive.net/sub2013/SeparateArticles/Bary&Maier.pdf>.
- Bary, Corien & Emar Maier. 2021. The landscape of speech reporting. *Semantics and Pragmatics* 14(8). <http://dx.doi.org/10.3765/sp.14.8>.
- Bateman, John & Janina Wildfeuer. 2014. A Multimodal Discourse Theory of Visual Narrative. *Journal of Pragmatics* 74. 180–208. <http://dx.doi.org/10.1016/j.pragma.2014.10.001>.
- Charnavel, Isabelle. 2024. Free Indirect Discourse as logophoric context. <https://lingbuzz.net/lingbuzz/007439>.
- Clements, George. 1975. The logophoric pronoun in Ewe: Its role in discourse. *Journal of West African Languages* 10. 141–177.
- Cumming, Samuel, Gabriel Greenberg & Rory Kelly. 2017. Conventions of Viewpoint Coherence in Film. *Philosophers' Imprint* 17(1). 1–29. <http://hdl.handle.net/2027/spo.3521354.0017.001>.

- Eckardt, Regine. 2014. *The Semantics of Free Indirect Speech. How Texts Let You Read Minds and Eavesdrop*. Leiden: Brill.
- Fabricius-Hansen, Cathrine & Kjell Johan Sæbø. 2004. In a mediative mood: The semantics of the German reportive subjunctive. *Natural Language Semantics* 12(3). 213–257. <http://link.springer.com/article/10.1023/B:NALS.0000034514.27887.d9>.
- Fludernik, Monika. 1995. *The Fictions of Language and the Languages of Fiction: The Linguistic Representation of Speech and Consciousness*. London: Routledge.
- Frege, Gottlob. 1892. Ueber Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* 100(1). 25–50.
- Greenberg, Gabriel. 2013. Beyond Resemblance. *Philosophical Review* 122(2). 215–287. <http://dx.doi.org/10.1215/00318108-1963716>.
- Hacquard, Valentine. 2010. On the event relativity of modal auxiliaries. *Natural Language Semantics* 18(1). 79–114. <http://dx.doi.org/10.1007/s11050-010-9056-4>.
- Hintikka, Jaakko. 1969. Semantics for propositional attitudes. In *Models for Modalities*, 87–112. Dordrecht: Reidel.
- Hunter, Julie. 2016. Reports in Discourse. *Dialogue & Discourse* 7(4). <http://dad.uni-bielefeld.de/index.php/dad/article/view/3695>.
- Kamp, Hans. 1981. A theory of truth and semantic representation. In Jeroen Groenendijk, Theo Janssen & Martin Stokhof (eds.), *Formal Methods in the Study of Language*, 277–322. Amsterdam: Mathematical Centre Tracts.
- Kamp, Hans, Josef van Genabith & Uwe Reyle. 2003. Discourse Representation Theory. In Dov Gabbay & Franz Guenther (eds.), *Handbook of Philosophical Logic*, vol. 10, 125–394. Heidelberg: Springer.
- Kaplan, David. 1989. Demonstratives. In Joseph Almog, John Perry & Howard Wettstein (eds.), *Themes from Kaplan*, 481–614. New York: Oxford University Press.
- Kratzer, Angelika. 2006. Decomposing attitude verbs Handout. Honoring Anita Mittwoch on her 80th birthday. The Hebrew University of Jerusalem. <https://semanticsarchive.net/Archive/DcwY2JkM/>.
- Lakoff, George. 1968. Counterparts, or the problem of reference in transformational grammar <https://escholarship.org/content/qt28b7k6sb/qt28b7k6sb.pdf>.
- Maier, Emar. 2015. Quotation and unquotation in free indirect discourse. *Mind and Language* 30(3). 235–273. <http://dx.doi.org/10.1111/mila.12083>.
- Maier, Emar. 2019. Picturing words: The semantics of speech balloons. *Proceedings of the Amsterdam Colloquium* 22. 584–592. <https://philpapers.org/rec/MAIPWT>.
- Maier, Emar. 2023. Attribution and the discourse structure of reports. *Dialogue & Discourse* 14(1). 34–55. <http://dx.doi.org/10.5210/dad.2023.102>.
- Maier, Emar. 2024. Pictorial language and linguistics. <https://lingbuzz.net/lingbuzz/>

- 007603.
- Maier, Emar & Sofia Bimpikou. 2019. Shifting perspectives in pictorial narratives. *Sinn und Bedeutung* 23(2). 91–106. <http://dx.doi.org/10.18148/sub/2019.v23i2.600>.
- McCloud, Scott. 1993. *Understanding Comics*. Tundra Publishing.
- Ninan, Dilip. 2008. *Imagination, Content, and the Self*: MIT dissertation.
- Ozyildiz, Deniz. 2021. *The Event Structure of Attitudes*. Amherst: UMass dissertation. <https://hdl.handle.net/20.500.14394/18513>.
- Pearson, Hazel. 2015. The interpretation of the logophoric pronoun in Ewe. *Natural Language Semantics* 23. 77–118. <http://dx.doi.org/10.1007/s11050-015-9112-1>.
- Pearson, Hazel & Jeruen Dery. 2013. Dreaming de re and de se: Experimental evidence for the Oneiric Reference Constraint. *Sinn und Bedeutung* 16. 322–339.
- Percus, Orin & Uli Sauerland. 2003. Pronoun Movement in Dream Reports. In Makoto Kadowaki & Shigeto Kawahara (eds.), *Proceedings of NELS 33*, Amherst, MA: GLSA.
- Quine, Willard V O. 1956. Quantifiers and propositional attitudes. *Journal of Philosophy* 53. 101–111.
- Roberts, Craige. 1987. *Modal subordination, anaphora, and distributivity*. Amherst: UMass dissertation. <http://scholarworks.umass.edu/dissertations/AAI8710499/>.
- Roberts, Craige. 1989. Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy* 12(6). 683–721. <http://dx.doi.org/10.1007/BF00632602>.
- Roberts, Craige. 2020. Modal Subordination. In *The Wiley Blackwell Companion to Semantics*, 1–36. John Wiley & Sons, Ltd. <http://dx.doi.org/10.1002/9781118788516.sem102>.
- Sæbø, Kjell Johan. 2013. Reports of Specific Indefinites. *Journal of Semantics* 30(3). 267–314. <http://dx.doi.org/10.1093/jos/ffs015>.
- Schlenker, Philippe. 2003. A plea for monsters. *Linguistics and Philosophy* 26(1). 29–120. <http://dx.doi.org/10.1023/A:102225203544>.
- Schlenker, Philippe. 2004. Context of thought and context of utterance: A note on Free Indirect Discourse and the Historical Present. *Mind and Language* 19(3). 279–304. <http://dx.doi.org/10.1111/j.1468-0017.2004.00259.x>.
- Schlöder, Julian J. & Daniel Altshuler. 2023. Super Pragmatics of (linguistic-)pictorial discourse. *Linguistics and Philosophy* <http://dx.doi.org/10.1007/s10988-022-09374-x>.
- Schneider, Adam & Bill Domhoff. 2024. DreamBank. <https://dreambank.net/>.
- Sharvit, Yael. 2008. The puzzle of free indirect discourse. *Linguistics and Philosophy* 31(3). 353–395. <http://dx.doi.org/10.1007/s10988-008-9039-9>.
- von Stechow, Arnim. 1982. Structured propositions .