

Metaphor Identification beyond Discourse Coherence

Inés Crespo

NYU Paris

Andreas Heise

Institut Jean Nicod (CNRS, EHESS, ENS-PSL)

Claudia Picazo

UNED

Abstract

In this paper, we propose an account of metaphor identification on the basis of contextual coherence. In doing so, we build on previous work by Nicholas Asher and Alex Lascarides that appeals to rhetorical relations in order to explain discourse structure and the constraints on the interpretation of metaphor that follow from it. Applying this general idea to our problem, we will show that rhetorical relations are sometimes insufficient and sometimes inadequate for deciding whether a given utterance is a case of metaphor. They are insufficient, since rhetorical relations fall short at times of providing a basis for disambiguating between literal and metaphorical interpretations. In such cases, contextual information other than previous discourse needs to enter the picture. To this effect, we bring the idea of external consistency into play. Beyond that, though, we will argue that rhetorical relations are sometimes inadequate to account for coherence, if conceived as relations among sentences only. The reason is that extra-linguistic elements of the situation in which the sentence is uttered may be crucial for getting at the preferred interpretation. To account for these cases, we allow rhetorical relations to connect both with previous discourse and with extra-linguistic situations. In our final refinement of the notion of contextual coherence, we forfeit any appeal to rhetorical relations in favour of Questions Under Discussion (QUD). We defend the view that this account does not only explain the same sort of cases. What is more, it solves the issue of metaphor identification in impoverished contexts.

Keywords: Metaphor identification, Context, Coherence, Rhetorical relations, Questions under discussion (QUD).

1. Introduction

What clues can interpreters rely on in deciding whether a sentence uttered in a specific context had better be interpreted in a metaphorical sense, rather than in a literal sense? Our aim here is to address this question about metaphor identification. When

faced with an utterance of a sentence, one of the things the interpreter might need to determine is whether it is best understood as a metaphor or as a literal claim. This choice is particularly difficult when it comes to sentences that do not involve any category mistake but that are nonetheless reasonably interpreted in a metaphorical sense—so-called twice-true metaphors¹—as illustrated by Disraeli’s utterance of (1):

(1) I have climbed to the top of that greasy pole.

We propose that considerations about coherence play a crucial role. Following Asher and Lascarides, we hold that coherence governs interpretation and that discourse structure constrains metaphorical interpretation.² However, we go beyond their approach and focus on how discourse structure often needs to be supplemented with extra-linguistic information available to the interpreter, without being properly construed as part of previous discourse.

We make two claims. First, although discourse structure constrains metaphorical interpretation with regard to previous discourse, as Asher and Lascarides show, we argue that rhetorical relations are ultimately insufficient to identify metaphor. We think that background information, comprising what Asher and Lascarides call world knowledge (like knowing that a political career for someone normally starts at the level of clerk and might culminate in becoming Prime Minister) and perceptual information also play a role in deciding whether a sentence should be interpreted as a metaphor or as a literal claim. Second, we argue that rhetorical relations, understood as relations among sentences, are sometimes not just insufficient but simply inadequate to account for metaphor identification via contextual coherence. We go beyond Asher and Lascarides’s view to accommodate examples where an utterance is identified as a metaphor, not by virtue of discourse relations established with previous discourse, but by how it stands in relation to perceptual information or background knowledge. In some cases, extra-linguistic elements of the situation in which the sentence is uttered can play a role analogous to that of the linguistic context. Our view integrates Stone et al.’s (2013) account of the role perceptual information plays in making an utterance coherent. This is to say that contextual coherence is sometimes determined by discourse relations established between an utterance and perceptually accessible features of a situation. We also consider cases where metaphor identification takes place in impoverished contexts, where we know relatively little aside from the topic under discussion. We finish with a suggestion about how Questions Under Discussion (QUD) could be used in metaphor identification, given that QUD are a way of checking whether an utterance addresses the topic under discussion.

We thus unfold and refine a notion of contextual coherence based on Asher and Lascarides’s work, but broader than the notion of discourse coherence they develop. This allows us to account for metaphor identification in cases they would not be able to account for.

2. Metaphor Identification and Interpretive Ambiguity

As mentioned, our focus in this paper is on how to tell whether a particular sentence, uttered in a specific context, is an instance of metaphor. We call this the

¹ This expression was coined by Cohen (1976) so as to draw attention to a range of cases that serve as counterexamples to deviance theories of metaphor identification, which characteristically relied on category mistakes as a cue.

² See for instance Asher and Lascarides 1995, 2001 and 2003.

problem of metaphor identification.³ Stern notes that identifying a sentence as a metaphor is not the same as interpreting a metaphor (Stern 2000: 7). One might have good reasons for believing that an utterance is best interpreted as a metaphor without knowing what precisely it says. For instance, we may suspect that an utterance of (2) is metaphorical:⁴

(2) Sam is a pebble.

At the same time, we may fail to draw the relevant comparison between Sam and a pebble. And, conversely, we may have a pretty good idea of what a certain sentence would say if it were to be taken metaphorically, while being ignorant about whether a given utterance of that sentence is best construed as a case of metaphor. This could happen with example (1).

In order to understand an utterance, interpreters typically have to do a number of things, such as determining the proposition expressed, which includes resolving syntactic and lexical ambiguities, and determining the referents of indexicals and context-sensitive expressions; identifying the speech act performed; or inferring implicatures. In some cases, the interpreter must also decide whether the utterance is to be understood as literal or not. Take again example (2), “Sam is a pebble”. This sentence can be used as a metaphor, but also as a metonymy (if Sam is a child dressed as a pebble, in the context of a school play) and perhaps even as a literal claim (in a fictional work as a cartoon, or a surrealist poem, for instance).

In this sense, the situation is similar to that of lexical ambiguity. In a case of *lexical* ambiguity (homonymy), the interpreter has to decide, using contextual information, which of a list of potential meanings is the intended or otherwise correct one. For instance, if faced with an occurrence of the word “bank”, the interpreter has to find out whether the utterance is about a financial institution or the side of a river. Similarly, when faced with an occurrence of “Sam is a pebble”, the interpreter has to find out which of different possible types of interpretation is most appropriate for that sentence, given the context in which it was uttered: Is the utterance to be understood as a metaphor, a metonymy, or a literal claim? In what follows we restrict the discussion to the distinction metaphorical versus literal interpretation, and proceed as if the interpreter only had to choose between these two options of resolving what we might call *interpretive* ambiguity.

It might seem that, in the case of (2), the figurative ambiguity between literal and metaphorical interpretation is very easy to resolve. Assuming that Sam is the name of a person, the sentence would be semantically deviant, and thus either

³ Stern (2000: 3) officially calls it the question of metaphor “recognition”, while describing it as addressing “the conditions, heuristics, clues, cues, trains of reasoning, or steps followed by speaker-hearers by which they identify or recognize particular utterances as metaphors”. We find it more intuitive to speak of metaphor identification here, thus adopting for our own theoretical purposes a terminology that is established in cognitive linguistics (cf. e.g. Pragglejaz Group 2007).

⁴ This example appears in Asher and Lascarides 2001. It is an example of a so-called nominal metaphor (of the form ‘A is B’). Metaphors come, of course, in a host of other syntactic shapes. This is noteworthy insofar as syntactic structure exerts distinctive constraints on intrasentential coherence. There is an interesting literature on the procedures, such as coercion (cf. e.g. Pustejovsky 1995, Asher 2011), by which intrasentential coherence is established when the specifications of semantic types conflict, as in category mistakes or metaphor (cf. Magidor 2020). Our focus, however, is on intersentential or contextual coherence. As Prandi (2021: 64) notes, the latter wins out in case the demands of the two types of coherence conflict.

meaningless or false. The interpreter could use a Gricean-style reasoning to conclude that the speaker must mean her words in a metaphorical sense. However, it is important to note that many metaphors are not semantically deviant. Contrary to what deviance theorists might have thought, no internal feature of the sentence is a reliable guide to metaphoricity, and the interpreter has to make a choice using the available contextual information. Consider a little story surrounding example (1):

After being appointed Prime Minister, Disraeli said: “I have always despised politics. But I have climbed to the top of that greasy pole”.

Example (1) does not involve any category mistake, yet it can be identified as a metaphor. Such an utterance does not have internal features that can be used as a reliable mark of metaphoricity. Our proposal, which we flesh out in sections 4 and 5, is that the key to the identification problem might be in the structure of the discourse in which the target sentence is embedded—as Asher and Lascarides hold—together with extra-linguistic elements of context such as perceptual information and background knowledge available to the interpreter. Before putting it forward, we review Asher and Lascarides’s view and argue that it needs to be supplemented with extra-linguistic information.

3. Asher and Lascarides’s View on Metaphor

In their paper “Metaphor in Discourse”, Asher and Lascarides (2001) argue that it is possible to specify the principles of metaphorical interpretation. In their view, lexical rules and discourse structure—in the form of rhetorical relations—constrain metaphorical interpretation. Interestingly, their account works for metaphors involving category mistakes and for those that do not bear this mark. In this section, we will briefly present their view, by assessing its merits in regards to the question of metaphor identification.

Asher and Lascarides have put forward a theory of interpretation, the Segmented Discourse Representation Theory (SDRT),⁵ in which the principle governing interpretation is *Maximize Discourse Coherence*. The underlying idea is that utterances in discourse are connected to one another. The principle establishes that we should prefer interpretations that maximize coherence—roughly, those allowing for as many connections as possible. The connections at stake are called rhetorical relations, and include Elaboration, Narration, Contrast, Question-Answer-Pair, etc. Rhetorical relations describe rhetorical roles utterances play in discourse context (Asher and Lascarides 2003: 3), for instance, constituting an explanation of a previous utterance. What is of interest for us here is that rhetorical relations can be used to explain constraints on interpretation. In SDRT, rhetorical relations are used to model the semantics/pragmatics interface. According to this theory, the logical form of discourse contains rhetorical relations, and these rhetorical relations can have truth-conditional effects on the sentences they link. Consider the following example, in which “bar” is ambiguous and the interpreter must find out which sense is correct (Asher and Lascarides 1995):

- (3) a. The judge asked where the defendant was.
 b. The barrister said he was in the pub drinking.
 c. The bailiff found him slumped beneath the bar.

⁵ See Asher and Lascarides 1995 and 2003.

How does discourse structure constrain interpretation—by resolving a lexical ambiguity in this case—and what role does coherence play? Asher and Lascarides’s account goes roughly as follows. The discourse structure in (3) is narration. Utterance (3b) is connected to (3a) by the relation Narration, and so is (3c) to (3b). This relation imposes spatio-temporal constraints in the events described. The narrative links are tighter when the interpretation of the sentence—here, utterance (3c)—complies with expectations created. In this example, the expectation would be that the situation described in (3c) is set at the place introduced in (3b)—namely, the pub. “Bar” is disambiguated accordingly. Hence, Asher and Lascarides hold that discourse structure influences ambiguity resolution.

In their 2001 paper, they apply this framework to metaphor and identify two types of constraints on interpretation. First, lexical rules predict that some word occurrences take metaphorical meaning. For instance, there is a rule that establishes that “rock” can take an argument of the type *human* and that, when this happens, the original physical object meaning applies in a metaphorical sense to the relevant human, as happens in “John is a rock”.⁶ Second, rhetorical relations can trigger metaphorical interpretations. One of the examples discussed involves (1), now slightly expanded by being preceded by the sentence (1*a), which couches the original example in a little story:⁷

- (1*) a. I have always despised politics.
 b. But I have climbed to the top of that greasy pole.

“But” in (1*b) signals the rhetorical relation of Contrast. This relation connects (1*b) with (1*a), and ensures that they match at the structural and semantic level. Since the two sentences are thus connected, the referent of the anaphoric expression “that greasy pole” is provided by the first sentence: “that greasy pole” refers to politics. Moreover, given that one cannot climb politics in a literal sense, the resolution of the anaphora forces a metaphorical interpretation. As Asher and Lascarides write, “the relation, together with the proposition expressed by the first sentence, triggers the metaphor in the first place” (2001: 285).⁸

Discourse coherence rendered via rhetorical relations with previous discourse, however, cannot be the whole story. Asher and Lascarides note that “computing rhetorical relations involves nonmonotonic reasoning on the reader’s background knowledge” (2001: 283). Deciding which rhetorical relation(s) hold(s) between the sentence and its preceding discourse might not be as automatic as the presence of “but” suggests.⁹ Furthermore, even if the rhetorical relation Contrast is clearly identified, background knowledge plays a role in the identification of metaphor. Asher and Lascarides (2001) invoke knowledge about “fundamental values in our culture” which allow the interpreter to construct a scale, from clerk to prime minister, at stake in the spatial projection forced by the change of location denoted by the verb

⁶ “Predicts” is probably too strong, since it is possible to use “John is a rock” in its literal sense, or as a metonymy.

⁷ (1*) includes (1) in the sense that (1*) expresses a conjunction of two propositions, one of which is the proposition semantically expressed by (1) (ignoring the complications created by the indexical).

⁸ This is Asher and Lascarides’ explanation of the example. However, as a reviewer has noted, the presence of ‘but’ could also suggest a different interpretation. It could establish a contrast between politics and another domain that is as challenging as politics. This is precisely what we show in what follows.

⁹ In Asher and Lascarides’ view (2003), rhetorical relations are computed on the basis of composition and lexical semantics together with domain knowledge.

“to climb”. Yet they do not offer a systematic integration of this sort of background knowledge when it gets to establishing discourse relations between the proposition at hand and the utterances preceding it.

In the remainder of this section, we offer three examples showing that discourse structure established upon preceding utterances is insufficient to solve the metaphor identification problem. First, imagine two friends chatting about their holidays.

- (4) A: a. Did you enjoy your trip to Moscow?
 B: b. Not really.
 B: c. Moscow is a cold city.

B’s utterance is linked to A’s utterance via the relation Question-Answer-Pair, as per the Segmented Discourse Representation Theory (SDRT) framework. A coherent interpretation of B’s utterance would take it to address A’s question. This relation connects (4b) to (4a). What about (4c)? A plausible interpretation here is that c explains why the speaker did not enjoy Moscow. The link here is given by the rhetorical relation Explanation. Now, either a literal or a metaphorical interpretation of “Moscow is a cold city” would count as an explanation. Taken as a literal claim, B can be interpreted as saying that the temperatures in Moscow are very cold. This explains why she did not enjoy the trip. But a metaphorical interpretation would also work. Taken as a metaphor, “Moscow is a cold city” would be saying, roughly, that the inhabitants are unfriendly. This interpretation would also explain why she did not enjoy the trip.

What clues can the interpreter rely on that would help her decide between the literal and the metaphorical interpretation? In this first example, she could have recourse to what we call background information. Imagine that A knows that B lives in a place where temperatures are lower than in Moscow, and that B does not mind cold weather. With this information, it makes more sense for A to interpret B as speaking metaphorically.

As a second example, imagine a visibly restless teenager who tells her friend:

- (5) a. I need to get out.
 b. I can’t breathe.

Again, (5b) can be interpreted literally, as referring to a physical symptom, or metaphorically, as a description of feelings. In this example, the second sentence is again connected to the first via Explanation, which is again insufficient to decide between the metaphor-literal interpretations. Both the literal and the metaphorical interpretation could constitute an explanation of why the speaker needs to get out. However, the perceptual information available to the interpreter might supplement discourse structure and make the metaphorical interpretation preferable. Information about how the teenager behaves, gestures, or about her physical condition could tip the scale towards one interpretation or the other.

Finally, consider again (1*):

- (1*) a. I have always despised politics.
 b. But I have climbed to the top of that greasy pole.

Although discourse structure points towards a metaphorical interpretation, it is not difficult to come up with more complex examples in which extra-linguistic information suggests otherwise. Imagine, as historically implausible as it may sound, that Disraeli utters (1*) before becoming Prime Minister and after climbing an actual greasy pole. In this case, it would be reasonable to interpret (1*b) as a

literal claim. What about the whole fragment, (1*)? What role is (1*a) playing, and how does it constrain the interpretation of (1*b)? (1*a) has introduced the domain of politics. Given that the sentences are connected by “but”, which cues towards the rhetorical relation Contrast, and the matching congruity this relation imposes, it makes sense to interpret (1*b) as implicitly referring to the domain of politics introduced by (1*a). However, “that” could refer to the actual pole and the utterance would be literal. In this context, in order to make sense of the contrast relation, we could interpret (1*b) as a suggestion that Disraeli is capable of great achievements. The contrast here would be between the claim that the speaker despises politics and the suggestion that he is capable of great achievements in the domain of politics—between (1*a) and an implicature communicated by (1*b).¹⁰

4. Contextual Coherence

The examples above suggest that coherence should be understood as a notion that comprises preceding discourse as well as extra-linguistic information. Instead of understanding coherence as being based only on rhetorical relations that can be established between the target sentence and previous linguistic utterances, we think that coherence is established by also taking into account the surrounding, perceptually accessible context of utterance, and even background knowledge accessible to the interpreter. We develop a broader notion of contextual coherence that expands on Asher and Lascarides’s theory of discourse coherence. In our view, interpretation in general, and metaphor identification in particular, is governed by the following principle:

Principle of Contextual Coherence: The preferred interpretation of some utterance is (the) one that is contextually coherent. As a corollary of this principle, contextually incoherent interpretations should be avoided.¹¹

What do we mean by “contextually coherent”? The remainder of the paper elaborates on this. Here is a first sketch.

Contextual Coherence (basic definition): For an interpretation of some utterance to be contextually coherent, two conditions must be met:

- (1) It must cohere with the discourse in which the target sentence is embedded.
- (2) It must be externally consistent, by which we mean consistent with the context, that is, with the preceding discourse, perceptual information available to the interpreter, and background knowledge within her reach.

The first condition can be rendered more precise if we rely, as Asher and Lascarides and others do, on rhetorical relations, and understand ‘cohere with the discourse’ in (1) as meaning that the target sentence is connected to other parts of the discourse by rhetorical relations. (We shall see in the coming section,

¹⁰ Suppose that, in the context described, it is clear that (1*b) is literal, because Disraeli has just climbed to the top of an actual greasy pole, a pole that is still perceptually salient. This literal interpretation is in principle unrelated to (1*a), for, what relation is there between despising politics and climbing greasy poles? However, speaker and audience can exploit Grice’s maxim of relevance and derive the implicature that the speaker is capable of great achievements even in unpleasant domains, such as politics.

¹¹ This leaves room for there being more than one contextually coherent interpretation, something we explore in Crespo, Heise and Picazo (ms.).

however, that this strategy is too restrictive.) Thus viewed, condition (1) on contextual coherence would establish that the preferred interpretation is one in which the target sentence is connected to at least one previous sentence in the discourse via a rhetorical relation.¹² This already precludes certain interpretations and favours others. Asher and Lascarides's account of discourse (1*) shows that, in absence of further information, the Contrast relation precludes a literal interpretation of "I have climbed to the top of that greasy pole".

Nonetheless, we have argued in section 3 that, in some cases, rhetorical relations do not suffice to account for the preferred interpretation. We may encounter at least two kinds of scenarios in which the decision as to whether an utterance is best understood as a metaphor or as a literal claim is affected by other sources of information. On the one hand, some rhetorical relations do not resolve the interpretive ambiguity (examples 4 and 5 above). On the other, extra-linguistic information can override an otherwise coherent interpretation (example (1*)). This is why we need condition (2) on contextual coherence.

Condition (2) on contextual coherence holds that the interpretation must be consistent with the context. As suggested by our analyses of examples in section 3, we believe it is best to conceive of context as a rich and heterogeneous body of information. Here we adopt the perspective of the interpreter and include in this body of information all the information available to her. This encompasses information from three different sources. First, a context includes the information given by previous discourse. This includes what has been explicitly said, and also implicatures and presuppositions. It is difficult to decide what counts as previous discourse,¹³ but let us simply note that sometimes previous conversations might count as previous discourse. Second, the context includes perceptually accessible information given to the interpreter. This is especially relevant for oral exchanges, where the interpreter has information about the conversational setting, the context of utterance. Third, interpreters have what we will call background information, that is, general knowledge that can be used for the purpose of interpretation akin to what Asher and Lascarides call world knowledge, but also other sorts of background knowledge, for instance, knowledge about punctuation, register, tone, and other conventional elements of oral and written discourse.¹⁴ Here we include as well information about the speaker or writer (who she is, etc.) and more generic information, including, for instance, sociological, historical or literary knowledge that may allow an interpreter to classify a text or exchange according to general criteria.

All this presupposes an ideal interpreter, someone able to deal with large bodies of information. Actual interpreters have limitations and are often unable to reason on the basis of the complex body of information that, in our view, constitutes the context. Moreover, actual speakers might be inattentive, fail to reason adequately, commit mistakes, etc. Therefore, in our proposal, we work with an

¹² As in Asher and Lascarides's framework, condition (1) would thus not apply to the first sentence of some discourse.

¹³ To give a flavour of the difficulty in question, consider a chat between two friends: It is likely that their exchange draws in part on things they talked about in the past. This is not the case with two strangers who strike a conversation. In view of such problems of delimitation, various proposals have been elaborated on how to restrict the domain of previous discourse. Relevance Theory (Sperber and Wilson 1995) would be one way to go forward, but there are others too. We do not wish to commit here to one particular proposal.

¹⁴ Syntactic constraints would figure among the relevant linguistic conventions (cf. fn. 4).

ideal interpreter. Furthermore, presumably some pieces of previous discourse and extralinguistic elements, but not others, will be salient. In our proposal we assume that the context includes information about what is salient, without deriving salience from more basic contextual information.

Condition (2) on contextual coherence urges the interpreter to avoid interpretations that are inconsistent with contextual information. Here we mean “inconsistent” in a loose sense. Besides being logically inconsistent, an interpretation can be inconsistent because of material incompatibility or presupposition failure. In spite of this rather loose notion of inconsistency, our proposal is quite conservative in that we think of inconsistency as a relation between propositions. Thus, the three types of information that make up the context have propositional format, including perceptual information. The question of whether perceptual experience has propositional content is anything but settled, of course (see Crane and Craig 2021). Thus, we are making a substantive assumption here.

The assumption that inconsistency is a relation between propositions requires that we are able to consider metaphorical meaning in a propositional manner. But is there something equivalent to a proposition when it comes to metaphorical meaning? We propose to conceive metaphorical meaning in terms of minimal paraphrases.¹⁵ A minimal paraphrase is a propositional rendering of what the metaphor expresses. It is minimal in the sense of not being exhaustive. Some of the effects of a metaphor might involve mental imagery and thus be impossible to capture by means of a proposition. Moreover, the full significance of the metaphor might be so rich that it is very difficult to fully articulate its meaning. However, the minimal paraphrase captures enough so as to enable the interpreter to reason on the basis of the metaphor (to derive implicatures, for instance), and to assess consistency with contextual information.

The notion of contextual coherence we propose works well for cases as those considered in section 3 in which rhetorical relations are insufficient to account for the preferred interpretation of an utterance. Let us go back to example (4). In section 3 we noted that the rhetorical relation Question-Answer-Pair is insufficient to decide whether the best interpretation of “Moscow is a cold city” is metaphorical or not. Both the metaphorical and the literal interpretation provide an explanation as to why B did not enjoy her trip. However, we also noted that information about the speaker—including what we have called background information—might lead us to prefer one interpretation over the other. In particular, the information that the speaker does not mind low temperatures is inconsistent—in our loose sense—with interpreting her as meaning “Moscow is a cold city” in a literal sense. The two propositions that are inconsistent can be specified as follows: the proposition that the speaker didn’t enjoy her stay in Moscow because temperatures in Moscow are low (derived from the literal meaning of “Not really. Moscow is a cold city” and the rhetorical relation of Explanation), and the proposition that the speaker does not mind low temperatures. By contrast, the metaphorical interpretation of “Moscow is a cold city” makes sense of the rhetorical relation and is consistent with the context.

¹⁵ Non-cognitivist such as Davidson (1978) would disagree with the claim that metaphors have minimal paraphrases, and interaction theorists such as Black (1954) would hold that paraphrases come with a loss of cognitive force. However, there are linguistic reasons for positing such (minimal) cognitive contents for metaphorical utterances. For instance, metaphors can be used as vehicles for implicatures, see Bezuidenhout 2001 and Camp 2006.

Similarly, in example (5) perceptual information might preclude the literal interpretation. Suppose that the hearer looks around and sees nothing that could prevent the speaker from breathing in the literal sense—no smoke, no asthma symptoms, etc. In this case, the literal interpretation would be inconsistent with the perceptual information that there are no signs of a physical inability to breathe.¹⁶ The metaphorical interpretation is preferred. It is connected to previous discourse by a rhetorical relation and is consistent with perceptual information (the speaker is visibly restless).

5. Rhetorical Relations, Situations and Topic under Discussion

In this section, we discuss cases that lead us to modify the principle of contextual coherence sketched in section 4. Here, rhetorical relations, understood as relations between a target sentence and previous discourse, are not insufficient but inadequate to account for metaphor identification. One reason is that some sentences are connected to extra-linguistic context. Another reason is that we are sometimes able to identify metaphor even if all we have at our disposal is an impoverished context and no previous discourse at all. These cases will lead us to modify our conception of contextual coherence.

The first reason why we think it is necessary to go beyond the notion of rhetorical relation that we have relied on so far—relations established between an utterance and propositions available from previous discourse—is that rhetorical relations, or at least something akin to them, can also be established between an utterance and its extra-linguistic context. Let us go back to example (5):

- (5) a. I need to get out.
b. I can't breathe.

Imagine again that the utterer is a restless teenager with tears in her eyes. In the previous section, we noted that (5b) is an Explanation of (5a). But we think that this is not the whole story. In the scenario described, it seems that part of the reason why we take (5b) to be metaphorical has to do with how it is connected to the real-world situation, not merely by being consistent with it—as required by our condition (2) on contextual coherence—but as being an explanation thereof. It seems that here the proposition enters into a discourse relation (a rhetorical relation) with extra-linguistic elements of the context in which it is uttered, to the extent that they are perceptually accessible to the interpreter. Note that this relation of explanation would also hold between the situation and (5b) when not preceded by (5a).

Stone et al. (2013) let perceptually grounded elements intervene in the discourse relations that model discourse coherence. They consider the role of perceptual information in the interpretation of situated utterances, and bootstrap this input by means of a simple dynamic semantics which introduces perceptually grounded discourse referents. We do not delve into the formal details of their proposal, but simply point out that such a move allows them to integrate perceptual input into the establishment of discourse relations. In their view, rhetorical relations (what they call discourse relations) can be established between an utterance

¹⁶ Again, it would be inconsistent in our loose sense. In general, these loose inconsistencies can be translated into logical inconsistencies by adding the implicit assumptions that are used in reasoning. For instance, here we could add the premise that the physical inability to breathe always comes with certain visible signs.

and what happens in the real-world situation the utterance pertains to. Following this idea, we are in a position to suggest that (5b) yields an Explanation of what is happening in the corresponding real-world situation, which is perceptually available to the interpreter.

Something similar could happen with example (1). Imagine a recently appointed Prime Minister who receives an old friend in Downing Street. Looking around, she utters: “I have climbed to the top of the greasy pole”. We think that this utterance, in the context described, is rhetorically connected to the situation in which the exchange takes place. For instance, it might serve as an explanation of why the speaker is in Downing Street. If this line of reasoning holds good, then the relevant relation of Explanation is stronger than the notion of loose consistency ensured by our condition (2) on contextual coherence.

This raises a concern for our previous characterisation of contextual coherence. The restrictive view on condition (1) advanced in section 4 had it that the target sentence should stand in a rhetorical relation with other propositions made available by previous discourse. Our discussion above suggests that a more accurate condition would be:

Contextual Coherence (with situated utterances): For an interpretation of some utterance to be contextually coherent, two conditions must be met:

- (1*) It must cohere with previous discourse or the real-world situation in which it is embedded.
- (2) It must be externally consistent, by which we mean consistent with the context, that is, with the preceding discourse, perceptual information available to the interpreter, and background knowledge within her reach.

As in the basic definition, condition (1*) can be spelled out using rhetorical relations. These relations can now connect the target sentence with sentences in the previous discourse or with perceptually grounded discourse referents.

Another reason to revise our conception of rhetorical relations, in our attempt to account for metaphor identification via contextual coherence, is provided by a different kind of example. Let us go back to (1):

(1) I have climbed to the top of that greasy pole.

In section 1 we provided some contextual information for (1). Crucially, we provided some (very limited) previous discourse. Now, as it turns out, many interpreters of (1) only know that this sentence is a quote from Disraeli, something he said after being appointed Prime Minister, together with some minimal knowledge about who Disraeli was. For many interpreters, (1) comes in the form of a direct quote, as in the following example:

After being appointed Prime Minister, Disraeli said: “I have climbed to the top of that greasy pole”.

Previous discourse is not available to these interpreters, and yet they might be able to identify the sentence as a metaphor—as being about politics, and not about an actual greasy pole. Note that what they identify as a metaphor is Disraeli’s utterance of the sentence, that is, they take Disraeli’s words to be metaphorical. Their interpretation is not about someone else’s use of the same words, say a reporter.

We can speak here of metaphor identification in impoverished contexts.¹⁷ The situation is analogous to a well-known problem of inter-contextual communication. Inter-contextual communication includes all those instances of communication in which speaker and audience are not co-situated and lack a common ground. For example, reading a report of what someone said, either in the form of a direct quote or as a homophonic indirect report, would be such a case. The problem of inter-contextual communication is that the interpreter lacks some important contextual information. In particular, and going back to the notion of coherence introduced above, the interpreter lacks access to the previous discourse or real-world situation with which to establish rhetorical relations. Despite this, non-co-situated interpreters are sometimes able to identify metaphors, including twice-true metaphors.

Impoverished contexts motivate an even more minimal notion of contextual coherence. In our view, the non-co-situated interpreter has what we have called ‘background information’. For example, she might have information about the speaker. A promising option of filtering that information is to consider the topic under discussion. When interpreting example (1), we might lack knowledge of previous discourse and have no perceptual access to the situation that the utterance is related to. However, by drawing on background information, it seems possible to reconstruct at least the topic. Since we know that Disraeli was a politician, we can venture the guess that the utterance might be about politics. This piece of information plausibly plays a role in identifying the utterance as metaphorical.

As the foregoing shows, we think that the topic under discussion can sometimes be reconstructed from background information. Remember that this includes information about the speaker and known historical facts, for instance. Thus, background information can play two roles. First, it may be used to infer a topic under discussion. Second, it enters into the assessment of contextual consistency.

This motivates a third version of the notion of contextual coherence:

Contextual Coherence (even in impoverished contexts): For an interpretation of some utterance to be contextually coherent, two conditions must be met:

(1**) It must *address the topic under discussion*.

(2) It must be externally consistent, by which we mean consistent with the context, that is, with the preceding discourse, perceptual information available to the interpreter, and background knowledge within her reach.¹⁸

To finish, we would like to suggest that this notion of contextual coherence can be applied beyond impoverished contexts. Even in cases in which there is a rich previous discourse, it might make sense to use the topic under discussion, instead of relying only on rhetorical relations, to decide whether the target sentence is literal or metaphorical. We will illustrate this possibility with example (5).

¹⁷ See Picazo 2022 for a discussion of inter-contextual communication and the kind of information that is shared across contexts.

¹⁸ Note that an interpreter with little access to elements of the context of utterance has more chances of misinterpreting what the speaker. The less contextual information, the more interpretive failures can arise. This is so because the interpreter can lack contextual information that would defeat her interpretation of the utterance, or information that would support a different interpretation.

Consider another context for example (5), this time Tamara Bach's (2003) novel *Marsmädchen*. The relevant paragraph reads as follows:

I have to get out. I have to do something, now, immediately. The room is too small, it's afternoon again, and it's always the same: too similar, too small. Not just my room, heck, the whole house, this city. I have to get out, I can't breathe in here (Bach 2003: sect. 5).

As we mentioned above, (5b) is related to (5a) via Explanation. Other relations can be established in the text, and rhetorical relations can be used to decide that the target sentence is metaphorical. The explanation might go as follows. The two sentences before (5) seem to be related to (5) via Result or Cause-Effect—the room, the house and the city are too small, so the speaker feels confined and has to get out. We think that this relation, together with some background knowledge, constrains interpretation in a way that makes the metaphorical interpretation coherent, and hence preferred. In this fragment, the relations among the sentences impose some constraints on interpretation. First, 'here' in (5) must refer to a location introduced in the previous sentence, if the fragment is to be coherent. We conclude that it refers to the room, the house, the city. Using only the latter, the target sentence says that the speaker needs to get out of the city. We have split (5) into two sub-sentences connected by Explanation. This relation imposes a connection between the two sentences such that not being able to breathe must be interpreted as a reason for needing to get out of the city. At this point, we need to use background information. We know that the size of the city does not cause a physical inability to breathe. It can, however, constitute an oppressive environment. This metaphorical interpretation is more coherent.

Rhetorical relations may be appealed to here. However, we think that the topic under discussion can play the same role. The novel is about a teenager, Miriam, who falls in love with her friend Laura. One topic under discussion concerns Miriam's feelings—how she feels, how she reacts to an oppressive environment. The whole paragraph can be read as addressing this topic. If we think of the paragraph as addressing Miriam's feelings, then the metaphorical interpretation would be coherent and the literal not.

How to account for a topic under discussion? In related work, we have proposed to view the topic under discussion as being given by the Questions Under Discussion (QUD) in the context, as identified from the interpreter's point of view (Crespo, Heise, Picazo, ms.). As rhetorical relations, QUD is a framework to model discourse structure.¹⁹ The basic idea is that the utterances in a conversation are interpreted by the interlocutors relative to the question being addressed, that is, relative to the topic under discussion. We think of QUD as a partially ordered set of questions, i.e., a set of questions ordered by a binary transitive and antisymmetric relation. Peculiar to our view is that we consider that QUDs may issue from the three sources that make up context: previous discourse, perceptual information and background information. That the topic under discussion may be inferred from previous discourse, and also from the target sentence itself via accommodation, is standardly accepted. But we think that it is also important to pay attention to the role that perceptual information and background knowledge play. As for perceptual information, whatever is perceptually salient in the situation in

¹⁹ Originally proposed in von Stutterheim and Klein 1989 and van Kuppevelt 1995. Later elaborations can be found among others in Ginzburg 1996, Roberts 2012, and van Rooij 2013.

which a conversation takes place may become a topic under discussion. Gestures, for instance, may introduce a topic under discussion, as when pointing makes an object or individual salient. These moves can introduce brand new topics but also sub-topics that relate to the main topic under discussion. As for background information, we have already argued that it plays a role in the determination of the QUD, notably in impoverished contexts.

With this in mind, how to conceive of condition (1**) on contextual coherence? We believe that the target sentence should address an issue raised by the heterogeneous, threefold context in which the utterance is given to the speaker, conceived as a live question in the QUD, that is, a question that has not yet been downdated, that has not been addressed already. Rhetorical relations or discourse relations, featured in the previous formulations of condition (1), are perhaps more variegated than the kind of constraint given by (1**) in our third attempt. However, they seem to be too stringent in how rich the context needs to be for them to be established in the first place. Hence we suggest that, at least for many examples, QUD can replace rhetorical relations in condition (1). If successful, this option would allow us to treat the different kinds of examples we have considered so far in a uniform manner.

Contextual Coherence (with QUD): For an interpretation of some utterance to be contextually coherent, two conditions must be met:

(1***) It must *answer a live question in the QUD*.

(2) It must be externally consistent, by which we mean consistent with the context, that is, with the preceding discourse, perceptual information available to the interpreter, and background knowledge within her reach.

It lies beyond the scope of this paper to assess whether QUD are a better tool overall to solve the problem of metaphor identification. Our aim with this fourth definition is more modest: We suggest that this unified definition can be used to deal with many examples, including metaphors whose identification proceeds via previous discourse (cases where previous discourse introduces a question to the QUD), metaphors that involve real-world situations (cases where the question is introduced via perceptual information), and impoverished contexts (cases where background information suggests a QUD).

6. Conclusions

We have proposed a solution to the problem of metaphor identification based on the notion of contextual coherence. Following Asher and Lascarides, we hold that coherence guides the identification of figures of speech. We agree with Asher and Lascarides that coherence is primarily a notion that concerns discourse structure. However, we have shown that in order to decide whether a given utterance is a case of metaphor, rhetorical relations are sometimes insufficient and sometimes inadequate. The reason why they are insufficient is that sometimes rhetorical relations fall short of providing a basis for deciding between literal and metaphorical interpretations. Contextual information other than previous discourse, such as background information about the speaker, may then play a crucial role in identifying metaphors. We have captured this in a first notion of contextual coherence (the basic notion). According to this first notion, coherent interpretations must cohere with the discourse they are embedded in and must be externally consistent,

that is, consistent with the contextual information available to the speaker. Coherence with previous discourse could in principle be explained using rhetorical relations. However, we have argued that rhetorical relations, conceived as relations among sentences, are sometimes an inadequate tool to account for coherence. The reason is that extra-linguistic elements of the situation in which the sentence is uttered may be crucial in getting at the preferred interpretation. In order to account for these cases, we proposed that rhetorical relations may be established both with previous discourse and extra-linguistic situations (second notion, contextual coherence with situated utterances). Finally, we have suggested another possible implementation of the notion of contextual coherence. This alternative notion makes use of QUD instead of rhetorical relations (fourth notion, contextual coherence with QUD). What speaks in its favour is the fact that this alternative notion promises not only to explain the same kind of cases as the version based on rhetorical relations, but it also accounts for metaphor identification in impoverished contexts.²⁰

References

- Asher, N. and Lascarides, A. 1995, "Lexical Disambiguation in a Discourse Context", *Journal of Semantics*, 12, 1, 69-108.
- Asher, N. and Lascarides, A. 2001, "Metaphor in Discourse", in Bouillon, P. and Busa, F. (eds.), *The Language of Word Meaning*, Cambridge: Cambridge University Press, 262-89.
- Asher, N. and Lascarides, A. 2003, *Logics of Conversation*, Cambridge: Cambridge University Press.
- Asher, N. 2011, *Lexical Meaning in Context: A Web of Words*, Cambridge: Cambridge University Press.
- Bezuidenhout, A. 2001, "Metaphor and What is Said: A Defense of a Direct Expression View of Metaphor", *Midwest Studies in Philosophy*, 25, 1, 156-86.
- Bach, T. 2003, *Marsmädchen*, Hamburg: Verlag Friedrich Oetinger.
- Camp, E. 2006, "Contextualism, Metaphor, and What is Said", *Mind & Language*, 21, 3, 280-309.
- Cohen, T. 1976, "Notes on Metaphor", *Journal of Aesthetics and Art Criticism*, 34, 3, 249-59.
- Crane, T. and Craig, F. 2021, "The Problem of Perception", in Zalta, E.N. (ed.), *The Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/archives/fall2021/entries/perception-problem/>.
- Crespo, I., Heise, A., and Picazo, C. (ms.), *Metaphor and Contextual Coherence: It's a Match!*
- Davidson, D. 1978, "What Metaphors Mean", *Critical Inquiry*, 5, 1, 31-47.
- Ginzburg, J. 1996, "Dynamics and the Semantics of Dialogue", in Seligman, J. and Westerståhl, D. (eds.), *Logic, Language and Computation*, Vol. 1, Stanford: Center for the Study of Language & Information (CSLI), 221-37.

²⁰ Andreas Heise wishes to acknowledge support by the grant ANR-17-EURE-0017. Claudia Picazo wishes to acknowledge support by the Ministerio de Ciencia e Innovación (Government of Spain) and European Union; grant PID2019-105728GB-I00 (MINECO/FEDER, EU).

- Magidor, O. 2020, “Category Mistakes”, in Zalta, E.N. (ed.), *The Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/archives/win2020/entries/category-mistakes/>.
- Picazo, C. 2022, “Homophonic Reports and Gradual Communication”, *Pacific Philosophical Quarterly*, 103, 2, 259-79.
- Pragglejaz Group 2007, “MIP: A Method for Identifying Metaphorically Used Words in Discourse”, *Metaphor and Symbol*, 22, 1, 1-39.
- Prandi, M. 2021, “Syntaxe formelle et cohérence textuelle: Deux sources pour le conflit conceptuel”, in Paissa, P., Conoscenti, M., Druetta, R., and Solly, M. (eds.), *Metaphor and Conflict. Métaphore et conflit*, Bern: Peter Lang, 55-75.
- Pustejovsky, J. 1995, *The Generative Lexicon*, Cambridge, MA: MIT Press.
- Roberts, C. 2012, “Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics”, *Semantics & Pragmatics*, 5, 1-69.
- Sperber, D. and Wilson, D. 1995, *Relevance Theory*, Oxford: Blackwell.
- Stern, J. 2000, *Metaphor in Context*, Cambridge, MA: MIT Press.
- Stone, M., Stojnic, U., and Lepore, E. 2013, “Situated Utterances and Discourse Relations”, *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013)*, – *Short papers*, Potsdam: Association for Computational Linguistics, 390-396.
- van Kuppevelt, J. 1995, “Discourse Structure, Topicality and Questioning”, *Journal of Linguistics*, 31, 109-47.
- van Rooij, R. 2013, “Questioning to Resolve Decision Problems”, *Linguistics and Philosophy*, 26, 727-63.
- von Stutterheim, C. and Klein, W. 1989, “Referential Movement in Descriptive and Narrative Discourse”, in Dietrich, R. and Graumann, C.F. (eds.), *Language Processing in Social Context*, Amsterdam: Elsevier, 39-76.