

Discourse Relations and Default Inferences in Metaphor Interpretation

Rodrigo Agerri
School of Computer Science
University of Birmingham, UK
email: r.agerri@cs.bham.ac.uk

Abstract

This paper focuses on the interpretation of metaphor in discourse. We build on previous work (Agerri, Barnden, Lee, & Wallington, 2007) to offer a formalization in a computationally-oriented formal semantic framework of a set of mappings that we claim are required for the interpretation of *map-transcending* metaphor. Such mappings are domain-independent and are identified as invariant adjuncts to any conceptual metaphor. In this paper we claim that the invariant adjunct mappings allow us to account for metaphors where inferring discourse structure is not sufficient. Moreover, these mappings interact with rhetorical relations in order to explain cases in which recognizing their structure is not sufficient to provide an interpretation of a metaphorical discourse.

Keywords: Discourse Structure, Metaphor Interpretation, Computational Semantics.

1 Introduction

We assume the general view that metaphor understanding involves some notion of events, properties, relations, etc. that are transferred from a source domain into a target domain. In this view, a metaphorical utterance conveys information about the target domain. We are particularly interested in a type of metaphorical utterances that we call *map-transcending*. A characteristic of *map-transcending* metaphor is that finding a target correspondent for every aspect of the source domain is a difficult task which, in some cases, seems to be plainly impossible. Thus, this type of metaphor poses great difficulties for correspondence-based approaches (Lakoff, 1993) which require to establish a parallelism between the source and target domains to explain metaphor.

We believe that an account of metaphor interpretation ought to explain what extra information *map-transcending* entities convey and it should provide a viable (computational) mechanism to explain how this transfer of information occurs. Moreover, it should do so by taking into account the fact that metaphor is a highly contextual phenomenon.

This paper addresses these two issues: Firstly, it builds on Agerri et al. (2007) to provide a formal set of invariant mappings that we call View-Neutral Mappings Adjuncts (VNMAAs) for the interpretation of *map-transcending* metaphor. Secondly, it grounds the invariant mappings on a (modified) computationally-oriented formal semantic framework for the interpretation of metaphor in discourse (Asher & Lascarides, 2003).

In order to do so, we first discuss the problems of correspondence approaches to deal with *map-transcending* metaphor. In section 3 we argue that inferring discourse structure is not sufficient to interpret certain metaphors. Sections 4 and 5 briefly describe our approach to metaphor interpretation. Section 6 describes a number of VNMA's that are particularly useful to interpret *map-transcending* metaphor. In section 7 we propose to adapt Segmented Discourse Representation Theory (SDRT) (Asher & Lascarides, 2003) to our purposes of providing a formal account of metaphor interpretation based on the ATT-Meta approach. Finally, in section 8 we present some conclusions and discussion on further work.

2 Missing Correspondents

We do not address in this paper the issue of when an utterance is to be considered metaphorical. Instead, we aim to offer an explanation of how a metaphorical utterance such as (1) can be interpreted.

- (1) McEnroe starved Connors to death.

If we infer, using our knowledge about McEnroe and Connors, that (1) is used to describe a tennis match, it can be understood as an example of the conceptual metaphors (or, in our terminology, ‘metaphorical views’) DEFEAT AS DEATH and NECESSITIES AS FOOD. However, these metaphorical views would not contain any relationship that maps the specific *manner* of dying that constitutes *being starved to death* (we say that “starving” is a map-transcending entity as it goes beyond known mappings). Yet one could argue that the *manner* of Connors’s death is a crucial part of the informational contribution of (1).

A possible solution would be to create a new view-specific mapping that goes from the form of killing involved in *starving to death* to some process in sport, but such enrichment of mappings would be needed for many other verbs or verbal phrases that refer to other *ways* in which death is brought about, each requiring a specific mapping when occurring in a metaphorical utterance. Thus, finding adequate mappings could become an endless and computational intensive process. Moreover, there are even cases in which we may not find a plausible mapping. Consider the following description of the progress of a love affair:

- (2) We are spinning our wheels.

It is not very clear what could be a target correspondent for ‘wheels’; the unavailability of a correspondent would therefore prevent the source to target transfer of information needed for the interpretation of the metaphorical utterance. Thus, an account of metaphor ought to explain what extra information map-transcending entities provide. Furthermore, how the transfer of information occurs should be accounted for in a viable computational manner.

3 Metaphor in Discourse

Consider the following example:

- (3) Sam is a pebble.

Asher and Lascarides (2001) claim that it is not possible to calculate the meaning of an utterance such as (3) on the basis of the domain information about pebbles, but that it is possible to process it if it is discourse related to other utterance such as in the discourse

(4) John is a rock but Sam is a pebble.

Specifically, they argue that inferring the *Contrast* discourse relation would help us to work out the metaphorical meaning of (3). A similar point is made by Hobbs (1996):

(5) John is an elephant.

Which Hobbs argue can only be interpreted if we add extra information such that the example now consists of:

(6) Mary is graceful but John is an elephant.


Hobbs also infers *Contrast* in order to work out the meaning of “John being an elephant” as opposed to “Mary being graceful”. Thus, not only it is necessary the availability of a whole discourse, as opposed to a single sentence, in order to interpret metaphor, but also these authors claim that recognizing the discourse structure is a required step to reach a full interpretation.

In this paper we claim that, in some cases, the inference of the discourse structure is not sufficient to to interpret the metaphor, and that additional means of analysis should be considered. It seems that some rhetorical relations such as *Contrast* provide enough information to infer the metaphorical meaning in examples (6) and (4). These issues raise quite easily when we consider discourse from which a *Coordinator* relation is inferred.

Gomez-Txurruka (2003) argued that when the conjunction ‘and’ relates two clauses a *Coordinator* discourse topic is inferred which leads to the recognition of a *Coordinator* rhetorical relation. A *Coordinator* discourse relation marked by ‘and’ is not cancelable. Consider the sentence below:

(7) Mary is a fox and John is an elephant.

We can represent the discourse structure of this discourse as follows (we simplify the event-analysis for clarity of exposition):



We infer a *Coordinator* discourse relation to account for the conjunction of the two segments. However, in our view inferring *Coordinator* would not be enough to address the fact that the information conveyed by (7) may be related to attributes of Mary (e.g., being cunning) and John (possessing a good memory).

Discourse-based approaches to metaphor such as Hobbs (1996) and Asher and Lascarides (2001) do not account for map-transcending entities, but they usually

assume that there is some straightforward correspondence between the concepts in the source and target domains. Moreover, it seems that in some cases the inference of discourse relations is not enough to interpret some utterances. At the same time, a computational account of metaphor should address the occurrence of metaphor in discourse.

In the following sections we argue the existence of other type of information that must be inferred if we want to explain in a systematic way metaphors such as the ones discussed up to this point.

4 VNMA's in ATT-Meta

Previous work (Wallington, Barnden, Glasbey, & Lee, 2006) has shown evidence that there are metaphorical aspects (relations between events such as *causation* and event properties such as *rate* and *duration*) that, subject to being called, invariantly map from source to target whatever metaphorical view is being used. We refer to these type of mappings as VNMA's. The VNMA's are a central component of the ATT-Meta approach and AI System to metaphor interpretation (Barnden & Lee, 2002).

ATT-Meta (Barnden & Lee, 2002) is an AI System and approach to metaphor interpretation that, apart from providing functionalities such as uncertainty and conflict handling, introduces two features central to the interpretation of metaphorical utterances such as (1) and (2): Instead of attempting the creation of new mappings to extend an existing metaphorical view, ATT-Meta employs query-driven reasoning within the terms of the source domain using various sources of information including *world* and *linguistic knowledge*. The nature of source domain reasoning in metaphor interpretation has not previously been adequately investigated, although a few authors have addressed it to a limited extent (Hobbs, 1996; Martin, 1990; Narayanan, 1997).

By means of VNMA's and source domain reasoning it is possible to reach an interpretation of (3) without necessarily needing a rhetorical relation such as Contrast to guide the reasoning. Thus, linguistic knowledge and source domain reasoning about 'pebbles' may establish that they are small, and a very frequent association of unimportant entities with "small size" allows the defeasible inference that something is low, inferior, limited in worth (see Wordnet or any other lexical database). Using a Value-Judgment VNMA to express that "Levels of goodness, importance, etc., assigned by the understander in the source domain map identically to levels of goodness, etc.", we can convey the meaning that Sam is limited in worth (worthless). Of course, the interpretation of (3) will vary if we change the discourse context.

Following this, and subject to the appropriate contextual query to be provided by the discourse, size-related features might be transferred in our approach by a Physical Size VNMA; in an appropriate context (7) could also be used to convey that John has a good memory and that Mary is cunning. In this case, forgetfulness could be seen a tendency to perform a mental act of a certain type and non-forgetfulness could be handled by a Negation VNMA, Mental states VNMA and a Event-Shape VNMA (for tendencies).

It may well be possible that studying the interaction between VNMA's and discourse relations may allow us to naturally extend the study of metaphor to discourse. For example, in cases such as (7) both VNMA's and rhetorical relations would be needed in order to give a full account of its interpretation. The interaction between VNMA's and rhetorical relations is particularly clear when we consider cases of temporal metaphor (see Glasbey, Barnden, Lee, & Wallington, 2004 for details on temporal

metaphor and discourse structure).

“With **the end of the digging season fast approaching**, we were forced to assess how we should proceed. “Richard and I stared at the site **over lunch** and decided that if the rest of the skeleton is there, it is dispersed over a wide area”, recorded Alan, **on the evening of September 19, three days before we pulled out.**” (Leakey & Lewin, 1992)

Given the prevalence of temporal metaphor in discourse, even if we are able to infer correctly the discourse structure of the text above, we would still be far from having a complete account of the temporal relations conveyed. At least until we are able to account for the contribution that temporal metaphors make to the structure of discourse.

5 Source Domain Reasoning and VNMAAs

Let us go back to example (1):

- (1) McEnroe starved Connors to death.

Assuming a commonsensical view of the world and if (1) is being used metaphorically to describe the result of a tennis match, a plausible target interpretation would be that McEnroe defeated Connors by performing some actions to deprive him of his usual playing style. In the ATT-Meta approach, source domain inferencing produces a proposition to which we may apply a mapping to transfer that information. Thus, and assuming a commonsensical view of the world, a source domain meaning would be that McEnroe *starved* Connors to death in a biological sense. The source domain reasoning can then conclude that McEnroe *caused* Connors’s death by *depriving* or disabling him. Leaving some details aside, the partial logical form (in the source domain) of the metaphorical utterance (1) may be represented as follows (without taking into account temporal issues):

$$(i) \exists x, y, e (McEnroe(x) \wedge Connors(y) \wedge starve - to - death(e, x, y))$$

This says that there is an event e of x starving y to death (we use the notion of event à la Hobbs (1996) to describe situations, processes, states, etc.). It may be suggested that if we were trying to map the partial expression (i), its correspondent proposition in the target could be expressed by this formula:

$$(ii) \exists x, y, e (McEnroe(x) \wedge Connors(y) \wedge defeat(e, x, y))$$

According to this, the event of x defeating y in the target would correspond to the event of x starving y to death in the source. However, by saying “McEnroe starved Connors to death” instead of simply “McEnroe killed Connors” the speaker is not merely intending to convey that McEnroe defeated Connors, but rather something related to the manner in which Connors was defeated. Following this, *starving* may be decomposed into the cause e_1 and its effect, namely, “being deprived of food”:

$$(iii) \exists x, y, z, e_1, e_2, e_3 (McEnroe(x) \wedge Connors(y) \wedge food(z) \wedge starve(e_1, x, y) \wedge death(e_2, y) \wedge deprived(e_3, y, z) \wedge cause(e_1, e_3))$$

Note that by factoring out “starving to death” in this way we not only distinguish the cause from the effect but doing so allows us to establish a relation between “death” in the source to “defeat” in the target using the known mapping in DEFEAT AS

DEATH (and possibly “starving” to “McEnroe’s playing” although we will not press this issue here).

Now, by means of lexical information regarding “starving”, it can be inferred that McEnroe deprived Connors of a necessity (see, e.g., Wordnet), namely, of the food required for his normal functioning (the NECESSITIES AS FOOD metaphorical view would provide mappings to transfer food to the type of shots that Connors *needs* to play his normal game). In other words, Connors is defeated by the particular means of depriving him of a necessity (food) which means that being deprived causes Connors’s defeat. This fits well with the interpretation of (1) where McEnroe’s playing deprived Connors of his usual game. Moreover, linguistic knowledge also provides the fact that starving someone to death is a gradual, slow process. The result of source domain inferencing may be represented as follows:

$$(iv) \exists x, y, z, e_1, e_2, e_3 (McEnroe(x) \wedge Connors(y) \wedge food(z) \wedge starve(e_1, x, y) \wedge death(e_2, y) \wedge deprived(e_3, y, z) \wedge cause(e_1, e_3) \wedge cause(e_3, e_2) \wedge rate(e_1, slow))$$

‘Slow’ refers to a commonsensical source domain concept related to the progress rate of *starving*. Now, the existing mapping DEFEAT AS DEATH can be applied to derive, outside the source domain, that McEnroe defeated Connors, but no correspondences are available to account for the fact that McEnroe *caused* the defeat of Connors by depriving him of his normal play. Furthermore, the same problem arises when trying to map the slow progress *rate* of a process like starving.

In the ATT-Meta approach to metaphor interpretation, the mappings of *caused* and *rate* discussed above are accomplished by the type of invariant mappings that we specify as VNMA’s (the Causation and Rate VNMA’s, respectively; see Wallington *et al.* (2006) for an informal but detailed description of a number of VNMA’s).

VNMA’s account for the mapping of aspects of the source domain that do not belong to a specific metaphorical view but that often carry an important informational contribution (or even the main one) of the metaphorical utterance. These source domain aspects can be captured as relationships and properties (causation, rate, etc.) between two events or entities that, subject to being called, identically transfer from source to target.

Summarizing, the following processes, amongst others, are involved in the understanding of map-transcending utterances:

1. Construction of source domain meaning of the utterance.
2. Source-domain reasoning using the direct meaning constructed in 1. with world and linguistic knowledge about the source domain.
3. Transfers by application of specific mappings in metaphorical views and often invariant mappings specified as VNMA’s.

6 Description of VNMA’s

By using VNMA’s and source domain inference, we do not need to extend the mappings in the metaphorical view to include information about “depriving of a necessity”, “food” or “causing Connors’s death”. VNMA’s transfer those properties or relations between mappées that are *view-neutral*. Moreover, VNMA’s are *parasitic* on the metaphorical views in the sense that they depend on some mappings to be established for the VNMA’s to be triggered. That is why VNMA’s are merely “adjuncts”.

VNMA can also be seen as pragmatic principles that guide the understanding of metaphor by transferring aspects of the source domain that remain invariant.

In example (1), there are two VNMA involved in the transfer of the causation and the “slowness”, namely, the Causation and Rate VNMA which are described below. Additionally, we also discuss a VNMA related to the temporal order of events (others are described in Barnden *et al.* (2003) and Wallington *et al.* (2006)).

6.1 Causation/Ability

The idea is that there are relationships and properties (causation, (dis)enablement, etc.) between two events or entities that identically transfer from source to target. We use the \mapsto symbol to express that this mapping is a default.

Causation/Ability VNMA: “Causation, prevention, helping, ability, dis-enablement and easiness/difficulty relationships or properties of events between events or other entities in the source domain, map to those relationships between their mappees (if they have any) in the target.” The invariant mapping involved in the interpretation of (1) could be represented as follows:

$$\text{Causation: } \forall e_1, e_2 (\text{cause}(e_1, e_2)_{\text{source}} \mapsto \text{cause}(e_1, e_2)_{\text{target}})$$

As an additional note, the specific mapping of each event or state variable does not depend on the VNMA but on the metaphorical view in play. For example, if we consider the contemporary situation in which McEnroe and Connors are tennis pundits on TV, we may need a metaphorical view such as ARGUMENT AS WAR to interpret the utterance “McEnroe starved Connors to death”. In other words, VNMA do not themselves establish the mappees between source and target.

6.2 Rate

Rate: “Qualitative rate of progress of an event in the source domain maps identically to qualitative rate of progress of its mappee. E.g., if an event progresses slowly (in the context of the everyday commonsensical world), then its mappee progresses slowly (in the target context)”.

Consider the following utterance:

(8) My car gulps gasoline.

Briefly, the metaphorical view involved is MACHINES AS CREATURES, that maps biological activity to mechanical activity. Source domain reasoning may be performed along the following lines: It can be inferred that gasoline helps the car to be alive, therefore, it helps the car to be biologically active. The Causation/Ability VNMA (which deals with helping) combined with the above metaphorical view provide the target domain contribution that gasoline helps the car to run. Given that we can assume that an act of gulping is normally moderately fast the use of the Rate VNMA allows us to conclude that the car’s use of gasoline is moderately fast. The logical form of this VNMA is could be expressed as follows:

$$\text{Rate: } \forall e, r (\text{rate}(e, r)_{\text{source}} \mapsto \text{rate}(e, r)_{\text{target}})$$

If the rate an event e in the source is r , then the rate maps to the mappee event in the target, that is, it also has rate r ; r refers to the qualitative rate of progress or duration of an specific event e .

6.3 Time-Order

Time-Order: “The time order of events in a source domain is the same as that of their mappee events, if any”.

Time-order is quite useful for map-transcending examples such as

- (9) McEnroe stopped hustling Connors.

We might infer in the source domain that McEnroe was once hustling Connors which would be transferred by the Time-Order VNMA. For the formalization of this VNMA, we say that if event e_1 precedes event e_2 in the source, then the mappee events in the target exhibit the same ordering.

Time-Order: $\forall e_1, e_2 (precede(e_1, e_2)_{source} \mapsto precede(e_1, e_2)_{target})$

7 Metaphor in a Semantic Framework

Embedding the VNMAs in a semantic framework for metaphor interpretation is useful as a first step towards their implementation as default rules in the ATT-Meta system, but it is also interesting in its own right to show the contribution that the ATT-Meta approach can make towards a semantics of metaphor. In the somewhat simplified discussion on the source domain reasoning and VNMAs employed in the interpretation of (1), we have not stressed the fact that actually the source domain reasoning performed by the ATT-Meta system is query-driven. Although in previous sections we used various sources of contextual information to license certain source domain inferences, we have only considered isolated metaphorical utterances, and metaphor understanding has been illustrated as a process of forward reasoning from the direct meaning of utterances (in the source domain) and then the application of various metaphorical mappings to the result of source domain reasoning to arrive at the informational contributions in the target. Moreover, other possible inferences that could be drawn were ignored without specifying any principles or criteria whereby the reasoning could be guided towards the particular informational contributions discussed. The notion of discourse-query-directed reasoning provides such a guidance. When analyzing previous examples, we assume that the surrounding discourse context supplies queries that guide source domain reasoning in broadly the reverse order to that in which we described them in section 5. Other authors such as Hobbs (1996) and Asher and Lascarides (2001) also acknowledge the importance of context-derived reasoning queries play an important role in the interpretation of metaphorical utterances.

We are not claiming that query-directed reasoning may be the only type of reasoning involved in the processing of metaphor, but it seems to be particularly important in the processing of connected discourse. Although the ATT-Meta system at present works with single-sentence utterances (albeit with the aid of discourse-query-directed reasoning), an aim for future versions is to extend it to the processing of *discourse*, and the semantic framework will need to allow for this.


We have been using various sources of contextual knowledge that interact in the processing of the utterance:

- a) View-specific mappings provided by the relevant metaphorical views (DEFEAT AS DEATH and NECESSITIES AS FOOD);
- b) Linguistic and contextual information necessary for source domain reasoning;

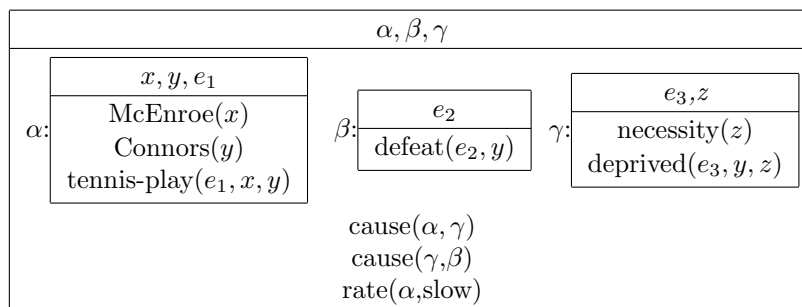
- c) Relations and properties between events such as *causation* and *rate* that are inferred in the source;
- d) VNMA's that transfer event relations and properties from source to target; and finally,
- e) Rhetorical relations that take into account the structure of discourse. In our view, a suitable approach to metaphor in discourse should include at least these five components.

7.1 Semantics for the ATT-Meta approach

Metaphor is a highly contextual phenomenon, and one of the most interesting semantic approaches that model context are dynamic semantics such as Segmented Discourse Representation Theory (SDRT) (Asher & Lascarides, 2003). Specifically, we adapt the semantic representation procedure of SDRT to build Segmented Discourse Representation Structures (SDRSs) consisting of the result of source domain reasoning. The conclusion of source domain inference can in turn be mapped to the target by using various view-specific mappings and VNMA's. In other words, we can see the source SDRS as the input for what the ATT-Meta system does when interpreting metaphor – it will reason with it, producing an output of inferred target facts which we may also represent by means of an SDRS. The result of reasoning in the source domain to interpret (1) would now look as follows:



where α and β are labels for DRSs representing events and \mapsto mappings (VNMA's and central mappings) needed in the interpretation of the metaphorical utterance. Importantly, the VNMA's would pick upon aspects such as causation and rate from the source to transfer them to the target producing an output which could also be represented as a SDRS:



Note that this formal representation integrates the systematicity of mapping invariantly certain aspects of metaphorical utterances by formulating them as relations and properties of events that can be represented as relations and properties of DRSS. For this purpose we will need to modify the construction rules of SDRSs to be able to infer properties and relations involving individuals (x, y, \dots) and not only DRSSs' labels such as α and β . In addition to this, we need to capture the interaction of the various sources of information used (linguistic knowledge, world knowledge, etc.) to infer causation and rate in the source domain. Thus, we partially adopt SDRT formal framework to represent ATT-Meta's source domain reasoning, event relations, event properties and VNMA's with the purpose of developing a semantic account of metaphor interpretation.

7.2 Discourse Contexts

Source domain reasoning partially relies on inferences provided by the discourse context and linguistic and world knowledge. In the ATT-Meta system, world knowledge roughly corresponds to source domain knowledge. On the one hand, we have been using our commonsensical knowledge about McEnroe and Connors to interpret example (1) as metaphorically describing a tennis match. On the other hand, linguistic knowledge is used to *pretend* that the direct meaning of the metaphorical utterance is true, which allows us to derive *causation* and *rate*. Thus, we assume that the understander possesses some world knowledge that provides information about “starving someone to death”:

- If e_3 where y is deprived and e_1 where x starves y are connected, then by default, e_1 *causes* e_3 .
- If e_2 where y dies and e_3 where y is deprived are connected, then by default, e_3 *causes* e_2 .
- If e_1 where x starves y , then by default, the rate of progress of e_1 is *slow*.

Furthermore, common sense about causation tells us that “if e_1 causes e_3 then e_3 does not occur before e_1 ”. Following this, the knowledge needed to interpret example (8) needs to include the that the drinking rate is fast:

If e where x gulps, then by default, x in e drinks moderately fast.

SDRT specifies where in the preceding discourse structure the proposition introduced by the current sentence can attach with a discourse relation. In order to do that, it is necessary to provide a set of rules for the understander to infer which discourse relation should be used to do attachment. We adopt a similar notation to represent discourse update (see Asher & Lascarides, 2003 for details on the discourse update function) so that defeasible knowledge about causation, rate, temporal order, etc., allows the inference of source domain event relations and properties.

Let us suppose that in a context (source domain) ω we want to attach some event denoted by β to α , such that $\langle \omega, \alpha, \beta \rangle$. This update function can be read as “the representation ω of a text so far is to be updated with the representation β of an event via a discourse relation with α ” (Asher & Lascarides, 2003). Let \rightsquigarrow represent a defeasible connective as a conditional, and let $ev(\alpha)$ stand for “the event described in α ”; although $ev(\alpha)$ is quite similar to the notion of main eventuality *me* defined by Asher and Lascarides (2003), we do not commit to other assumptions of their theory.

Thus, some of the source domain knowledge about causation in (1) discussed above could now be represented as follows:

$$\langle \omega, \alpha, \beta \rangle \text{dies}(\text{connors}, \text{ev}(\beta)) \wedge \text{starves}(\text{mcenroe}, \text{connors}, \text{ev}(\alpha)) \rightsquigarrow \text{cause}(\text{ev}(\alpha), \text{ev}(\beta))$$

We can then infer in the source a *causation* relation between α and β if the event represented in α normally causes β :

$$\textbf{Causation: } \langle \omega, \alpha, \beta \rangle \wedge (\text{cause}(\text{ev}(\alpha), \text{ev}(\beta))) \rightsquigarrow \text{causation}(\alpha, \beta)$$

Note that ‘cause’ refers to the epistemic notion of one event causing another, whereas ‘causation’ refers to an inferred semantic relation between segments of discourse or, in other words, between semantic representation of events by means of DRSs. In order to include properties (and not only relations) in this framework, we assume a conceptualist point of view and consider that properties such as rate or value-judgement denote *concepts* (fast, slow, good, bad) which may correspond to the absolute rate in a commonsensical view of the world. Its representation in our semantic framework could be defined by adding an extra clause to the definition of DRS-formulae:

- If P is a property symbol and α and r are an episode label and a property label respectively, then $P(\alpha, r)$ is an DRS-formula (see Asher & Lascarides, 2003 for the complete definitions of DRS-formulae and SDRS construction).

Thus, a rule encoding contextual knowledge to infer rate in the source would look as follows (note that when considering event properties we only need to consider one DRS α in our rules, even though a discourse usually consists of one or more DRSs):

$$\langle \omega, \alpha \rangle \text{gulps}(\text{car}, \text{gasoline}, \text{ev}(\alpha)) \rightsquigarrow \text{fast}(\text{ev}(\alpha))$$

Supported by this rule we can then infer an event property in the source for its subsequent transfer to target via the Rate VNMA (when the Rate VNMA is instantiated):

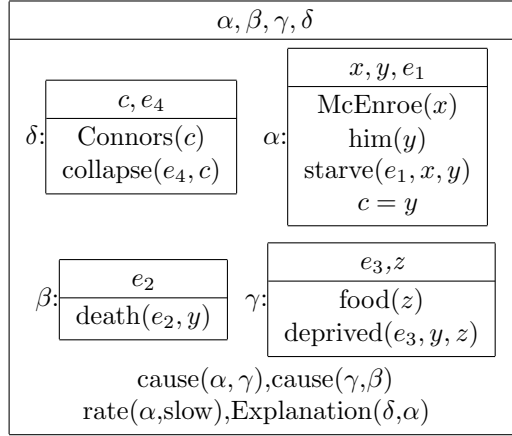
$$\textbf{Rate: } \langle \omega, \alpha \rangle (\text{fast}(\text{ev}(\alpha))) \rightsquigarrow \text{rate}(\alpha, \text{fast})$$

7.3 VNMA and Rhetorical Relations

We are now ready to extend the use the VNMA introduced in section 6 and the above points about source domain inferencing and contextual knowledge to offer SDRT-based semantic representations, based on the ATT-Meta approach to metaphor, for discourse examples. For simplicity of exposition, we leave out any details not directly relevant to the discussion on VNMA. Consider the following variation of (1):

(1b) Connors collapsed as McEnroe starved him to death.

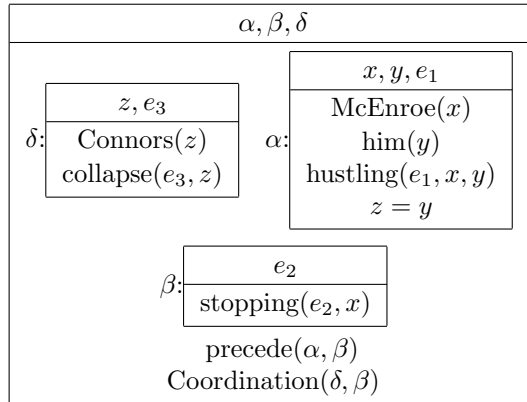
(1b) suggests that the cause or explanation of the collapsing of Connors is the “starving of Connors” which precedes Connors’s collapsing. Without going into many details (see Asher & Lascarides, 2003) the two main clauses of the discourse are linked by the Explanation rhetorical relation. This means that in order to interpret fully (1b) we need to take into account both the metaphorical aspects and its structure. Leaving aside the specific metaphorical meaning of ‘collapse’ and pro-nominal issues, the result of source domain reasoning for (1b) could be represented as follows:



Thus, inferring Explanation allows us to conclude that being starved to death *explains* Connors collapsing. The following example will allow us to show how our approach deals with coordination and temporal precedence discussed in examples (7) and (9):

(1c) Connors collapsed and McEnroe stopped hustling him.

In terms of discourse structure, we follow Gomez-Txurruka's (2003) approach: 'and' marks a Coordination relation between the conjuncts, blocking any other plausible interpretation of (1c) such as Result (the second conjunct will be the result of the first one). If we also consider the metaphorical analysis offered while discussing example (9), the result is the following semantic structure representing the conclusion of source domain reasoning:



Summarizing, the semantic framework outlined in this section consists of:

- (i) DRSs and SDRSs consisting of events, individuals, states, etc. They can be thought of as situations or as representation structures as in dynamic semantics. A context consists of one or more DRSs, DRSs relations and properties.
- (ii) Event relations and properties such as *causation*, *rate*, *time-order*, etc inferred in the source domain for the systematic transfer of certain type of information conveyed by metaphorical utterances. The transfer of this type of information via VNMA is a contribution of the ATT-Meta approach to metaphor interpretation (Barnden et al., 2003; Wallington et al., 2006).

- (iii) Rhetorical relations to address the structure of discourse and provide a more complete analysis of metaphor occurring in discourse.

8 Concluding Remarks

This paper investigates the formalization and semantic representation of the ATT-Meta approach to metaphor interpretation combined with ideas of discourse structure theories and, in particular, of SDRT. We claim that inferring the discourse structure is not sufficient to analyze metaphors where a Coordinator discourse relation is inferred.

The ATT-Meta approach to metaphor interpretation is backed up by a powerful implementation that performs sophisticated reasoning to interpret metaphorical utterances. We have focused on description and formalization of several VNMA, mappings for the systematic transference of invariant aspects from source to target. We have shown how a dynamic semantic approach can be adapted for these purposes to offer an unified semantic representation of ATT-Meta’s view of metaphor interpretation.

Map-transcending entities pose a problem for several analogy-based approaches to metaphor interpretation, both from a computational and a theoretical point of view. With respect to the computational approaches, theories of metaphor interpretation based on analogy (Falkenhainer, Forbus, & Gentner, 1989; Holyoak & Thagard, 1989) usually require a conceptual similarity between the source and the target domains. Map-transcending entities need to be mapped by extending on the fly the metaphorical views with new correspondences. We have argued that this strategy is both computationally expensive and in some cases, plainly impossible.

Formal semantic approaches (Asher & Lascarides, 2003) do not account for metaphorical utterances including map-transcending entities. Other works (Hobbs, 1996; Martin, 1990; Narayanan, 1997) have addressed source domain reasoning to a limited extent, but its role in metaphor interpretation has not previously been adequately investigated. Moreover, map-transcending entities pose a problem for analogy-based approaches to metaphor interpretation (Falkenhainer et al., 1989), which usually require a conceptual similarity between the source and the target domains.

Acknowledgments

The author is currently employed by the EPSRC EP/C538943/1 project in the Figurative Language group (John Barnden, Mark Lee and Alan Wallington) at the School of Computer Science, University of Birmingham, UK; a substantial part of the work has been done in collaboration with my colleagues.

To Isabel Gómez Txurruka for her influence and help both on my beginnings *and* on my current work.

References

- Agerri, R., Barnden, J., Lee, M., & Wallington, A. (2007). On the formalization of invariant mappings for metaphor interpretation. In *Proceedings of ACL 2007 Demo and Poster Sessions*, pp. 100–113, Prague. Association for Computational Linguistics.

- Asher, N., & Lascarides, A. (2001). The semantics and pragmatics of metaphor. In Bouillon, P., & Busa, F. (Eds.), *The Language of Word Meaning*, pp. 262–289. Cambridge University Press.
- Asher, N., & Lascarides, A. (2003). *Logics of Conversation*. Cambridge University Press.
- Barnden, J., Glasbey, S., Lee, M., & Wallington, A. (2003). Domain-transcending mappings in a system for metaphorical reasoning. In *Companion to the 10th EACL*, pp. 57–61.
- Barnden, J., & Lee, M. (2002). An artificial intelligence approach to metaphor understanding. *Theoria et Historia Scientiarum*, 6(1), 399–412.
- Falkenhainer, B., Forbus, K., & Gentner, D. (1989). The structure-mapping engine: algorithm and examples. *Artificial Intelligence*, 41(1), 1–63.
- Glasbey, S., Barnden, J., Lee, M., & Wallington, A. (2004). Temporal metaphors in discourse. Tech. rep. CRSP-04-03, School of Computer Science, University of Birmingham. <ftp://ftp.cs.bham.ac.uk/pub/tech-reports/2004/CSRP-04-03.ps.gz>.
- Gomez-Txurruka, I. (2003). The natural language conjunction ‘and’. *Linguistics and Philosophy*, 26(3), 255–285.
- Hobbs, J. (1996). An approach to the structure of discourse. In *Discourse: Linguistic, Computational and Philosophical Perspectives*. www.isi.edu/hobbs/discourse-references/discourse-references.html. Probably not to be published.
- Holyoak, K., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13(3), 295–355.
- Lakoff, G. (1993). The contemporary theory of metaphor. In Ortony, A. (Ed.), *Metaphor and Thought*, 2nd edition. Cambridge University Press, Cambridge (MA).
- Leakey, R., & Lewin, R. (1992). *Origins Reconsidered*. Little, Brown and Company, London.
- Martin, J. (1990). *A computational model of metaphor interpretation*. Academic Press, New York.
- Narayanan, S. (1997). *KARMA: Knowledge-based action representations for metaphor and aspect*. Ph.D. thesis, Computer Science Division, EECS Department, University of California, Berkeley.
- Wallington, A., Barnden, J., Glasbey, S., & Lee, M. (2006). Metaphorical reasoning with an economical set of mappings. *Delta*, 22(1).