On the formalization of Invariant Mappings for Metaphor Interpretation

Rodrigo Agerri, John Barnden, Mark Lee and Alan Wallington
School of Computer Science, Univ. of Birmingham
B15 2TT Birmingham, UK
r.agerri@cs.bham.ac.uk

Abstract

In this paper we provide a formalization of a set of default rules that we claim are required for the transfer of information such as causation, event rate and duration in the interpretation of metaphor. Such rules are domain-independent and are identified as invariant adjuncts to any conceptual metaphor. We also show a way of embedding the invariant mappings in a semantic framework.

1 Introduction

It is generally accepted that much of everyday language shows evidence of metaphor. We assume the general view that metaphor understanding involves some notion of events, properties, relations, etc. that are transferred from the source domain into the target domain. In this view, a metaphorical utterance conveys information about the target domain. We are particularly interested in the metaphorical utterances that we call *map-transcending*. Consider the following example:

(1) "McEnroe starved Connors to death."

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. 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). 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 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're spinning our wheels."

It is not very clear what could be a target correspondent for 'wheels'. We have developed an AI system called ATT-Meta for metaphor interpretation (Barnden et al., 2002) that employs reasoning within the terms of the source domain using various sources of information including *world* and *linguistic knowledge*. The reasoning connects unmapped ideas used by utterances, such as wheels and starving, to other source-domain ideas for which a mapping is already known. These known mappings may be constituents of particular metaphorical view, but previous work (Barnden et al., 2003; Wallington et al., 2006) has

109

shown evidence that there are metaphorical aspects (such as causal relations between events) that, subject to being called, invariantly map from source to target (we call these mappings View-Neutral Mapping Adjuncts or VNMAs) irrespective of whatever specific metaphorical views are in play. These allow many mapping effects, which would otherwise have to be duplicated across all view-specific mappings, to be factored out into separate mappings. In our approach, source domain reasoning takes place in a special, protected computational context that we call the "pretence space". We use the term 'reality' to refer to the space outside the pretence where propositions are about reality as the understander sees it.

Currently ATT-Meta implements the VNMAs by including them in view-specific rules, but we plan to make the system more modular and its view-specific mappings more economical by implementing VNMAs as separate default rules. The first step towards that goal is to provide a formalization of these mappings and to show their role in metaphor interpretation. In order to do so, we provide a semantic representation of how these VNMAs work by adopting Segmented Discourse Representation Theory (Asher and Lascarides, 2003) to capture the main aspects of the ATT-Meta approach.

2 Knowledge and Inference

If (1) is being used metaphorically to describe the result of a tennis match, a plausible target interpretation would be that McEnroe defeated Connors in a slow manner by performing some actions to deprive him of his usual playing style. Assuming a commonsensical view of the world, a within-pretence meaning would be that McEnroe starved Connors to death in the real, biological sense. The inferencing within the pretence can then conclude that McEnroe *caused* Connors's death by *depriving* or disabling him. Leaving some details aside, the partial logical form (in the pretence) 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 also use the notion of event 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 reality would correspond to the event of x starving y to death in the pretence. 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) $\exists x, y, z, e_1, e_2, e_3(McEnroe(x) \wedge Connors(y) \wedge food(z) \wedge starve(e_1, x, y) \wedge death(e_2, y) \wedge deprived(e_3, y, z) \wedge cause(e_1, e_3))$

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 within-pretence 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 concept in the pretence related to the progress rate of *starving*. Now, the existing mapping DEFEAT AS DEATH can be applied to derive, outside the pretence, 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. We appear to have a problem also to map the slow progress *rate* of a process like starving.

3 VNMAs in a Semantic Framework

In the ATT-Meta approach to metaphor interpretation, the mappings of *caused* and *rate* discussed above are accomplished by a type of default mappings that we specify as VNMAs (the Causation and Rate VNMAs, respectively; see (Wallington and Barnden, 2006) for an informal but detailed description of a number of VNMAs). The idea is that there are relationships and properties (causation, rate, etc.) between two events or entities that identically transfer from the pretence to the reality. We use the $\mapsto$ symbol to express that this mapping is a default. The VNMAs involved in the interpretation of (1) can be represented as follows:

Causation: $\forall e_1, e_2(cause(e_1, e_2)_{pret} \mapsto cause(e_1, e_2)_{rlt})$

The Rate VNMA transfers the qualitative rate of progress of events in the source domain to the qualitative rate of progress of its mappee:

Rate: $\forall e, r(rate(e, r)_{pret} \mapsto rate(e, r)_{rlt})$

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 the semantics of metaphor. In the somewhat simplified discussion on the within-pretence reasoning and mappings necessary to interpret metaphorical utterances such as (1), we have been using various sources of information 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 reasoning in the pretence; c) Relations and properties between events such as *causation* and *rate* that are inferred in the pretence; d) VNMAs that transfer within-pretence event relations and properties to reality.

There are two prominent computationally-oriented semantic approaches (Hobbs, 1996) and (Asher and Lascarides, 2003) that take into account contextual and linguistic information and stress the importance of relations between text segments in discourse interpretation. In fact, the incorporation of the above types of information ties in well with the SDRT (Asher and Lascarides, 2003) view of language understanding. For example, we can think of the pretence space as a Segmented Discourse Representation Structure (SDRS) representing the result of within-pretence inference which can be mapped by using various view-specific and invariant mappings to reality. In other words, we can see the pretence SDRS as the input for what the ATT-Meta system does when interpreting metaphor – it will reason with it, producing an output of inferred reality facts which we may also represent by means of an SDRS. The result of reasoning in the pretence to interpret (1) would now looks as follows:
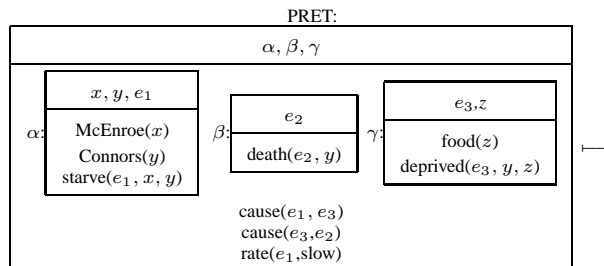
where α and β are labels for DRSs representing events, PRET for a pretence space and $\longmapsto$ mappings (VNMAs and central mappings) needed in the interpretation of the metaphorical utterance. Importantly, the VNMAs would pick upon aspects such as causation and rate from pretence to transfer them to reality producing an output which could also be represented as a SDRS:
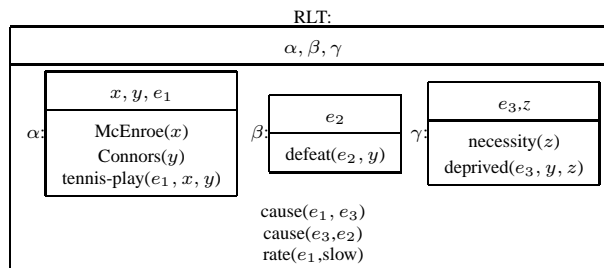
Note that this formal representation integrates the systematicity of mapping invariantly certain aspects of metaphorical utterances by formulating them as relations between events that can be represented as

relations and properties of DRSs. For this purpose we need to modify the construction rules of SDRSs to be able to infer properties and relations involving individuals and not only DRSs' labels. In addition to this, we have shown in the previous section how ATT-Meta source domain reasoning captures the interaction of the various sources of knowledge used to infer causation and rate in the pretence. Furthermore, studying the interaction between VNMAs and discourse relations may allow us to extend the study of metaphor to discourse.

4   Concluding Remarks

Following the ATT-Meta claim metaphors often convey crucial information via VNMAs, we can re-analyze example (1) so that the effects of the NE-CESSITIES AS FOOD mapping are obtained by VNMAs. In the pretence, the food is something Connors needs for proper functioning: i.e., it is necessary that Connors have the food in order to function properly. The necessity here is covered by the Modality VNMA, which maps relative degrees of necessity, possibility, obligation, etc., from pretence to reality. Moreover, the functioning properly would be covered by the Function and Value-Judgement (levels of goodness, importance, etc. map identically to levels of goodness, etc.). So all that is left is the possession which could be covered by a STATE AS POSSESSION mapping.

Formal semantic approaches (Asher and Lascarides, 2003) do not account for metaphorical utterances including map-transcending entities. Other works (Carbonell, 1982; Hobbs, 1990; 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 require the discovery of an elaborate structural similarity between the source and target domains and/or the imposition of un-mapped source domain structures on the target domain, whereas part of our approach is that the un-mapped source domain structure introduced by the utterance is by default not carried over.

References

Nicholas Asher and Alex Lascarides. 2001. The semantics and pragmatics of metaphor. In P. Bouillon and F. Busa, editors, *The Language of Word Meaning*, pages 262–289. Cambridge University Press.

Nicholas Asher and Alex Lascarides. 2003. *Logics of Conversation*. Cambridge University Press.

John Barnden, Sheila Glasbey, Mark Lee, and Alan Wallington. 2002. Reasoning in metaphor understanding: The att-meta approach and system. In *19th Conference on Computational Linguistics (COLING-2002)*.

John Barnden, Sheila Glasbey, Mark Lee, and Alan Wallington. 2003. Domain-transcending mappings in a system for metaphorical reasoning. In *Conference Companion to the 10th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2003)*, pages 57–61.

Jaime Carbonell. 1982. Metaphor: An inescapable phenomenon in natural-language comprehension. In W. Lehnert and M. Ringle, editors, *Strategies for Natural Language Processing*, pages 415–434. Lawrence Erlbaum, Hillsdale, NJ.

BrianFalkenhainer, Kenneth Forbus, and Dedre Gentner. 1989. The structure-mapping engine: algorithm and examples. *Artificial Intelligence*, 41(1):1–63.

Jerry Hobbs. 1990. *Literature and Cognition*. CSLI, Lecture Notes, Stanford.

Jerry Hobbs. 1996. An approach to the structure of discourse. In D. Everett, editor, *Discourse: Linguistic, Computational and Philosophical Perspectives*.

James Martin. 1990. *A computational model of metaphor interpretation*. Academic Press, New York.

Srini Narayanan. 1997. *KARMA: Knowledge-based action representations for metaphor and aspect*. Ph.D. thesis, Computer Science Division, EECS Department, University of California, Berkeley, August.

Alan Wallington and John Barnden. 2006. Similarity as a basis for metaphor: Invariant transfer and the role of VNMAs. Technical Report CSRP-06-02, School of Computer Science, Univ. of Birmingham, December.

Alan Wallington, John Barnden, Sheila Glasbey, and Mark Lee. 2006. Metaphorical reasoning with an economical set of mappings. *Delta*, 22(1).