

Towards a Grounded Model for Ontological Metaphors

Sushobhan Nayak

Indian Institute of Technology, Kanpur, India

snayak@iitk.ac.in

Abstract

A novel approach to learning metaphors without any prior knowledge is proposed, in which ideas are acquired as concrete concepts and later on develop their abstraction. A grounded model of linguistic concepts and a hierarchical probability map is used to interpret/generate ontological metaphors.

1 Introduction

Consider the following sentences:

- You are *wasting* my time. (TIME IS MONEY)(Lakoff and Johnson(1980))
- We need to *combat* inflation. (INFLATION IS AN ENTITY)(Lakoff and Johnson(1980))

These usages are so ingrained in everyday conversation that we fail to recognize that they are actually metaphors that try to describe abstract concepts in terms of concrete ones(Lakoff and Johnson(1980)). A proper understanding of language and thought therefore calls for an increased and focused research into metaphors and the way they are acquired, interpreted and generated.

As such, there have been many attempts at interpreting metaphors over the years. The earlier works like Wilks(1975); Fass(1991) are based on a violation of selectional restrictions in a given context. Kintsch(2000) effectively uses LSA to interpret metaphors like “My lawyer is a shark”. However, these models are incapable of handling Lakoff and Johnson(1980))’s view of metaphors. The works that are closest to the modern view, in their attempt at interpreting common text from financial or other domains, encompass Narayanan(1997); Narayanan(1999). Shutova et al.(2010) also show metaphor paraphrasing using noun-verb clustering. However, they both need a hand-coded seed set or metaphor repository to do further learning. The model proposed here assumes no prior knowledge. From

multimodal input, we first ground some basic concepts and using them, exploit the language syntax to learn/interpret/generate metaphorical mappings in a natural way that emulates language learning in an early learner.

2 Proposed Model

The models till date have looked at metaphors as something that is acquired/interpreted differently than common language. However, the treatment by Lakoff and Johnson(1980)) suggests that we should look at metaphor acquisition as we look at language acquisition, and not something that is interpreted/acquired *after* we have acquired the basic nuances of the language. There is ample evidence in literature to suggest that basic linguistic forms can be grounded(Roy and Reiter(2005)). Consider an early learner who has acquired the grounded concepts of the verb *impale* and understands that, based on its experience till date, only a living entity is capable of executing it. Then it comes across this:“*Arrogance impaled him.*” Now based on this linguistic input only, in the absence of any other physical stimuli, what is to stop the learner from interpreting ‘arrogance’ as a living entity? It is only when it comes across other usages of the concept ‘arrogance’ that its initial idea that ‘arrogance is an entity’ might be modified to ‘arrogance is an abstract concept’. However, as one might notice, ‘arrogance is an entity’ is actually a well established metaphor. This leads us to look at metaphor acquisition in a different light. Why look at ‘arrogance’ etc. as abstract concepts that *need* to be understood through grounded concepts *later*? Why not look at them as *grounded* concepts, which *later* acquire the abstractness, there by being imparted with metaphorical mappings, which is suggested by the example alluded to before? The proposed model takes this approach, where we start with grounding some very basic verbs/nouns

and then go on to acquire/interpret/generate ontological metaphors (a metaphor in which abstract notions are projected through concrete concepts of objects/entities/substance etc., i.e. some physical entity.), just as we would do any other words/concepts. In the discussion that follows, it's assumed that language representation incorporates Langacker(1987)'s *image schema*, and that language understanding incorporates Embodied Construction Grammar(ECG) (Bergen et al.(2004)).

2.1 Grounded Forms

It is more or less established in literature that linguistic concepts are cognitively characterized in terms of *image schemas*, which are schematized recurring patterns from the embodied domains of force, motion, and space(Langacker(1987); Lakoff and Johnson(1980)). Before we go into learning ontological metaphors, we create a grounded system that helps modify/nurture image schemas of new concepts as they come along.

Almost all types of ontological metaphors come under three broad categories of an Object, a Substance or a Container. In fact, these concepts emerge directly for an early learner through physical experience(Lakoff and Johnson(1980)). A probable scenario for an artificial learner can be to try grounding concepts from multimodal inputs of image, sound and written transcripts. Consider the work by Mukerjee et al.(2011), where they try to discover coreference using image grounded verb argument structure. From multimodal input of a video, and associated narration, they have been able to learn the verb structure CHASE(A,B). This can be further extended to derive the image schemas of CHASE and the actors A and B. Mukerjee et al.(2011) use velocity features of the objects in the video to unsupervisingly cluster them to find the cluster of CHASE(A,B). This cluster that contains much of the velocity related information for concept CHASE can be presented as an image schema for CHASE(A,B); and since CHASE is a two-party interaction, learning CHASE means the concept of agents A and B are also learned. In the present model, the centroid, maxima, and minima of the feature cluster have been stored to represent CHASE¹, whereas

¹This is just the chosen representation. One might well like to store the whole cluster for representing the action. The claim here is not that the present representation is the one that is actually there; the idea is just to demonstrate that even such a crude model works.

A and B are being treated as *point objects*, i.e. entities whose behavior would remain unchanged if they are replaced by geometrical points in the visual space they act in. Storing maxima/minima also helps us create image schemas of adjectives SLOW() and FAST(). In this model, these functions take an action, say CHASE, as argument, and output the minimum or maximum of the action's velocity feature. With this sort of grounding at hand, we can handle the mental simulation part in ECG(Bergen et al.(2004)) used for understanding linguistic occurrences. A *fast(chase)* would mean the simulation runs with the image schema of CHASE, with the velocity features being maintained at their maximum.

We now have grounded forms of CHASE(A,B), SLOW(),FAST() and entities(point-objects). We next take cues from Mukerjee and Sarkar(2007) and learn IN(A,B), OUT OF(A,B), INTO(A,B) and the 'container'. Mukerjee and Sarkar(2007) use Voronoi-based features to distinguish space into the interior or exterior of an object. In this model, the image schemas of IN(A), OUT OF(A), INTO(A) and the 'container' are interpreted as being interconnected. While the boundaries of the object-container in the visual input are taken as the boundaries of concept 'container', IN(A,B) is represented by substance/entity A inside container B. INTO(A,B) and OUT OF(A,B) are schemas in two states, where the object/substance A is in/out of the container B in one state, and changes its position in the other. Substance is crudely grounded as something that can't be represented by a point object, i.e. something that is not executing rigid body motion. Essentially, combined with, say INTO(A,B), if in the motion schema simulation of action INTO(A,B), A can't be represented by a point, it is taken as a Substance. This further allows us to ground adjectives MORE(A) and LESS(A) for Substances, based on the change they bring about in the volume of the Voronoi interior.

To reiterate, the objective of this section is not to claim that a proper image schema has been developed for the above concepts. The goal here has been to show that even from a simple multimodal input like a video and the associated commentary, an intelligent agent can get a crude grounded model of linguistic concepts. This can only mean that an early human learner will be much better at this job. To summarize, the model has at its disposal, the grounded concepts of Entity/Object,

Substance, Container, CHASE(A,B), INTO(A,B), IN(A,B), OUT OF (A,B), MORE(A), LESS(A), SLOW() and FAST(). However, these alone are insufficient to show how metaphors are acquired ('chase' verb is sparsely used in general literature), and we therefore *assume* the availability of GIVE(AgentA, AgentB, Entity/Substance C), SOME(Substance A) and SPEND(Substance A) hereafter.

2.2 Concept Acquisition

The vast majority of our vocabularies are learned later purely from the linguistic input (Bloom(2000)). The goal from here-on will thus be to acquire language concepts with the help of the aforementioned grounded concepts and a text corpus. To determine how far language usage alone can help shape the concept of metaphors, we compiled a list of sentences from Lakoff and Johnson(1980)) and Lakoff et al.(1991) that correspond to the metaphor-mappings for Containers, Objects and Substances. The salient findings are:

- Of the 85 sentences denoting Container metaphors, in 65, the abstract idea was imparted the image schema of a container based only on the prepositions *in/out*. In the rest 20, adjectives (*full, empty*) and verbs (*explode, erupt, fill*) took the mantle.
- In all of the 63 sentences for Object metaphor, the Object property was imparted to the concept because VERB(A,B) took object arguments, i.e. verbs were the primary basis of metaphor mapping.
- Of the 42 sentences for Substance metaphors, 17 mappings were done based on adjectives (*more, less*) while the rest were of the type Container contains Substance, i.e. first the Container property was imparted, and then whatever was supposed to be inside the container was called a substance.

Based on this observation, we construct a model of concept acquisition which incorporates the following bold (unproven) claims:

Claim 1 Verbs, adjectives, nouns and prepositions all play roles in concept acquisition and have varying importance in different forms of concept acquisition

Claim 2 There are only a limited representative verbs/adjectives that are grounded (i.e. have stable and distinct image schema), and the rest import their schemas and modify them to suit themselves. So, this structure is hierarchically organized.

Claim 3 They can be represented in terms of a probability map, whence we can get an idea of the interplay between different concepts.

Claim 4 Abstract concepts are acquired as concrete concepts first, just like any other grounded concept, and later they acquire their abstraction due to emergence of future evidence.

Claim 1 is supported by the observation that precedes it. Claim 2 is somewhat self-explanatory, and better understood through examples. Consider verbs *impart, provide, shower, bombard, donate* etc. A close look will reveal that they can all be derived from GIVE(). For instance, SHOWER() (as in 'shower somebody with praises') can be a combination of verb GIVE() and adjective MORE(). In fact this kind of representation seems more memory-efficient. If we are required to store image schemas of the millions of words that we come across, our memory will be a mess. Storing only a select few and combining them to derive the others is a more structured, systematic and efficient mechanism. Furthermore, we should also notice that we had to take help of an adjective to describe a verb, which reinforces Claim 1. The first level of hierarchy consists of these grounded forms which are distinguishable. The second level is the derived one that draws from all the nodes of the first level. Claim 3 asserts that this interplay can be represented through a probability map. For instance, for a single verb GIVE(), the adjectives can be assigned probabilities based on how frequently they modify GIVE() to produce an understandable image schema, so that we have an idea of which ones are more probable of appending to the verb when a new concept involving the same emerges. We will later see how this map can incorporate many aspects of the metaphor acquisition task.

2.2.1 The Model

We now describe the metaphor acquisition model based on Claim 4. We first have a repository of grounded concepts. Then as the learner is exposed

to more sentences, the sentences are searched for contexts similar to the ones already learned. The noun arguments, which might be new concepts, are assigned a dynamic probability of belonging to one of the classes of Object/Substance/Container. With more evidence, these probabilities are modified within a reward/penalty scenario. All concepts are treated concrete unless evidence to the contrary crops up.

The model is better understood through examples. Let the learner come across the sentence – ‘I can’t give you much time’ \equiv GIVE(MUCH(time)). Now MUCH() takes a Substance as an argument. So time is assigned the schema of Substance with probability 1. Then GIVE(time), which takes either a Substance or an Object as its argument, dynamically changes the probabilities to 2/3 for Substance and 1/3 for Object. When it further comes across “In time, you will understand”, i.e. IN(time), the probabilities are modified to 1/4, 2/4 and 1/4 for container, substance and object respectively. This assignment helps us in two ways – firstly, it prevents us from exclusively assigning the concept to any single class, thereby allowing us to model metaphors that contextually take up the properties of different classes. And secondly, it also gives us an idea of the affinity of the abstract concepts for different classes. To avoid confusion, as of now, from the above example, we can only assume that time is a concrete concept that has properties of Object, Container and Substance ².

ACQUISITION: We next tackle the problem of distinguishing between abstract and concrete concepts. Consider an early learner who comes across the following:

- Llama is a four-legged animal.
- Anger is a red-eyed demon.

How does the subject distinguish between concrete Llama and abstract Anger? One way to look at it would be that as soon as the learner comes across these, it incorporates the features in the image schema of the concept. The schema field is then searched for possible conflicts. If two properties are in conflict, they are brought to the CONTEST field in the image schema, where ‘voting’

²Which creates no conflict since Substances can act as agents/point-objects. Similarly containers and objects are interchangeable based on the context.

takes place between the conflicting scenarios. Voting is done to take care of two possibilities. First, the conflict that arises might be due to a false evidence. If one of the properties is discarded based on some false evidence, the schema might become erroneous. So both are kept, but they are assigned probabilities of expected occurrence. Secondly, it gives room for metaphorical descriptions. For instance, suppose Anger has been assigned Object and Substance properties before this occurrence. Now when it acquires the schema of ‘red-eyed demon’, the properties it assumes are, say, physical appearance and adjectives describing a demon. The conflict arises between physical appearance and Substance (because a substance can’t have eyes). So they are brought to the CONTEST field. Based on how often these concepts occur in the corpus, they are assigned probabilities. For example, in this case, the probability previously assigned with Substance property is converted to the equivalent vote and red-eye is given vote 1. On next occurrence of ‘anger is a substance’, we follow a reward/penalty scenario ³, where the vote of Substance is increased by one, and that of ‘red-eye’ is reduced by one. When the vote of one concept is reduced to zero, the CONTEST field is cleared of the two. Assuming that false alarms are very less in number compared to correct usages, this process will reach stability. One might also note that this process doesn’t in anyway harm the objective of the second sentence. The idea that it wanted to convey about Anger remains there inside Anger’s image schema in the form of the adjectives – only the physical appearance schema is eliminated, which is the ideal behavior. Given enough time to settle down, the concrete concepts would thus have some sort of physical characteristic in their image schema that is NOT derived from Object/Substance/Container. The abstract concepts, on the other hand, would only be linked to the basic schemas, without any distinguishing and particular physical characteristics. For example, ‘Box is a container’ and ‘Love is a Container’. Both will imbibe concepts of enclosure, boundary etc. from the container. But BOX would have additional schemas of a lid, wooden material etc. (which would actually vary subjectively). Since the concept of ‘Love is a Container’ is ingrained

³It is to be noted that this scenario is only followed for schemas that are in the CONTEST class, not for all the schemas. This prevents unnecessary removal of non-conflicting schemas.

in the learner, we will say that the subject has *acquired* the metaphor. This method of metaphor acquisition eliminates any argument regarding a violation of selectional restrictions and the need for basic seed metaphors to understand others. In this scenario, metaphors are learned like any other linguistic concept. The ‘concrete first-then abstraction’ should thus score over ‘abstract first-concrete grounding’ approach.

INTERPRETATION: This representation also helps in understanding of metaphorical occurrences. Previous works like KARMA(Narayanan(1997)), when they come across a metaphorical occurrence, search in a repository of metaphorical mappings to understand the statement. Understanding in the present system can naturally flow through the ECG approach. The ECG asserts that when we are interpreting ‘Harry fell in water’, we actually simulate Harry falling in water to understand the utterance. In the present approach, the occurrence ‘Harry fell in love’, when simulated, will behave like this – the sentence will first be converted to concept FALL(Harry, love), then using FALL(Object,Container), HARRY and LOVE would import those schemas. HARRY would also import its own physical characteristic schemas while LOVE would have no such schema. Then motion schema of FALL() would be brought along and these three will be composited to produce the final simulation. Metaphorical mapping would thus be understood as any other linguistic concept, the only distinguishing factor would be that while concrete concepts would bring in their associated physical properties, abstract concepts will be described by a bare-bones schema. The idea of assigning probabilities of a concept’s association to different base classes helps us in another elegant way. Previously, to understand ‘Harry is in love’ and ‘Love led to his demise’, the models had to invoke two different metaphorical mappings of ‘Love is a container’ and ‘Love is an entity’. Whereas, in the present model, the concept Love has already been assigned to both the Container and Object class, and based on the context one of the assignments gets highlighted. This reduces the memory inefficient and crude method of having a repository of metaphor-maps.

GENERATION: Claim 3 helps us generate new metaphors or use established metaphors just as natural language is used, without conscious effort.

Once the system needs to convey an abstract idea, it has at its disposal a probability map through which the idea is connected to other concrete concepts. Those concepts are further connected to verbs/adjectives etc. with certain other probabilities. A path through this map, which can represent a coherent structure, would lead to a metaphorical mapping. Which metaphorical mapping is more culturally accepted would of course depend on the rating of the path (in this case, simply the multiplication of probabilities). This model however would be extremely good in interpretation of newly created metaphors. The understanding process would just involve simulating the ideas by ECG. The older models, since they rely on a metaphor list, would have a hard time understanding new metaphors because they might not fit into the established scenarios.

3 Experiments

The Brown Corpus was used to test the ideas and derive some possible metaphorical mappings. All the occurrences of the grounded concepts, viz. CHASE(A,B), INTO(A,B), IN(A,B), OUT OF(A,B), MORE(A), LESS(A), SLOW() and FAST() were found out and the sentence structure was converted to these functional structures using a very crude method – the first occurrence of a singular or mass noun(NN) in the tagged corpus was assigned to the concept. For example, the sentence fragment ‘into a hot cauldron’ is converted to INTO(cauldron). Using this very basic method, some of the possible metaphor mappings that were found were:

Container The following concepts were common between IN(), INTO(), OUT OF(), leading to a strong Container metaphorical map—*future battle fight mission darkness violence chaos silence water mind religion language*

Substance The following concepts were common between GIVE() and MORE(), leading to a strong affinity for Substance mapping: *affection information emphasis interest protection time*

To have a flavor of how an abstract concept is connected to the base classes, we examined all occurrences of noun LOVE in the corpus. While 90% of the time it acted like Object/Substance, 10% of the time, it acted like container, hinting

that the affinity of Love for Container is minimal.⁴ The Object and Substance cases are almost indistinguishable except for when substance-specific adjectives like ‘more/less’ are used. Otherwise, Love is considered a physical material, and it is not usually distinguishable whether its an Object or a Substance.

To look into how deeper mappings like ‘Time is Money’ might be deciphered, we also looked into SPEND(), and WASTE(). 60% of usages of SPEND() were SPEND(Time)(in various forms like day/year etc.), while the rest were SPEND(Money), with very minimal(two or three occurrences of the 200 odd) SPEND(*other substance*). Similarly, of the 10 occurrences of verb WASTE(), 9 are concerning Time and the rest concerns Substance. The trend also points to one important assumption we have made – that is, abstract concepts are first learned as concrete. As we see, these verb usages correspond to abstract concepts much more readily than they do to their concrete counterparts. So it’s but natural on the part of an early learner to assume them as concrete ideas.

4 Conclusion and Future Work

While the above description pointed towards a new approach of handling metaphors that is closer in spirit to the view that metaphors are an integral part of thought and language usage and not just poetic devices, the work might still look incomplete. This is so because even if the basic ideas and claims have been supported, the system is still not fully functional. To be precise, as of now we have the grounded concepts described in Section 2.1 and based on that extremely small test set, we have tried to learn some metaphorical mappings. As more concepts are grounded in some way or other, the system will be better equipped to handle other mappings.

The ultimate aim would be to finally simulate all this in an ECG framework to show that the model is capable of emulating human behaviour. However, this must again be reiterated that the aim of this paper was not to show a working model fully capable of handling ontological metaphors, which is under construction, but that the new approach might be better and more natural than previous works that depended on hand-coded knowl-

⁴The container metaphor arose almost exclusively in the usage ‘in love’.

edge in some form or other.

Acknowledgement

I would like to thank Dr. Amitabha Mukerjee for his constant guidance throughout this work.

References

- Bergen, B, N Chang, and S Narayan. 2004. Simulated action in an embodied construction grammar. In *Proc. of the 26th Annual Meeting of the Cognitive Science Society*.
- Bloom, Paul. 2000. *How Children Learn the Meanings of Words*. MIT Press, Cambridge, MA.
- Fass, D. 1991. Met*: A method for discriminating metonymy and metaphor by computer. *Computational Linguistics* 17(1):49–90.
- Kintsch, W. 2000. Metaphor comprehension: A computational theory. *Psychonomic Bulletin and Review* pages 257–266.
- Lakoff, George, Jane Espenson, and Alan Schwartz. 1991. *Master metaphor list: 2nd draft copy*.
- Lakoff, George and Mark Johnson, editors. 1980. *Metaphors We Live By*. University of Chicago Press.
- Langacker, RW, editor. 1987. *Foundations of Cognitive Grammar I: Theoretical Prerequisites*. Stanford University Press.
- Mukerjee, Amit, Kruti Neema, and Sushobhan Nayak. 2011. Discovering coreference using image-grounded verb models. In *Proceedings of RANLP*.
- Mukerjee, Amitabha and Mausoom Sarkar. 2007. Grounded perceptual schemas: Developmental acquisition of spatial concepts. In *Spatial Cognition V Reasoning, Action, Interaction*, Springer Berlin / Heidelberg, volume 4387, pages 210–228.
- Narayanan, S. 1997. *Knowledge-based Action Representations for Metaphor and Aspect (KARMA)*. Ph.D. thesis, CS, UC Berkeley.
- Narayanan, S. 1999. Moving right along: A computational model of metaphoric reasoning about events. In *Proc. of NCAI*. pages 121–129.
- Roy, Deb and Ehud Reiter. 2005. Connecting language to the world. *Artificial Intelligence: Special Issue on Connecting Language to the World* 167:112.
- Shutova, E, L Sun, and A Korhonen. 2010. Metaphor identification using verb and noun clustering. In *Proc. of COLING*. pages 1002–1010.
- Wilks, Y. 1975. A preferential pattern seeking semantics for natural language inference. *Artificial Intelligence* 6(1):53–74.