The Knowledge Engineering Review, Vol. 31:3, 261–277.
doi:10.1017/S0269888916000060

Data fusion and abductive inference for metaphor resolution: a bridging discussion

GIOVANNI FERRIN, LAURO SNIDARO and GIAN LUCA FORESTI

Department of Mathematics, Computer Science and Physics, University of Udine, Via delle Scienze 206, 33100 Udine, Italy; e-mail: giovanni.ferrin@uniud.it, lauro.snidaro@uniud.it, gianluca.foresti@uniud.it

Abstract

Since the 1980s, metaphor has been recognized as a pervasively diffused phenomenon in communication, absolutely not restricted to rhetoric and linguistic phenomena, involving structured concepts, relations, and matching 'rules'. Metaphor resolution, that is metaphor understanding, as well as metaphor creation, has become an issue in automated processing and understanding of natural language as well as of mixed visual communication. It can be showed as a process of structure finding and mapping procedure between conceptual denotation–connotation structures necessary for interpretation. Creative abduction is then showed to be the pattern inference required to work out structure-mappings in corresponding nodes as present in metaphors. In this paper, we review some key issues (definitions, typologies, theoretical problems) involving the concept of 'metaphor' and survey some definitions and concepts emerging in contemporary debate on abductive inference. Finally, we argue that metaphor understanding process can be recognized as a fusion tractable problem, allowing the exploitation of frameworks and algorithms of such domain.

1 Introduction

With the advances in digital technologies and the advent of the Web, information generated, analyzed, stored, accessed and transmitted has exponentially increased in quantity and complexity. Most of this data not only are multimedia in nature, including digital images, video, audio, text data, hierarchically structured documents, web pages, and transmitted and delivered through multimodal communication channels, but also are largely user generated and therefore structured following both intersubjective and personal 'grammars' and semantic rules.

To make use of this vast amount of multimedia data, techniques have been developed to efficiently manage indexing, processing, interacting, and retrieving information over large multimedia repositories based on their content. Multimedia-processing deals with live spatio-temporal data which is redundant, dynamic in time, and often pervasively and contextually connected with other data, like visual data with audio, music, text, etc.

1.1 Content understanding

The challenge, however, is content extraction, that is generating a (possibly) correct interpretation of what is being represented. Figure 1 demonstrates that 'understanding', an image in this case, is not just a matter of recognizing persons and other objects, but requires moving through different levels: from the level of bare data, collected in sensing processes, to the levels of structured features, properties, objects, relations, events, and situations.

At all different levels, any process faces uncertainty problems: real world is 'abundant' (Feyerabend, 1999) but its representations necessarily leave out a lot: translating real objects into numeric data only

Figure 1 Image interpretation—climbing a mountain of paperwork?

'some' elements are considered (e.g. three-dimensional (3D) to two-dimensional (2D) projection) and translating back numeric data to recognized objects (thus generating information) only 'relevant' elements are considered (e.g. through feature selection) and this constitutes already an *a priori* interpretation.

Many different approaches for handling multimedia data exist. On one hand, based on implicit or explicit ways of modelling reality and on the other hand using a wide range of algorithms and tools, from purely logical ones (first-order logic based) to probabilistic Bayesian ones. Rule based or logic approaches, based for example on ontologies, better model very restricted domains. Probabilistic approaches on the other hand show human 'comparable' results in some domains: real world seems to be self-structuring, structures can show relations (like causal ones) and can be discovered (Kemp & Tenenbaum, 2008).

Focussing our attention on a natural language fragment like a dialog or on any other form of linguistic interaction between more than two persons, if we are trying to understand the topic of the conversation and possibly use the 'extracted knowledge', we have to take under consideration the main 'meaning-bearing' features or phenomena. General linguistics, semiotics, formal semantics frameworks, and tools succeed in understanding and accounting for purely linguistic phenomena like implicature, presupposition, focus (Rooth, 1985; Van der Sandt, 1992; Kadmon, 2001). Nevertheless, the 'fusion' of linguistic content and visual analysis, where each of the two modes constitute consistent context of the other, is definitely necessary in dealing with other linguistic phenomena like indexicals, or some definite descriptions, and paralinguistic phenomena like proxemic relations.

Fusion of video and linguistic information is necessary to deal with intentional or unintentional ambiguities present in the linguistic and/or in the visual component. 'Fusion' techniques have shown to be a viable method to improve the automatic understanding of an oral communication (Ferrin *et al.*, 2008) or visual information (Snidaro *et al.*, 2011), but these techniques must be developed within a methodological framework which considers as a primary issue the concept of 'context'.

Context has been investigated for a long time now in diverse domains like Linguistics, Epistemology, Cognitive Psychology, and Artificial Intelligence (AI) aiming at formally defining and representing context, and discovering and reasoning about context relevancy for particular applications. A recent survey covering different domains and focussing on context for Information Fusion can be found in Snidaro *et al.* (2015).

Information extraction from multimodal communication is a task further complicated by the fact that it is not always clear what is the real focus, and any communication stream (or mode) can play the role of context for the others. Information is generated by correct interpretation of data. But the huge amount of diverse communication, very often noisy and inaccurate (e.g. contained in UGC data streams, that is User-Generated Content coming from 'human sensors') has to be properly processed to give us the 'comprehension' (*comprehendere* in Latin means 'to grasp' and 'to retain together') of a situation and its possible interpretations.

1.2 Metaphor, pervasively present in human communication

The conceptual pathway starting from the analysis of the intercontextual relation between communication modes, to meaning 'discovery' issues through 'construction' mechanisms of conceptual structures and causal chains, has made emerge some methodological arguments about meaning analysis of multimedia streams representing real-world situations. Attention has been drawn on a well-known phenomenon, which has been first studied by Aristotle in the *Poetics* (Aristotle, 1995), namely metaphor. Up until most recently, metaphor has been primarily studied by philosophers, rhetoricians, literary critics, psychologists, and linguists, such as Hume, Locke, Vico, Herder, Cassirer, Buhler, I. A. Richards, Whorf, Goodman, Max Black, and beside them thousands of people did work on metaphor over the past 2000 years. Today, an increasing number of cognitive scientists engage in the research on metaphor. The reason is that metaphor, together with its influence in the creation of our social, cultural, and psychological reality, plays a very important role in human thought, understanding, and reasoning.

Within science, traditionally viewed as a neutral and objective activity, metaphor tends to be regarded as at best irrelevant and at worst detrimental. Hobbes, for example, famously argued in his *Leviathan* (Hobbes, 1651) that 'metaphors and senseless and ambiguous words, are like *ignes fatui*; and reasoning upon them, is wandering amongst innumerable absurdities' ('Of reason and science', ch. 5) and should therefore be excluded from reasoning, demonstration, and 'all rigorous search of truth' ('Of the virtues commonly called intellectual; and their contrary defects', ch. 8). This view has progressively shifted by the recognition that the use of metaphor is both pervasive and essential in science as in everyday life, as thoroughly analyzed in Semino (2008). Many scientists have themselves written on the use of metaphor in their specific disciplines.

According to a common, traditional view of science and technology, scientists engage in the direct observation of phenomena, and use language to report their findings neutrally and objectively. This view has been shown to be inadequate in many studies which have pointed out that the phenomena studied by scientists are often not directly accessible, and can be observed in more than one way (e.g. can only be observed via instruments whose output requires interpretation), and scientific 'knowledge' or 'facts' are constructed, for example, via social (and linguistic) processes involving negotiation, argumentation, and persuasion, as well as reasoning with visual metaphors (McAllister, 2013).

These considerations show how metaphors can also be found in technical reports as well as in any other message on any medium vehicle.

Being prevalent in human linguistic discourse, metaphor is 'an increasingly looming obstacle for Engineering AI, as attempts are made to bring better automated human-language processing into commercial products, to develop ever more advanced computer interfaces and virtual reality systems, to develop automated understanding and production of emotional expression given that this is often conveyed explicitly or implicitly by metaphor' (Barnden, 2008); this is the main reason why metaphor has been investigated also in the field of AI, and draws our attention.

Some systems have been developed which show emphasis on uncertainty and gradedness in metaphorical reasoning and focus on mechanisms for exploiting context. Among the others, the *MIDAS* system (*Metaphor Interpretation, Denotation, and Acquisition System*) produced by Martin (1990, 2000); the *meta5* system, implemented by Fass (1997), also measuring the degree of disanalogy between source and target structures, and using this measure in rating the aptness of the metaphor; Barnden's *ATT-Meta* (Barnden & Lee, 1999) recently included in the *Gen-Meta* project (Gargett & Barnden, 2015), a system whose purpose is generating metaphors exploiting AI reasoning and corpus-based modelling; Narayanan's (1999) metaphor understanding system that has mostly been applied to interpreting metaphorical statements about economic policy; Veale's *Sapper system* (Veale & Keane, 1997; Veale, 1998) an hybrid symbolic/connectionist model for finding structural analogies and the more recent service-oriented architecture for metaphor processing (Veale, 2014); and Hobbs' *TACITUS* system (Hobbs *et al.*, 1993) investigating how knowledge is used in the interpretation of discourse in developing procedures using a large knowledge base of commonsense and domain knowledge.

All of the cited systems perform textual/linguistic metaphor interpretation and, as far as we know, not much effort has been devoted to trying to deal with pictorial or multimodal metaphor.

As stated earlier, both generating and understanding information intertwined in complex and often ambiguous texts[1] requires a creative form of reasoning that we identify with abduction.

1.3 Finding explanations and integrating information: abduction and Information Fusion

In recent decades, abduction has gained attention in several fields as philosophy of science, logic and AI. It has been widely recognized that the abductive style of reasoning is in fact inherent in many different tasks, from problem-solving to knowledge acquisition. Much research has been devoted to elaborating a better understanding of the theoretical grounds underlying abductive reasoning (Hanson, 1958; Harman, 1965; Flach & Kakas, 2000; Aliseda, 2006), the purpose being the implementation of a vast scope of applications from language and multimedia interpretation, through design problems, to planning and many others. Philosophers and logicians often disagree about the character of abduction, arguing from different premises, and the discussion is particularly extended also in contemporary works. The same is not happening in other domains where abduction plays a very important if not vital role, within which, usually, the only definition and structure of abduction is 'Inference to the Best Explanation' (IBE). For example, in Linguistics, in Mathematics, in AI, generally speaking, and not to mention specific tasks like metaphor-processing tools themselves.

Within the domain of AI, we will focus our interest on the field of research of Information Fusion (see Liggins *et al.*, 2008) because in its technological domain, which deals with the real world, the faced problematic issues imply processes which follow the sequences of scientific discovery and proof as well as our everyday life experience.

In the field of our ordinary perception, for example, abductive inference clearly shows to apply to fusion processes. Within the human system of visual perception, unconscious abduction processes perform the recognition of physical objects (3D in nature) based upon many images (2D projections) and the completion of objects partly concealed behind others. A fusion task can be also recognized in the inter-sensual correlation between visual perceptions and tactile perceptions: when the visual appearance of an object, for example, is not enough to build a satisfying knowledge, we will probably get next to the object and try to touch it to collect more information about it.

Aim of the present paper is providing a survey of concepts, bridging the domains of metaphor, abduction and data fusion, and structuring the problem. We argue that their integration can lead to designing some framework for automatic metaphor understanding of real-life multimodal, intrinsically ambiguous, communication. Our survey will show that metaphor resolution and creation process actually is a fusion process requiring several 'levels' and several tasks to be performed. At the semantic level, the tasks consist in structuring and mapping 'information' as well as abductively reasoning in a 'creative' way.

2 Metaphors

Many different approaches can be found in literature of metaphors, but the following characteristics are universally recognized:

1. Metaphor is a phenomenon universally present in communication domain, both linguistic and paralinguistic, in everyday common speech as well as in scientific and specialized discourse.
2. Metaphor is a phenomenon whose working principle presupposes the existence or the construction of conceptual structures or causal chains in an other but mechanistic way with wide uncertainty ranges.
3. Metaphor resolution (that is correct interpretation), as well as its ideation and construction, needs an inferential process of structure generation, identification, and mapping, very often a 'creative' one.

2.1 Definition

Looking for 'metaphor' in dictionaries or other resources it is possible to find many different definitions, possibly influenced by different theoretical backgrounds. A quite 'neutral' one to start from (the translation

[1] We use the word 'text' in the semiotic sense of a message which can exist in any medium.

of the original Greek word *metaphora* is ‘transfer’, from *metà*—‘over’ and *pherein*—‘to carry’) could be the following definition from Wordnet 3.0:

1. A figure of speech in which an expression is used to refer to something that it does not literally denote in order to suggest a similarity.

or the following, which can be found in the *Encyclopaedia Britannica*

2. Metaphor [is a] figure of speech that implies comparison between two unlike entities, as distinguished from simile, an explicit comparison signalled by the words ‘like’ or ‘as’.

For example, we would consider the word ‘lion’ to be a metaphor in the sentence ‘Spartacus was a lion in the fight’. There is a widely shared view—the most common conception of metaphor in the popular mind (sometimes the only view of metaphor), according to which it is said that the word ‘lion’ is used metaphorically in order to achieve some artistic and rhetorical effect, to communicate eloquently, to impress others with ‘beautiful’, aesthetically pleasing words, or to express some deep emotion.

To characterize this commonly accepted concept, some of its features can be pointed out: metaphor is just a linguistic phenomenon, used for rhetorical purposes, based on some similarity between the two entities. Metaphor is a conscious use of words mastered by talented people like poets and, as such, not necessary in everyday communication.

George Lakoff and Mark Johnson in their 1980 study *Metaphors We Live By* (Lakoff & Johnson, 1980) developed a novel view of metaphor, the Cognitive Linguistic, which coherently and systematically challenged all the cited aspects of the traditional theory. They basically claimed that metaphor is a property of concepts, and not of words, whose function is highlighting certain concepts. Ordinary people usually take advantage of metaphors in everyday language and not just for aesthetic purposes and is an inevitable process of human thought and reasoning.

2.2 Key points, terminology

A linguistic metaphor consists of two elements. The first element is the ‘thing’ (object or concept) we intend to talk about, but we represent it with signs pertaining to another ‘thing’. This second element is then used to convey concepts pertaining to the first element.

A group of metaphor scholars, from a variety of academic disciplines, has recently worked on and proposed an explicit method for identifying metaphorically used words in spoken and written language (Pragglejaz Group, 2007). The procedure aims to establish, for each lexical unit in a set of utterances (part of a discourse), whether its use in the particular context under investigation can be described as metaphorical and to provide a reliable research method for determining whether words in contexts convey metaphorical meaning. The general purpose being to provide a research tool relatively simple to use and flexible for adaptation by scholars interested in the metaphorical content of natural discourse.

The ‘metaphor identification procedure’, is as follows:

1. Read the entire text–discourse to establish a general understanding of the meaning.
2. Determine the lexical units in the text–discourse.
3.
    a. For each lexical unit in the text, establish its meaning in context, that is, how it applies to an entity, relation, or attribute in the situation evoked by the text (contextual meaning). Take into account what comes before and after the lexical unit.
    b. For each lexical unit, determine if it has a more basic contemporary meaning in other contexts than the one in the given context. For our purposes, basic meanings tend to be

        - more concrete (what they evoke is easier to imagine, see, hear, feel, smell, and taste);
        - related to bodily action;
        - more precise (as opposed to vague);
        - historically older.

Basic meanings are not necessarily the most frequent meanings of the lexical unit.

c. If the lexical unit has a more basic current-contemporary meaning in other contexts than the given context, decide whether the contextual meaning contrasts with the basic meaning but can be understood in comparison with it.

4. If yes, mark the lexical unit as metaphorical.

Throughout the studies in the field the terminology to denote the lexical units involved in a metaphorical substitution, has evolved. We will employ the two terms which are now commonly used: 'target' and 'source'.

Thus in the sentence

(3) Theories are buildings

'theories' is the target and 'buildings' the source.

Target and source are part of a complex network of related meanings conveyed by words. 'Theories' is part of a conceptual network including, for example, 'theorists', 'construct (argument)', 'foundation', 'support', etc. but it can be infinitely extended. The source, 'buildings', is part of a network of concepts as well. It includes 'construction', 'foundation', 'ground', 'demolition', 'framework', 'architecture', etc.

Not only denotations (that is the strict and literal meanings, as found in a dictionary) are part of the networks but also words' connotations (the suggestive emotional contents associated with the word and additional to its explicit literal meaning) and also pragmatic considerations. Such networks, therefore, connect concepts, attitudes, potential actions, and so forth, and are usually referred to as the 'target domain' and the 'source domain'.

In a metaphor, at least one feature linked to the source (belonging to the source domain) is mapped onto the target exploiting some resemblance between the target and the source. In THEORIES ARE BUILDINGS, for example, a minimal resemblance between the two domains is that in both someone is working on connecting 'parts' which 'support' one another.

A structural relationship between elements found in the source domain and the corresponding elements in the target domain must exist to enable the mapping from source to target, but also the context in which a metaphor appears must provide the interpreter with the details about which features are to be mapped (Figure 2).

Usually only a few features or characteristics of a source are mappable, and which features are to be mapped depends on many circumstances, namely on context.

Sometimes a metaphor suggests the mapping of a single feature. When Spartacus is called a lion, it is very likely that the only feature intended for the interpreter to map from the source domain LION to the target domain SPARTACUS is 'fierceness' and not having mane and a tail (Forceville, 2005).

Metaphors can also provide views of a given target domain creating links with an unexpected source domain, or by mapping unusual features from a familiar source domain to the target. This happens when the chosen

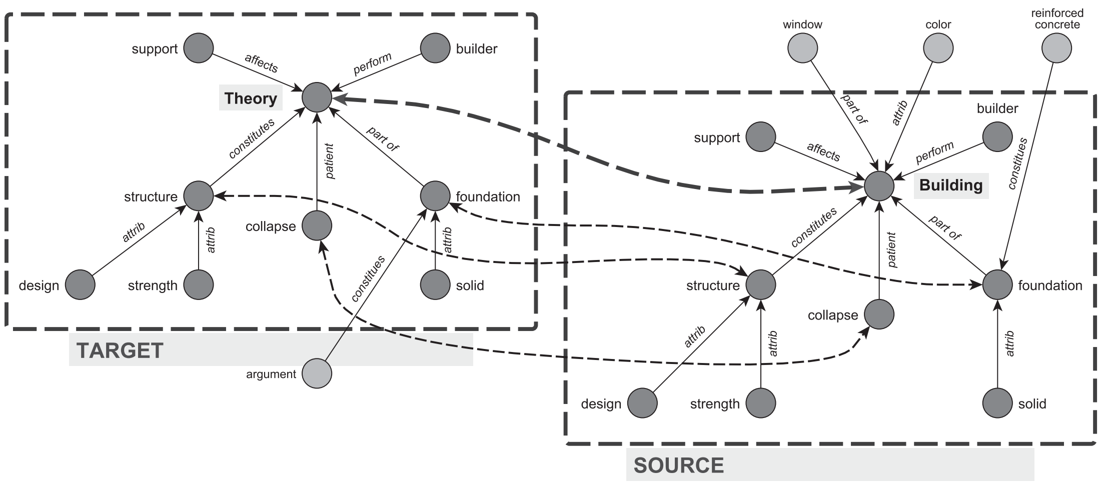

Figure 2 Structure mapping in metaphor

source domain marks a similar latent structure in the target domain. This is what very interestingly happens in the use of metaphor in scientific theory-making and in the teaching or explication of same theories.

Since the 1980s, conceptual 'revolution' of Lakoff and Johnson, some different theories have been carried out. The currently active ideas in the field of metaphor are covered by, among the others, the Discourse Dynamics Framework (Cameron & Deignan, 2006), the Context-Limited Simulation (Barsalou, 1999), the Lexical Concepts and Cognitive Models Theory (Evans, 2009), and the Conceptual Metaphor Theory (see Kövecses, 2002).

2.3 Multimodal metaphors

As we pointed out, most of the concepts and object representations that are exchanged through new media is multimedia in nature, including images, video, audio, text data, and hierarchically structured documents. Metaphors consist of the same two cited elements in whatever medium.

Let us consider the concept NECKTIE form a semantic-pragmatic point of view as we did in Ferrin *et al.* (2011). A NECKTIE is a piece of clothing associated, in our culture, with some dress code, it has 'fashion' connotations. It has a form, dimension, and texture; it is worn on the body. It can be an element in a ceremonial outfit. In principle, all the connotations pertaining to an uttered word could be evoked, probably not in the same way for any person, and the same we expect to happen looking at a pictorial representation of the same object the word denoted. However, a picture can only depict some *specific* object, for example, the necktie, and that specific necktie, even if it is decontextualized, does not only share connotations with other decontextualized neckties but also has connotations differing from them.

The following example, whose structure we owe to (Forceville, 2005), will clarify the concept. Look at Figure 3(A (a)), it can be considered without doubt a prototypical necktie, which, despite the absence of any context, conveys some connotations of elegance, of visual pleasure, etc. Compare this to Figure 3(A (b)), which too looks like a rather prototypical necktie. This is not a photograph but a drawn picture of a necktie and the connotations proposed for the necktie in Figure 3(A (a)) do not adhere to the necktie in Figure 3(A (b)), a probable connotation here is arguably that it is an item in a list or a knot schematics.

In the next picture, Figure 3(A (c)), the represented necktie is similar to that in Figure 3(A (a)) in being once more a photographic representation, but different in the sense that it is a different kind of necktie. Being a football team tie seems to be connected to some uniform, thus evoking, together with 'elegance', the 'belonging to a team', the loyalty owed to it, the pride of wearing it, etc.

While the connotation of 'belonging', and somehow also of 'pride', seems to be conveyed by the necktie in Figure 3(A (d)), too, the connotation 'elegance' may be here reduced because we recognize that this is a geek necktie. So among the potential associations activated here are 'nerd', 'geek', etc.

In Figure 3(B) is another example: a series of representations of DOG COLLARS. While dog collars prototypically connote dogs, individual dog collars may elicit specific associations. The dog collar in

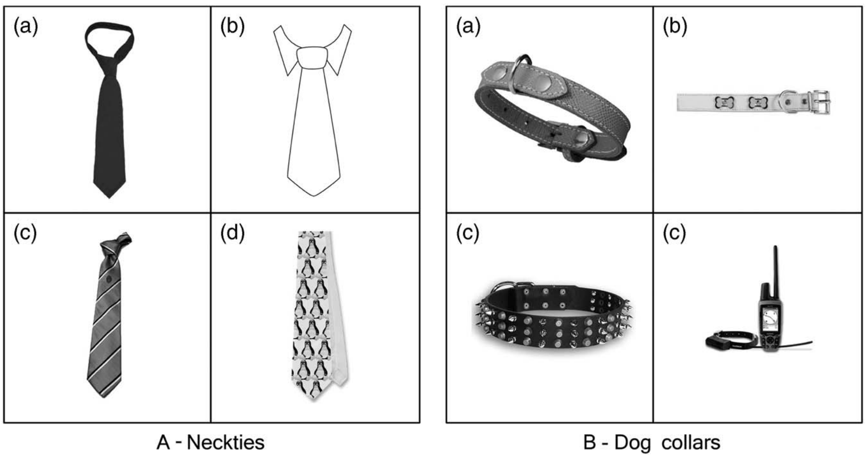

Figure 3 Representations denoting the same objects, Neckties on the left and Dog collars on the right, conveying different connotations

Figure 3(B (a)) looks very ordinary (it is worth noting that 'ordinariness' is in itself a potential connotation). Figure 3(B (b)) shows a dog collar that is less straight, it is coloured in pink and it looks like to be worn by a little 'pet' doggie.

The predominant characteristics of the dog collar in Figure 3(B (c)) are the spikes and pyramids and connote 'ferocity', most certainly the 'pet' dimension here is almost non-existent. The example in Figure 3(B (d)), finally, is definitely different from the others. It can be defined (and actually is defined so) a dog collar just because it is worn by dogs, but the tool functionality and purpose here changes from 'physical' to 'virtual' control.

With the images we have, we could start building metaphors, even if no context is available. The similarity between neckties and dog collars that serves as basis for a metaphor—the 'metaphorical tension' between target and source can arise only thanks to some degree of similarity between the two—can be recognized in the form of both neckties and dog collars.

Let us say we put the dog collar in Figure 3(B (c)) in the 'target' slot, and the necktie in Figure 3(A (a)) in the 'source' slot, so as to construe the metaphor DOG COLLAR in Figure 3(B (c)) IS NECKTIE in Figure 3(A (a)). The fluffiness of dog collar in Figure 3(B (b)) being a salient feature, the metaphor could be interpreted to mean something like 'the dog collar in Figure 3(B (c)) is part of a dress code as necktie in Figure 3(A (a)), a particular dress code indeed. Of course the metaphor we have just tried is a strained example. In 'real life', decontextualized examples like the one discussed above are very seldom met.

There is no 'natural' way in which features from both domains are to be matched. Which features can be matched will depend on the context in which the metaphor occurs and different combinations of features will lead to different 'emergent properties'. In Figure 4, for instance, it can be recognized, through the gesture context, the metaphor THIS NECKTIE IS A DOG COLLAR.

Research devoted to developing a model for the analysis of pictorial metaphor, which can be largely found in static advertisements, has been performed by Forceville (1994, 1996, 2000, 2005) leading to the recognition of different metaphor typologies or prototypes, on the basis of how target and source are represented, namely hybrid metaphor, contextual metaphor, pictorial simile, and integrated metaphor.

Anyway, only the awareness of incongruity and similarity of the phenomena represented in a picture guides the user to the recognition of elements to be identified as the target and source of a metaphor. As in pictures there is no equivalent to the verbal or textual 'is' or 'is like', it may be an issue whether a picture actually features a metaphor. Sometimes a metaphor identification is self-evident; sometimes only some viewers will 'see' a metaphor. For this reason, as Forceville (2005) argues 'it makes sense to say that a picture strongly or weakly invites a viewer to construe a metaphor, rather than to say that a picture contains a metaphor'. 'Strongly' or 'weakly' depending on intentions of the producer to force the audience, more or less, to perceive it.

We have so far separately considered verbal metaphors and pictorial ones. In these metaphor examples the targets and the sources pertain to the same communication mode, that is, natural language and images. However, any sign-bearing connotations upon which it is possible to build a metaphorical target or source,

Figure 4 Necktie—dog collar—a rather 'easy' example of NECKTIE IS A DOG COLLAR

Figure 5 Text and visual—verbo-pictorial metaphor

can belong to any 'mode', for example, music, sounds, or other stimuli. Modes appear to be linked to senses, as signs are perceivable items; however, a correlation between modes and the five senses is not proper as, for example, written language and pictures are both seen, and either spoken language, or music, or sounds are all heard. In metaphor analysis, therefore, the following modes are usually distinguished: language, both spoken and written, images, non-verbal sounds, music, smell, taste, and touch.

It frequently happens that a metaphorical term (no matter if target or source) is simultaneously conveyed by more than one mode. A multimodal metaphor, just like any other metaphor, can only function if the source domain (whatever it is) is perceived and recognized, and elicits one or more connotations which can be matched to elements in the target domain, through feature and structure mapping.

The research on multimodal metaphor is mainly focussed on the mass-media type, linguistic and pictorial, therefore 2D. Smell, taste, and touch are usually not dealt with in literature, though their use in metaphors show absolutely no issues.

A metaphor is to be considered multimodal (see Forceville, 1996), when its target and source are rendered 'exclusively or predominantly in two different modes' as happens in the combination of an image with a text in Figure 5, an example of one of the most frequent types of multimodal metaphors in everyday life.

The great amount of such diverse produced and available data, noisy, inaccurate, and ambiguous, is gathered from senses and sensors and must be processed to give us the 'comprehension' of a situation and its possible interpretations.

Studies in cognitive sciences show how achieving significant degree of success in 'comprehension' needs discerning the underlying regularities in the world, despite sparsity and noise (Tenenbaum *et al.*, 2006), also exploiting relational structures and direct similarity of objects like in the multi-constraint theory exposed in Holyoak and Thagard (1995). According to these cognitive theories, the best the human mind can do in inferring from available data is to make the 'best possible guess' guided by prior probabilities about which world structures are most likely. Prior knowledge provides some 'library' of structures that are available to fit any data set. As Hume stated 'All kinds of reasoning consist in nothing but a comparison, and a discovery of those relations, either constant or inconstant, which two or more objects bear to each other' (Hume, 2007).

Bayesian approach seems to best model human reasoning over structures, relations, and links, and it is possible to provide a detailed computational account of how a number of basic structural forms can be inferred from various types of data (feature sets, similarity matrices, relations) out of different areas of interest, covering higher-level problems like inferring causal structures, learning about hidden properties or objects, and interpreting the meaning of words.

This is the core problem at the very first stage of any processing and reasoning system: choosing 'significant' data to start the interpretation process. In logic-based approaches, talking about 'significant' data (or features) implies already having an 'ontological' (in philosophical sense) model according to which data can be judged as such, that is having some already structured *a priori* assumption over the situation which is the object of a semantic analysis and reasoning. The feedback consists in a hard revision of ontological structure of the model. A probabilistic approach, on the contrary, can succeed in identifying a structural model using just loose belief constraints. It does not need a pre-existing ontological model. The feedback consists in a recursive revision of our believes.

In any of the two cases, a backward reasoning task is needed, that is an abductive inference.

3 Abduction

We have seen as one of the key roles in metaphor recognizing is played by a mapping procedure between conceptual structures. Both structures and mappings are not stable, depending on context-dependent connotations, therefore, cannot be made available in ontological or lexical frameworks.

Mapping discovery is performed through an abductive inferential process. In this section, abductive inference will be surveyed in the domain of epistemology, mainly following the analysis worked out in Schurz (2008), and will show how some principles which can be found within the epistemological debate can be borrowed to be directly applied to a fusion domain, with the purpose of transferring them to metaphor analysis (Figure 6).

Charles S. Peirce, the 18th-century American pragmatist philosopher, in his writings, talks several times about abduction and gives to the reader a number of slightly different definitions, for a complete review see Niiniluoto (1999).

Studying Aristotle's doctrine of induction, which is considered the inference of the major premise of a syllogism, Peirce observed that there is 'a large class of reasonings' that are neither deductive nor inductive: for example, reasoning *a posteriori* to infer a cause from its effect. Peirce called *Hypothesis* this kind of reasoning, which can be represented as the inference of the minor premise of a syllogism. In an 1878 article, 'Deduction, induction, and hypothesis', Peirce (1931: CP 2.623) reports the two ways of permuting the arguments of a deductive inference.

Hypothesis is the inference of the cause from the rule and result:

(4) All the beans from this bag are white.
These beans are white.
∴ These beans are from this bag.

A typical example of *Hypothesis* in the sense of (4) has thus the following logical form:

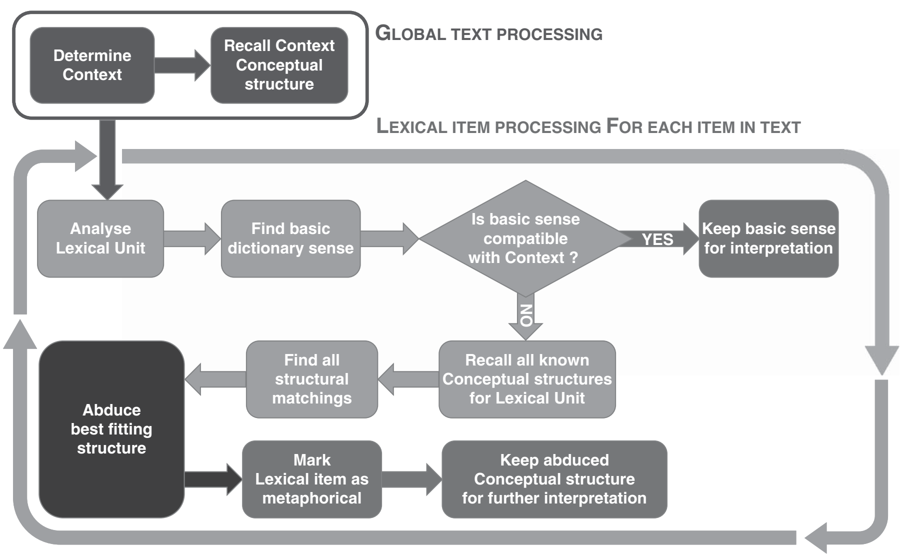

Figure 6 A simplified schema for metaphor recognition through abduction

(5) Given the law $\forall x(Fx \rightarrow Gx)$,
from Ga infer Fa.

Soon it became clear to Peirce that *Hypothesis* is an inference to an *explanation*:

> We find that light gives certain peculiar fringes. Required an explanation of the fact. We reflect that ether waves would give the same fringes. We have therefore only to suppose that light is ether waves and the marvel is explained' (Peirce, 1982–: 1:267).

and in 1866, the philosopher says that *Hypothesis* is the inversion of the explaining syllogism which has the typical structure of a deduction (Peirce, 1982–: 1:428, 425, 440, 452).

Deductions are certain: given true premises, the conclusion must be true, but non-ampliative. It is commonly recognized the tautological nature of deductions. On the other hand, inductions and abductions are ampliative, their conclusions enrich knowledge but in uncertain way: they may well be false even if the truth of the premises can be taken for granted, so further testing is necessary.

Many epistemologists have been using the term 'induction' to denote all kinds of non-deductive (therefore ampliative) inferences since the 19th century when John Stuart Mill despite 'several methods for discovery and demonstration of causal relationship' had been recognized.

Nowadays, philosophers and logicians (Aliseda, 2006; Schurz, 2008) recognize that induction and abduction are distinct inferences not reducible to each other. Both of them have the purpose of extending knowledge beyond empiric data but with rather different purposes. The goal of inductions consists in trying to foresee the future events. On the other hand, the goal of abductions consists in trying to attribute an unobserved (or unobservable) cause or explanation to some observed events.

The main reason according to which abductions cannot be reduced to inductions consists in that the latter do not introduce new concepts or conceptual models in the inferential process; their 'ampliative power' is confined to merely transferring concepts to newly observed instances. Some forms of abductions, on the other hand, do introduce new concepts or models. Following Magnani (2001), they are called *creative*, as opposed to *selective* abductions whose goal is to select the best possible explanation out of the available ones.

Besides a 'Peircean' use of the term *abduction*, basically seen as 'the third' mode of inference and further studied by Hanson (1958), who investigated the logic of discovery in the 1950s, another way of using it considers it a synonym for the IBE model. Such a model, which has a great diffusion in the technological domains, from multimedia automated analysis (Möller & Neumann, 2008) to the fusion domain itself (Bharathan & Josephson, 2006; Josephson, 2008) is based on the studies of Harman's (1965), who formulated the idea in the 1960s.

Most of the authors working on explanation agree that nobody can know all possible explanations for a given phenomenon, therefore IBE needs to be modified: what can be really reached is an inference to the best *available* explanation, a so-called IBAE. Some philosophers have advised, however, that the best available explanation within an inference could happen to be not rationally acceptable. If we come across some new and not well-understood phenomenon, then the best explanation we can think of is usually a mere conjecture.

In Gabbay and Woods (2005) can be found an insightful analysis of abduction from a logical point of view, and, as declared by authors a very prudent one indeed. Among the other key concepts exposed in the book, very important are some definitions covering the concepts of 'ignorance problem' (IP) which somehow triggers and drives an abduction process, 'ignorance preservation' ('whereas deduction is truth-preserving and induction is probability-enhancing, abduction is ignorance-preserving') and 'ignorance mitigation' ('although a solution to an abduction problem preserves the ignorance that gave rise to it, it may also contribute to the solution of the originating problem by identifying candidates for the status of new knowledge').

A cognitive agent faces an *IP* whenever she has a cognitive target T that cannot be accomplished exploiting his current knowledge base K. At this point she is usually given two possible options: (1) acknowledging that T is not attainable from K, constituting the pair $\{K, T\}$ an *insolubilium* or

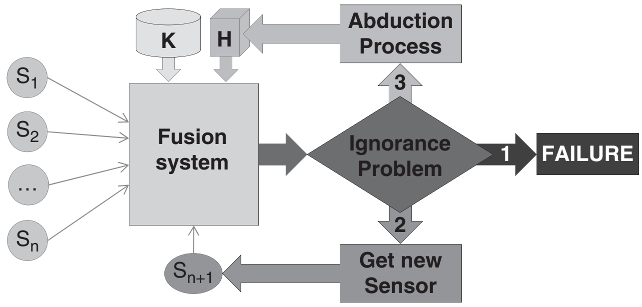

Figure 7 Facing an ignorance problem—possible options

(2) acquiring new information with the purpose of enabling *T* to be attained (through new and different sensors extending his *K*).

Actually, the agent has a third option (3) which represents, according to Gabbay and Woods, 'the founding datum of abduction'. The agent finds an hypothesis *H* which, if he knew it, together with *K* would solve his *IP*; and from that fact he conjectures *H*. In the fusion domain, as it is shown in Figure 7, the hypothesis *H* generated by an abductive inference can be considered as an input to the fusion process, just like any other sensor. Option (3) essentially incorporates the element of conjecture. But the critical point is that this does not solve the original problem. The agent's problem is that *T* is attainable only on the basis of what he senses or knows (*K*) or can manage to know. After turning to option (3), his situation is that *T* cannot be attained either way. Selecting an *H* such that the truth of *K* revised by *H* (*K* (*H*) for now on) would hit *T*, then conjecturing *H* does not extend *K*; *K* (*H*) is not a knowledge base for the agent and *X*'s *IP* is not logically solved.

This highlights a second unavoidable element of conjecture embedded in option (3). *K* (*H*) does not directly hit *T*, but it does it *presumptively*. Accordingly, option (3) does not offer the agent a solution of his IP, but an attainment of a minor target. It proposes a conjectural variant of a target which admits of only *epistemic* attainment: a target that provides *presumptive* attainment. A thorough investigation about ignorance-preserving and 'mitigating' character of abduction can be also found in Magnani (2009, 2013), and a new model contending a knowledge-enhancing role for abduction is proposed in Magnani (2015).

All considered authors agree upon the fact that all inferences have, to a different degree, a *justificational* (or 'inferential') and a *strategical* (or 'discovery') function (Schurz, 2008). The justificational function of an inference consists in the justification of the conclusion, under condition of the justification of the premises. The strategical function consists in finding a most promising conjecture (conclusion) which stimulates further empirical tests, or new questions (Hintikka, 1998).

In any deduction the justificational function is maximal, because the inference structure guarantees the truth of the conclusion, provided that the premises are true. Deductions may have also non-trivial strategical functions, as from the same premises can be derived a number of different conclusions. In inductions, instead, search strategy does not play a predominant role, because, given the premises, the possible conclusions are only drawn through a process of generalization over the collected instances. The primary function of inductions is therefore justificational, but, by its nature, the justificational value of an induction lacks the character of certainty.

The strategical function appears to be central in abductions. In abduction problems, cognitive agents often deal with a huge number of possible explanations (conjectures or conclusions). Abductions essentially consist in actual search strategies which drive the agent in choosing an explanation to further investigate or which suggest the agent a route, short and promising enough (but without any guarantee of success), among the exponential complexity of possible explanations. However, in abductions the justificational function is secondary. Peirce (1931: CP 5.171) himself mentioned that hypotheses abductively inferred are merely possible, unlike the probable ones inferred by induction. An abductive hypothesis may gain probability only after being further confirmed.

A good search strategy leading an agent to an optimal conjecture must give results in a finite and possibly acceptable time. In this regard, the IBE or IBAE procedures fail. IBE or IBAE can tell the agent which is the best available conjecture only after evaluating them all, IBE follows just the justificational function but it rules out the strategical one which is essential in the nature of abductions. The rule of IB(A)E is therefore quite uninformative from the epistemic point of view and definitely not creative. Peirce himself remarked that phenomena under scientific investigation can be explained by an extremely high number of hypotheses, nevertheless, researchers usually find the correct explanations without evaluating them all, but after only a rather small number of *guesses* (Peirce, 1931: CP 6.5000). On the contrary, Peirce did not write a line about any abductive rules for conjecturing new theories; he explained the ability of human minds just talking about abductive *instincts*. The central issue seems to be the existence of any 'logic' underlying explanation processes. According to Schurz (2008), patterns such rules exist and every kind of abduction pattern, classified and labelled by the author, constitutes such a rule, each pattern fitting a particular kind of conjectural situation.

Analogical abduction is defined as a second-order pattern of abduction. In this kind the *explanandum* consists of one or more empirical phenomena or laws. The conclusion of this abductive inference is a new property or kind concept dependent on a new, or partly new, theoretical law. An extrapolation, an analogy, or a simple unification can drive such an abduction depending on the grade of novelty of the concept.

In an analogical abduction a partially new concept is inferred together with new laws which relate the same concept with other already known concepts, in order to explain the law-like phenomenon under investigation. The concept novelty is only partial because of its analogy to familiar concepts, and its discovery is just due to the analogy. Analogical abduction is *driven* by analogy. Like in the following scientific example:

(6) *Background knowledge*: laws of propagation and reflection of water waves.
Phenomenon to be explained: propagation and reflection of sound.
Abductive conjecture: sound consists of atmospheric waves in analogy to water waves.

According to Schurz (2008), an analogical abduction involves a central process of *conceptual abstraction* based on a *mapping* between *isomorphic* (or homomorphic) structures. The hypothesis abduced from the analogy present in Example 8 not only consists in the combined concept of 'sound *is a* wave', but at the same time the abstract (or *theoretical*) concept of *wave* itself.

Gentner (1983) gave a lucid analysis of analogy based on mapping and conceptual abstraction. According to Gentner's analysis, an analogy is a *partial isomorphic* mapping between two relational structures, the *source structure* and the *target structure*. Both structures are constituted by monadic predicates, that is properties, and relations connecting them. Gentner argues that in an analogical mapping only the relations of the two structures are preserved, while monadic properties are not. This variability of properties characterize the distinction between an *analogy* and a *literal similarity*.

It is worth noting that the use of 'metaphor'-related terminology (*source structure* and *target structure*) which preludes the following Gentner research on metaphor in Gentner (1988).

The problem of discerning the underlying regularities in the world, despite sparsity and noise in data and information, has been addressed to by recent studies in cognitive sciences and some good results have been achieved by taking into consideration some 'library' of structures (see Kemp & Tenenbaum, 2008), to fit most of the data set under investigation. Such ideas have been proposed to the high-level fusion community too (Ferrin *et al.*, 2009).

As presented in this section, abductive inference, whose goal is inferring the unobserved (or unobservable) causes or explanatory reasons of observed events, shows to be the unique possible form of 'creative' inference and can be considered from different points of view making emerge many interesting key concepts. Some of them are stuck in the subtleties of philosophy, but most of them are immediately disposable for a necessary theoretical discussion which preludes the development of tools to be used in a fusion process. IBE, for example, should not be considered as 'the' (almost unique) kind of available concept of abduction to deal with. 'Ignorance'-related definitions can have great influence on knowledge base update issues: an abduced conjecture H cannot be considered as a candidate to be introduced in the knowledge base itself in its 'floating' status waiting to possibly gain existence rights through a follow-up procedure.

4 Metaphor resolution as a fusion problem

Data fusion is a framework that enables the combination of data and information gathered from diverse sources with the purpose of creating a coherent and structured picture, thus facilitating further analysis. The framework and the tools were conceived primarily for military applications where there usually is an urgent need to assess complex situations making use of a large number of observations and to determine the optimal course of action. Many commercial applications nowadays exploit these technologies which deal with complex data and generate accessible information to decision makers. The application fields include, for example, situation assessment in the domain of security, robotics in the domain of manufacturing, and medical diagnosis in health care domain. A thorough review of well-established mathematical techniques used in a fusion process which involves interrelating and consolidating information at different levels of detail, can be found in Liggins *et al.* (2008). The US Department of Defense proposed a standard model for data fusion to facilitate discussion, component reuse, and system integration: the Joint Directors of Laboratories (JDL) data fusion model. The JDL constitutes a multi-level functional model that details the organization of the process in a military data fusion system. Generalizations of the model have been widely adopted both in academic research and in commercial applications beyond military purposes. Even though many authors have proposed expansions and more detailed specifications of the levels, generally to accommodate the requirements of specific applications or domains, the simplicity and generality of the initial model still stand, making it a *de facto* standard in the data fusion community.

A typical data fusion problem can be defined, according to the JDL fusion model proposed in Llinas *et al.* (2004) and its revisions, as a complex integration of processes working on defined functional levels (Figure 8):

- Level 0 (*signal assessment*): extracting features from raw data;
- Level 1 (*object assessment*): gathering information about individual entities;
- Level 2 (*situation assessment*): focussing on the relationships between entities and contextual implications;
- Level 3 (*impact assessment*): assessing consequences of applying known plans on the current situation;
- Level 4 (*performance assessment*): measuring the performance and effectiveness of the system to facilitate refinement.

If we draw our attention to some sentences like the following, easy to find in everyday news, examples of the metaphor ECONOMY IS A BUILDING:

(7) A stock market *collapse* would inflict *little damage* on the real economy.
(8) When asset prices increase to unsustainable levels, it will threaten to *destabilize* the economy.
(9) The X government has to start to *rebuild* the economy.

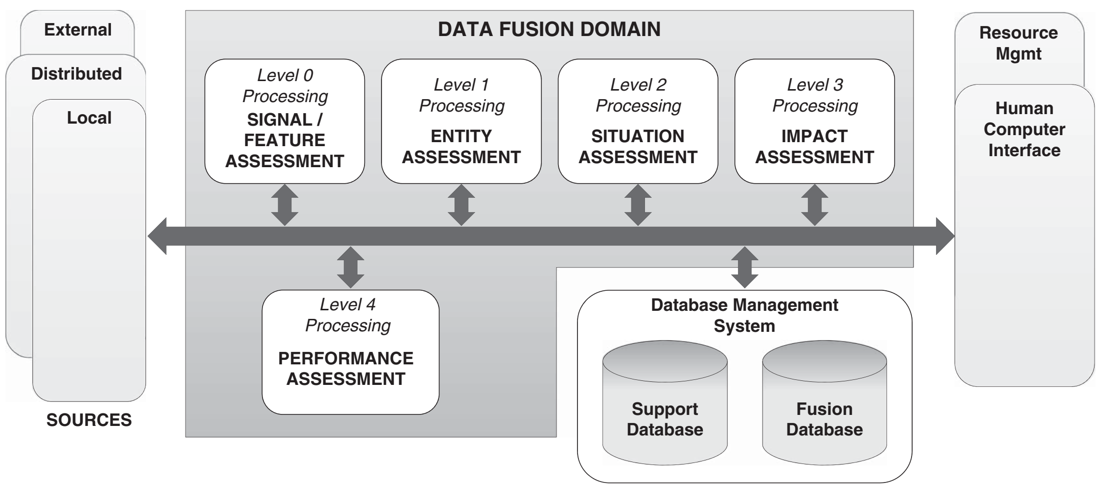

Figure 8 Joint Directors of Laboratories data fusion model—a model schematics from Steinberg *et al.* (2004)

(10) The *best-designed* market will deliver a better deal than a *poorly designed* one.
(11) The firms also *restructured* more vigorously.

and analyze the necessary processes to understand them, we can recognize the following steps:

- Level 0: text analysis, image processing, feature extraction, etc. Data are pre-processed and enhanced (e.g. noise reduction, filtering, etc.) in order to facilitate feature and descriptor extraction. In the case of text analysis, this would account to sentence splitting, part-of-speech tagging, dependency parsing.
- Level 1: referent recognition performed on lexical entry. In a metaphor, the problem consists in exploiting referents non-coherent with context. The system should keep any interpretation.
- Level 2: relations between detected entities in the text in the current context are established. These steps allow the understanding of the described situation and the resolution of metaphors within the current context.
- Level 3: the evolution of the current situation and/or context is formulated in order to improve the estimation process at the next step.
- Level 4: algorithms and parameters can be adjusted depending on the current situation/context to improve performances at all levels. In particular, a change of context could require different metaphor resolution algorithms.

In multimodal metaphors, for instance, the simple pictorial/textual metaphor in Figure 5, Level 2 mechanisms must process different mode of communications integrating the conceptual domain structures with all possible elements.

From the previous sections it is clear that the understanding of a metaphorical construction, whether it be only verbal or multimodal, requires an abductive form of pattern inference required to work out structure-mappings that could drive the formation of plausible hypotheses. Constructing plausible hypotheses both requires the fusion of available relevant information and the support for uncertain information management as hypothesized causes need to be properly handled to maintain the consistency of the knowledge base. It should be noted how in the 2004 revision (Llinas *et al.*, 2004), the abductive type of reasoning is mentioned as a key mechanism for the discovery of causes of observed patterns and should be taken into account in fusion processes paired with ‘knowledge discovery’ functions.

A possible mechanism for implementing the abductive process of hypothesis generation for fusion is reported in Burch (2000) describing the work of Professor Victor Konstantinovich Finn. In particular, the ‘JSM Method of Automatic Generation of Hypotheses’ (in honour of John Stuart Mill, the 19th century English philosopher) is a pattern-matching procedure with two stages. In the first one hypotheses are generated while in second hypotheses are either accepted or discarded. In Finn’s theory, all the entities, their properties, and relations are modelled through an ontology expressible in terms of semi-lattice. The JSM method describes several casual reasoning mechanisms, where the direct method of agreement is shown as the most straightforwardly implementable. This agreement method searches for the cause that has the maximal set of features that are common to all the cases in which the observed effect is present (Burch, 2000). The procedure takes then the form of a pattern-matching algorithm for (sub)structure mapping.

Due to space limitations, we will just mention two other important points for metaphor resolution: the proper grasping of the relevant contextual information necessary to its understanding and the concept of ‘artefact metaphor’.

The first is a relatively recent and relevant research topic in the fusion community (Snidaro *et al.*, 2015) but it is rapidly gaining popularity as the principled exploitation of the relations between the problem space variables and the context variables is likely to help explaining/constraining the problem variables of interest. This is particularly true in the case of metaphors where the context in which they are used is often the only key to decode their meaning via a ‘locally abductive’ process (Llinas *et al.*, 2004).

The second topic, which we are currently working on (Ferrin *et al.*, 2015), deals with the exploitation of ‘metaphorical’ artefacts (i.e. tools whose characteristics and features simply fit for the purpose), instead of the designed ones to achieve an intended goal. It is easily shown that the process of research and retrieval of artefacts which are possible substitutes for proper tools, involves a mechanism of similarity mapping between structures of features and functions, the same mechanism driving metaphor creation and resolution.

5 Conclusion

As Forceville (2002) argues in the core problems in general metaphor understanding can be set with three questions, namely:

1. Which are the two terms of metaphor, and how do we know?
2. Which is the metaphor's target domain and which is the metaphor's source domain, and how do we know?
3. Which features can/should be mapped from the source domain to the target domain, and how is their selection decided upon?

These questions are still far to be answered. Nevertheless, being metaphor pervasively present in every sort of communication, metaphor interpretation frameworks should be seriously taken under consideration.

Contemporary multidisciplinary research shows anyway that metaphor understanding is matter of 'recognition', 'knowledge' generation and mining, 'inferencing' and reasoning, and when holding an interpretation, of making use of gathered new 'hypotheses' to act and behave accordingly in further communication acts. These key points can only be addressed using abductive inferences setting the problem as a fusion process, so exploiting concepts and, above all, already robust algorithms and systems.

Aim of this paper is providing a survey of concepts, bridging the domains of metaphor, abduction and data fusion, and structuring the problem, in order to facilitate their integration.

References

Aliseda, A. 2006. *Abductive Reasoning. Logical Investigations into Discovery and Explanation (Synthese Library)*, Springer.

Aristotle, 1995. *Poetics*. William Heinemann Ltd.

Barnden, J. A. 2008. Metaphor and artificial intelligence: why they matter to each other. In *The Cambridge Handbook of Metaphor and Thought*, Gibbs, R. W., Jr. (ed.). Cambridge University Press, 311–338.

Barnden, J. A. & Lee, M. G. 1999. An implemented context system that combines belief reasoning, metaphor-based reasoning and uncertainty handling. In *Proceedings of the Second International and Interdisciplinary Conference on Modeling and Using Context*, CONTEXT '99, 28–41. Springer-Verlag.

Barsalou, L. W. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences* **22**(4), 577–660.

Bharathan, V. & Josephson, J. R. 2006. An abductive framework for level one information fusion. In *Proceedings of the 9th International Conference on Information Fusion*, 1–7.

Burch, R. W. 2000. Semeiotic data fusion. In *Proceedings of the Third International Conference on Information Fusion*, **2**, 67–73.

Cameron, L. & Deignan, A. 2006. The emergence of metaphor in discourse. *Applied Linguistics* **27**(4), 671–690.

Evans, V. 2009. *How Words Mean: Lexical Concepts, Cognitive Models and Meaning Construction*. Oxford University Press.

Fass, D. 1997. *Processing Metaphor and Metonymy*. Ablex.

Ferrin, G., Snidaro, L., Canazza, S. & Foresti, G. 2008. Soft data issues in fusion of video surveillance. In *Proceedings of the 11th International Conference on Information Fusion*, 1882–1889.

Ferrin, G., Snidaro, L. & Foresti, G. 2009. Structuring relations for fusion in intelligence. In *Proceedings of the 12th International Conference on Information Fusion*, 1621–1626.

Ferrin, G., Snidaro, L. & Foresti, G. L. 2011. Contexts, co-texts and situations in fusion domain. In *Proceedings of the 14th International Conference on Information Fusion*, 1–6. IEEE.

Ferrin, G., Snidaro, L. & Foresti, G. L. 2015. Artifact 'metaphors': gaining capability using 'wrong' tools. In *Proceedings of the 18th International Conference on Information Fusion*, 1176–1181.

Feyerabend, P. K. 1999. *Conquest of Abundance: A Tale of Abstraction Versus the Richness of Being*. University of Chicago Press.

Flach, P. A. & Kakas, A. C. (eds) 2000. *Abduction and Induction: Essays on their Relation and Integration*. Kluwer Academic Publishers.

Forceville, C. 1994. Pictorial metaphor in advertisements. *Metaphor and Symbolic Activity* **9**(1), 1–29.

Forceville, C. 1996. *Pictorial Metaphor in Advertising*. Routledge (Taylor and Francis).

Forceville, C. 2000. Compasses, beauty queens and other PCs: pictorial metaphors in computer advertisements. *Hermes, Journal of Linguistics* **24**, 31–55.

Forceville, C. 2002. The identification of target and source in pictorial metaphors. *Journal of Pragmatics* **34**(1), 1–14.

Forceville, C. 2005. *A course in pictorial and multimodal metaphor, Semiotics Institute*. `http://semioticon.com/sio/courses/pictorial-multimodal-metaphor`.

Gabbay, D. M. & Woods, J. 2005. *A Practical Logic of Cognitive Systems, Volume 2: The Reach of Abduction: Insight and Trial*. Elsevier Science Inc.

Gargett, A. & Barnden, J. 2015. Gen-Meta: generating metaphors by combining AI and corpus-based modeling. *Web Intelligence* **13**, 103–114.

Gentner, D. 1983. Structure-mapping: a theoretical framework for analogy. *Cognitive Science* **7**(2), 155–170.

Gentner, D. 1988. Metaphor as structure mapping: the relational shift. *Child Development* **59**, 47–59.

Hanson, N. R. 1958. *Patterns of Discovery*. Cambridge University Press.

Harman, G. H. 1965. The inference to the best explanation. *The Philosophical Review* **74**(1), 88–95.

Hintikka, J. 1998. What is abduction? The fundamental problem of contemporary epistemology. *Transactions of the Charles Sanders Peirce Society* **34**, 503–533.

Hobbes, T. 1651. *Leviathan, or, the Matter, Forme and Power of a Common Wealth Ecclesiasticall and Civil*. Andrew Crooke.

Hobbs, J. R., Stickel, M. & Martin, P. 1993. Interpretation as abduction. *Artificial Intelligence* **63**, 69–142.

Holyoak, K. J. & Thagard, P. 1995. *Mental Leaps: Analogy in Creative Thought*. MIT Press.

Hume, D. 2007. *A treatise of human nature*. Norton, D.F. & Norton, M.J. (eds). Clarendon Press.

Josephson, J. R. 2008. Abductive inferencing for integrating information from human and robotic sources. In *Proceedings of the 11th International Conference on Information Fusion*, 109–114.

Kadmon, N. 2001. *Formal Pragmatics: Semantics, Pragmatics, Presupposition, and Focus*. Blackwell Publishers.

Kemp, C. & Tenenbaum, J. B. B. 2008. The discovery of structural form. *Proceedings of the National Academy of Sciences of the United States of America* **105**(31), 10687–10692.

Kövecses, Z. 2002. *Metaphor: A Practical Introduction*. Oxford University Press.

Lakoff, G. & Johnson, M. 1980. *Metaphors We Live By*. University of Chicago Press.

Liggins, M. E., Hall, D. L. & Llinas, J. (eds) 2008. *Handbook of Multisensor data fusion: theory and practice, The electrical engineering & applied signal processing series, second edn*. CRC Press.

Llinas, J., Bowman, C., Rogova, G., Steinberg, A., Waltz, E. & White, F. 2004. Revisiting the JDL data fusion model II. In *Proceedings of the 7th International Conference on Information Fusion*, 1218–1230.

Magnani, L. 2001. *Abduction, Reason and Science. Processes of Discovery and Explanation*. Kluwer Academic Plenum Publishers.

Magnani, L. 2009. *Abductive Cognition—The Epistemological and Eco-Cognitive Dimensions of Hypothetical Reasoning*, Cognitive Systems Monographs **3**. Springer.

Magnani, L. 2013. Is abduction ignorance-preserving? Conventions, models and fictions in science. *Logic Journal of the IGPL* **21**(6), 882–914.

Magnani, L. 2015. The eco-cognitive model of abduction. *Journal of Applied Logic* **13**(3), 285–315.

Martin, J. H. 1990. *A Computational Model of Metaphor Interpretation*. Academic Press Professional, Inc.

Martin, J. H. 2000. Representing UNIX domain metaphors. *Artificial Intelligence Review* **14**, 377–401.

McAllister, J. W. 2013. Reasoning with visual metaphors. *Knowledge Engineering Review* **28**(3), 367–379.

Möller, R. & Neumann, B. 2008. Ontology-based reasoning techniques for multimedia interpretation and retrieval. In *Semantic Multimedia and Ontologies: Theory and Applications*, Kompatsiaris Y. & Hobson P. (eds). Springer, 55–98.

Narayanan, S. 1999. Moving right along: a computational model of metaphoric reasoning about events. In *Proceedings of the National Conference on Artificial Intelligence*, 121–128. AAAI Press.

Niiniluoto, I. 1999. Defending abduction. *Philosophy of Science* **66**, 436–451.

Peirce, C. S. 1931. *Collected Papers of Charles Sanders Peirce*. Harvard University Press.

Peirce, C. S. 1982–. *Writings of Charles S. Peirce: A Chronological edition* **1**. Indiana University Press.

Pragglejaz Group, T. 2007. MIP: a method for identifying metaphorically used words in discourse. *Metaphor and Symbol* **22**(1), 1–39.

Rooth, M. 1985. *Association with Focus*, PhD thesis, University of Massachusetts.

Schurz, G. 2008. Patterns of abduction. *Synthese* **164**(2), 201–234.

Semino, E. 2008. *Metaphor in Discourse*. Cambridge University Press.

Snidaro, L., Garcia, J. & Llinas, J. 2015. Context-based information fusion: a survey and discussion. *Information Fusion* **25**, 16–31.

Snidaro, L., Visentini, I. & Foresti, G. 2011. *Data Fusion in Modern Surveillance* (chapter 1), Studies in Computational Intelligence **336**, 1–21. Springer.

Steinberg, A., Bowman, C. & White, F. 2004. Rethinking the JDL data fusion model. In *Proceedings of NSSDF*. JHUAPL.

Tenenbaum, J., Griffiths, T. & Kemp, C. 2006. Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences* **10**(7), 309–318.

Van der Sandt, R. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* **9**(4), 333–377.

Veale, T. 1998. 'Just in time' analogical mapping, an iterative-deepening approach to structure-mapping. In *Proceedings of ECAI98, the Thirteenth European Conference on Artificial Intelligence*.

Veale, T. 2014. A service-oriented architecture for metaphor processing. In *ACL 2014*, 52.

Veale, T. & Keane, M. T. 1997. The competence of sub-optimal structure mapping on hard analogies. In *The Proceedings of IJCAI'97, the International Joint Conference on Artificial Intelligence*.