Available online at www.sciencedirect.com
ScienceDirect
Cognitive Systems Research 39 (2016) 42–57
Cognitive Systems RESEARCH
www.elsevier.com/locate/cogsys

Image schemas in computational conceptual blending

Maria M. Hedblom [a], Oliver Kutz [b,*], Fabian Neuhaus [a]

[a] Institute of Knowledge and Language Engineering, Otto-von-Guericke University of Magdeburg, Germany
[b] KRDB Research Centre for Knowledge and Data, Free University of Bozen-Bolzano, Italy

Received 10 April 2015; received in revised form 8 October 2015; accepted 16 December 2015
Available online 13 February 2016

Abstract

In cognitive science, image schemas are identified as fundamental patterns of cognition. They are schematic prelinguistic conceptualisations of events and serve as conceptual building blocks for concepts. This paper proposes that image schemas can play an important role in computational concept invention, namely within the computational realisation of conceptual blending. We propose to build a library of formalised image schemas, and illustrate how they can guide the search for a base space in the concept invention work flow. Their schematic nature is captured by the idea of organising image schemas into families. Formally, they are represented as heterogeneous, interlinked theories.

Keywords: Computational creativity; Conceptual blending; Concept invention; Image schemas; Embodiment; Spatial cognition

1. Introduction

Computational creativity has seen significant progress in the last decade. Using a variety of artificial intelligence techniques there are now a multitude of systems that paint, write poems and solve problem (see the recent overview Besold et al., 2015). In this field the notion of 'creativity' is typically understood as a cognitive process defined and evaluated based on *degree of novelty* and *usefulness* of the resulting artefact (Boden, 2009; Runco & Jaeger, 2012). While humans are creative on a daily basis, computer systems still struggle to consistently produce output that human evaluators would deem creative.

The cognitive mechanisms behind human concept generation and understanding, are still largely unknown. Cognitive psychology and developmental linguistics have yet to provide a holistic explanation of the human capacity to learn concepts and from these generate new ones. Naturally, this therefore becomes difficult to model computationally. However, there are promising approaches that describe aspects of it. This paper investigates two of these theories: *conceptual blending* and *image schemas*. Built on the cognitive mechanisms behind analogical thinking, the theories provide some of the fundamental parts to the puzzle of human concept formation.

Conceptual blending is presented as the cognitive process behind creative thinking and generation of novelty in Turner (2014). The idea is that novel concepts are created when already known (and potentially conflicting) conceptual spaces[1] are merged into a new conceptual space, which, due to the unique combination of information, exhibits emergent properties.

One critical step in blending is the identification of shared structure across the different input domains. While

* Corresponding author.

[1] These are also called mental spaces in Fauconnier and Turner (1998) and are not to be confused with the 'conceptual spaces' in the sense of Gärdenfors (2000).

http://dx.doi.org/10.1016/j.cogsys.2015.12.010
1389-0417/

humans do this more or less automatically, this is one of the more complicated aspects of modelling conceptual blending formally. The main hypothesis of this paper is that image schema may play a vital role in identifying such shared structure.

While conceptual blending deals with already established concepts and knowledge, the theory of image schemas aims to explain some of the fundamental properties of concepts. Stemming from the embodied mind theory, image schemas are hypothesised to capture abstractions that model affordances related to spatio-temporal processes and relationships (Kuhn, 2007). In the cognitive sciences, *image schemas* are identified as the fundamental patterns for the cognition of objects, which are perceived, conceptualised and manipulated in space and time (Mandler & Pagán Cánovas, 2014). Examples of image schemas, proposed in the literature, are Containment, Support and Source_Path_Goal.

In this paper, we argue that combining conceptual blending with image schemas may not only shed light on the phenomenon of concept generation and creative thinking in humans, but also provide a useful tool for computational concept invention in computational creativity (Kutz, Bateman, Neuhaus, Mossakowski, & Bhatt, 2014; Schorlemmer et al., 2014).

The paper is structured as follows: in Section 2, the theory of image schemas is introduced. This section also includes an illustration of the ubiquity of image schemas in existing applied ontologies, and a discussion of related (formal) work on image schemas. This is followed, in Section 3, by a brief introduction of conceptual blending and a discussion on how conceptual blending can be computationally modelled and implemented. Section 4 discusses how image schemas can provide heuristics in the computational blending process. As these heuristics are based on organising image schemas into families of closely related theories rather than seeing them as individual theories, this idea is discussed in more details in Section 5, including a discussion of formal and algorithmic aspects of the proposal. We conclude the paper with a short summary and outlook to future work.[2]

2. Image schemas

This section presents the basic theory of image schemas. We begin, in Section 2.1, by introducing the central ideas with the help of a number of examples. We continue in Section 2.2 with an analysis of definitions of the notion of image schema found in the literature. We then, in Section 2.3, illustrate the prevalence of concepts closely related to image schemas in existing applied ontologies, before we conclude this introduction to image schemas with a discussion of related (formal) work in Section 2.4.

2.1. The basic idea illustrated by examples

Embodied theories of cognition (Barsalou, 2008) emphasise bodily experiences as the prime source for concept formation. Based on this cognitively supported view (Gallese & Lakoff, 2005), the theory of image schemas suggests that our conceptual world is grounded in the perceptive spatial relationships between objects.

Founded on psychological research (Mandler, 2004), the theory states that image schemas are formed as infants have repeated perceptual experiences, e.g. a plate being placed on a table. From this a generalisation emerges, an image schema, capturing the spatial relationships between the objects involved in an event. In the mentioned example, the image schema of Support is learnt. The understanding that plates can be placed on tables can be generalised and analogically transferred to other situations and objects. This means that infants who have learnt the Support schema through exposure to a plate on a table also grasp the notion of a book lying on a desk, as this represent the same spatial relationship. The more experience an infant has with a particular image schema, the more it becomes fine-tuned to accommodate different situations. Mandler (2008) describes how children can be observed to mentally expand image schemas such as the Support schema by adding information, for example when understanding that a large part of an object needs to be on the supporting surface.

Another image schema example is the notion of Containment, the notion that an object can be within a border (two-dimensional), or inside a container (three-dimensional). The image schema also includes the events of entering and exiting.[3]

The Containment schema is one of the most investigated image schemas (Johanson & Papafragou, 2014) as it is one of the first to be developed (Mandler, 1992), and since the relationships of enclosure and containment are essential for understanding our physical surroundings. It forms early as infants are immediately exposed to many situations in which objects are contained within one another, e.g. an embrace, lying in a crib, going into a house, eating food, etc.

One important aspect of image schemas is that they can be combined with one another. The image schema Path can easily merge with the image schema Link, leading to the more complex image-schematic concept Linked_Path. As Path illustrates a movement through space, and Link illustrates the causal relationship between two (or more) objects, a Linked_Path represents joint movement on two paths; e.g., a truck and trailer moving along a highway, or the joint movement of two separate magnets.

The 'cognitive benefit' of image schemas is to provide a means for information transfer. The conceptual

[2] This paper is a revised and extended version of Hedblom, Kutz, and Neuhaus (2014).

[3] It can be argued that In and Out are by themselves image schemas, or spatial primitives. For now we include them under the umbrella schema of Containment.

abstraction that constitutes the image schema can be utilised to explain unknown relationships and affordances of objects. The core idea is that after an image schema has been formed, it can be generalised and the structure can be transferred through analogical reasoning to other domains with similar characteristics (Mandler, 1992). That is, an image schema structure may be used as a conceptual skeleton in an analogical transfer from the concrete spatial domain of the image schema to another domain. This target domain may involve quite abstract concepts. Traces of this can often be seen in how language is used to explain more abstract concepts. It can be argued that much of metaphorical language is based on sensory-motor experiences and, thus, involves image schemas.

For example, processes and time are often conceptualised as objects and spatial regions. Expressions such as 'we meet *on* Thursday', map information from a concrete situation such as 'a book on a table' to the abstract process and time period. Another example is our conceptualisation of relationships like love or marriage, which also are often based on spatial metaphors. For example, one way to view a marriage is as LINKED_PATH, where the PATH represents how two spouses move together trough time and the LINK between them is the bond they share. A sentence like *Their marriage chains them together* works only if one conceptualises the relationship as a LINKED_PATH, because it reinterprets the LINK as an element that constraints the movements of both lovers. Alternatively, marriage may also be conceptualised as CONTAINMENT. This is reflected by metaphors like 'marriage is a prison', 'marriage is a safe harbour', and 'open marriage'. Depending on whether one chooses CONTAINMENT or LINKED_PATH as a base for the conceptualisation of marriage, a different vocabulary and different metaphors are supported.

The examples illustrate how image schemas may be used to conceptualise an abstract domain. As mentioned above, the first image schemas are developed by infants at an early stage where abstract thought is not yet present. This illustrates how concrete reasoning involving physical objects can provide the basis for the conceptualisation of the world and the formation of more abstract concepts.

2.2. Defining "image schema"

The term "Image schema" is hard to define properly. Image schemas are studied in several disciplines and from various perspectives, including neuroscience (Rohrer, 2005), developmental psychology (Mandler, 1992), cognitive linguistics (Hampe & Grady, 2005) and formal approaches (St. Amant et al., 2003). This broad range of research has lead to incoherence in the use of terminology. Also, the disputed relationship between socio-cultural aspects and the neurobiology of embodied cognition (Hampe, 2005) complicates the literature on image schema research.

Oakley defines an image schema as "*...a condensed re-description of perceptual experience for the purpose of mapping spatial structure onto conceptual structure*" (Oakley, 2007, p. 215). Mark Johnson describes them as "*...a recurring, dynamic pattern of our perceptual interactions and motor programs that gives coherence and structure to our experience*" (Johnson, 1987, p. xiv). Kuhn (2007) considers image schemas as the pre-linguistic structures of object relations in time and space.

One issue of these explanations of image schemas is that they do not provide individuation criteria. Hence, it is hard to evaluate whether a proposed image schema qualifies as such or not. The situation is complicated by the fact that image schema may change and become more specialised during the development of a child (Mandler & Pagán Cánovas, 2014). It is sometimes not obvious whether two conceptual structures are just variants of the same image schema or whether they are different image schemas.

One important attempt to structure the technical terminology of image schemas is made by Mandler and Pagán Cánovas (2014). In their paper they suggest to refine the umbrella term 'image schema' by distinguishing three different levels (p. 17):

1. *Spatial primitives*. The first building blocks that allow us to understand what we perceive: PATH, CONTAINMENT, THING, CONTACT, etc.
2. *Image schemas*. Representations of simple spatial events using the primitives: PATH OF THING, THING INTO CONTAINER, etc.
3. *Schematic integrations*. The first conceptual representations to include non-spatial elements, by projecting feelings or non-spatial perceptions to blends structured by image schemas.

From our perspective, this terminology provides the benefit of clearly distinguishing between image schemas and their building blocks (the spatial primitives). An image schema always represents an event and, thus, has some temporal dimension. The spatial primitives are the components that are participating in the event. E.g., according to this terminology PATH is not an image schema but a spatial primitive. In contrast, MOVEMENT ON PATH is an image schema. Another benefit is that it provides a clear criterion for distinguishing two image schemas (or schematic integrations): if x and y involve different spatial primitives, then x and y are different.[4]

Mandler and Pagán Cánovas approach provides a useful way to explain how conceptualisations are refined: an image schema is a representation of some kind of spatial event involving a number of spatial primitives. Hence, an image schema may be enriched by adding spatial primitives, yielding a more complex image schema. E.g., by adding the spatial primitives CONTAINER and INTO to the image

[4] Note that this is a sufficient condition, but not a necessary one, since two different representations may involve the same spatial primitives arranged in different ways.

schema Movement on Path, we obtain the schema Movement on Path into Container. This new image schema is more specific and less universally applicable. However, it provides more specific information when it is utilised conceptualising analogous situations. It follows that image schemas can be ordered into a hierarchy ranging from general image schemas, which contain only few spatial primitives, to more specific image schemas, which contain more spatial primitives.[5] Hence, image schemas do not exists in isolation but can be organised (at least) with respect to their (shared) spatial primitives. This observation is discussed further in Section 5.

In the following we continue to use "image schema" as the umbrella term for the three levels of conceptualisations. To avoid any ambiguity, we will refer to image schemas in the sense of Mandler and Pagán Cánovas as *spatial schemas*.

2.3. The ubiquity of image schemas

Image schematic notions play a central role in many efforts aiming to capture common sense knowledge. In this section, we illustrate the ubiquity of notions closely related to image schemas in existing applied ontologies. We will focus on the notion of Containment and discuss some prominent ontologies that incorporate them in various ways. Similar overviews could be generated for other prominent image-schematic notions such as 'Path', 'Link' and 'Support'.

Image schematic notions can be found early on in efforts such as building the Cyc knowledge-base, (Lenat, Prakash, & Shepherd, 1985), or in the collections[6] of common sense modelling problems. Morgenstern's 'A Case Study in Egg Cracking' (Morgenstern, 2001) contains extensive axiomatisations of variants of containment, and Cyc includes a variety of notions of containment and path-following at its most general levels of knowledge modelling and categorisation.

'Containment' is a crucial notion in areas such as geography and transportation (Egenhofer & Mark, 1995), anatomy and bio-medicine in general (Smith et al., 2005), linguistics and cognition (Bateman, Hois, Ross, & Tenbrink, 2010; Reed & Pease, 2015), or indeed cooking (Krieg-Brückner, Autexier, Rink, & Nokam, 2015).

The relevant notions of containment range from down-to-earth notions of containment such as 'holding milk in a cup' to semiotic and information-theoretic notions such as 'signs holding information' or '.tex files holding UTF8 characters' to fully abstract versions such as the 'class containing all twin primes'. In addition to the variety of concrete versions of containment, also the levels of formalisation differ dramatically, namely from extensive (at least) first-order based axiomatisations such as in Morgenstern (2001), to more light-weight axiomatisations as found e.g. in GUM-Space[7] (Bateman et al., 2010), to mere annotation of concepts or relations, as it is common practice in bio-medical terminologies and ontologies (see below).

We will now present a number of concrete containment notions as they can be found in prominent ontologies. Namely, we will give examples of containment notions from the areas or architecture, natural language, biomedicine, and cultural heritage. We will list them roughly in the order of concrete to abstract.

The Industry Foundation classes IFC[8] is an object-oriented data model to support data exchange in (among others) the areas of architecture and built environments, and can be seen as an application ontology. Notions of contact, containment, and composition are central to such engineering contexts (see e.g. Bhatt, Hois, & Kutz, 2012; Hois, Bhatt, & Kutz, 2009). As described in Bazjanac et al. (2002):

> The compositional aspect of the connectivity model supports three possible relationships: aggregation, containment and nesting. [...] Containment implies a stronger form of composition, where the components cannot be considered independently: where the definition of the whole element depends on the definition of its parts and the parts depend on the existence of the whole element. [...]

When investigating formal reasoning approaches for such notions of containment, often qualitative spatial calculi such as the Region Connection Calculus (RCC8) are employed (Randell, Cui, & Cohn, 1992), as discussed in Bhatt et al. (2012).

Moving on from engineering to natural language, according to WordNet, a 'container' is described as a noun:

> container (any object that can be used to hold things).

In addition to this rather general version of containment (which however seems to exclude abstract containers), WordNet lists the most typical instantiation as:

> (especially a large metal boxlike object of standardised dimensions that can be loaded from one form of transport to another).

Moreover, the Synset for 'container' reports about 100 direct hyponyms, reflecting the prevalence of container-like terms in natural language, and includes terms such as 'basket', 'spoon', and 'time capsule'.

The linguistic ontology GUM (the 'Generalised Upper Model') (Bateman et al., 2010) specifies detailed semantics for linguistic spatial expressions. GUM-Space (the space-related module of GUM) specifies 'containment' as

[5] In their list of spatial primitives, Mandler and Pagán Cánovas include move, Animated Move, and Blocked Move. This seems to suggests that the spatial primitives are ordered into a subtype hierarchy, since both animated movement and blocked movement are a kind of movement.

[6] See http://www-formal.stanford.edu/leora/commonsense/.

[7] See http://ontologydesignpatterns.org/wiki/Ontology:GUM-Space.

[8] IFC is registered by ISO and is an official International Standard ISO 16739:2013.

a specific kind of 'FunctionalSpatialModality' exhibiting 'Control', namely:

> The reified functional relation holding between two spatial objects x and y, such that x functionally contains y; x need not spatially contain y. An example of an expression falling into this category is: ``The apple is *in* the bowl". Here, the apple does not necessarily need to be spatially contained in the bowl (no topological containment).

The GUM ontology was axiomatised both in first-order logic and description logic variants.[9]

Quite in contrast to this, though containment is an important notion in biomedicine and can be found in a large number of ontologies hosted on Bioportal,[10] here the meaning of 'containment' is typically specified via annotations. E.g., according to the Ontology of Biomedical Investigations[11]:

> A device that can be used to restrict the location of material entities over time.

Notice that unlike the variants of containment notions mentioned above, we here find an explicit conceptualisation of the *temporal aspects* of containment. A more detailed formal treatment of various relations relevant in biomedical ontologies can be found in Smith et al. (2005). There, 'containment' is characterised in first-order logic with the help of mereological notions. The informal characterisation is:

> Containment is location not involving parthood, and arises only where some immaterial continuant is involved. [...] Containment obtains in each case between material and immaterial continuants, for instance: *lung contained_in thoracic cavity*; *bladder contained_in pelvic cavity*. Hence containment is not a transitive relation (Smith et al., 2005).

We close this section with an example of an abstract, non-physical notion of containment, namely the idea of a document as a container of 'information'. The CIDOC (Doerr, 2003) Conceptual Reference Model (CRM) provides an ontology for concepts and information in the domain of cultural heritage and museum documentation.[12]

We here find the abstract notion of a document as a container of 'information', e.g. the CIDOC notion of an 'Information carrier':

> This class comprises all instances of E22 Man-Made Object that are explicitly designed to act as persistent physical carriers for instances of E73 Information Object.

Similarly, motivated by the need to cover concepts related to UML to model aspects of information systems, the Cyc knowledge base was enriched by adding notions of 'abstract containment' (Reed et al., 2002).

2.4. Related work on formalising image schemas

Image schemas are a well studied field in research on cognitive linguistics and developmental psychology. Recently, more computationally-oriented research has shown an increased interest in image schemas as a route to approach new (partial) solutions to the symbol grounding problem and to aid computational concept invention (Goguen & Harrell, 2010; Kuhn, 2002; Kutz, Bateman, et al., 2014; Morgenstern, 2001).

Lakoff and Núñez (2000) used image schemas extensively in their reconstruction of abstract mathematical concepts using blending and image schemas. Working from the perspective that all of mathematics can be deduced from the body's interactions with its environment, they give a detailed account on how image schemas provide some of the conceptual principles that provide a grounding of abstract concepts.

While Lakoff and Núñez's effort is not a formalisation of image schemas, their attempt to ground mathematics in embodied cognition has been further developed and formalised. Guhe et al. (2011) account for the ideas in Lakoff and Núñez (2000) by formalising in first-order logic some basic mathematical constructs such as the *measuring stick*, *motion along a path*, and *object construction*. Using the analogy engine Heuristic Driven Theory Projection, HDTP, they illustrate how generalisations such as image schemas could be used to transfer information in a computational system. Their system uses anti-unification to find the common structure in both source and target domain. This common structure is used to transfer information to the target domain from the source.

St. Amant et al. (2003) introduced the Image Schema Language, ISL, in which they discuss how image schemas can be represented and simulated computationally. They argue that their representation provides a structured image schema description of a situation. They use three different scenarios to discuss simulations of image schemas: a Chess game, military tactics, and a robot simulation. The Chess game example is particularly interesting as it accounts for the two-dimensional, spatial relationships between the pieces on the Chess board, constrained by the rules of the game. Using a combination of the image schemas Containment, Link and Path, they illustrate how the board configurations can be viewed from higher conceptual perspectives (rather than simply as spatial configurations).

Kuhn (2002) presented another approach where he used Wordnet to extract meaning from words, and employed the programming language Haskell to generate testable models. In Kuhn (2007), he extended his previous image schema research by presenting a method to account for spatial categorisation and developing an algebraic theory formalising image schemas. Here he argues that the image schemas capture the abstractions essential to model affordances. For example, a cup is a cup because it can contain liquid, or an object is a vehicle when it affords

---

[9] See http://www.ontospace.uni-bremen.de/ontology/gum.html.

[10] See http://bioportal.bioontology.org.

[11] See http://bioportal.bioontology.org/ontologies/OBI.

[12] It is registered as international standard (ISO 21127:2014) for the controlled exchange of cultural heritage information.

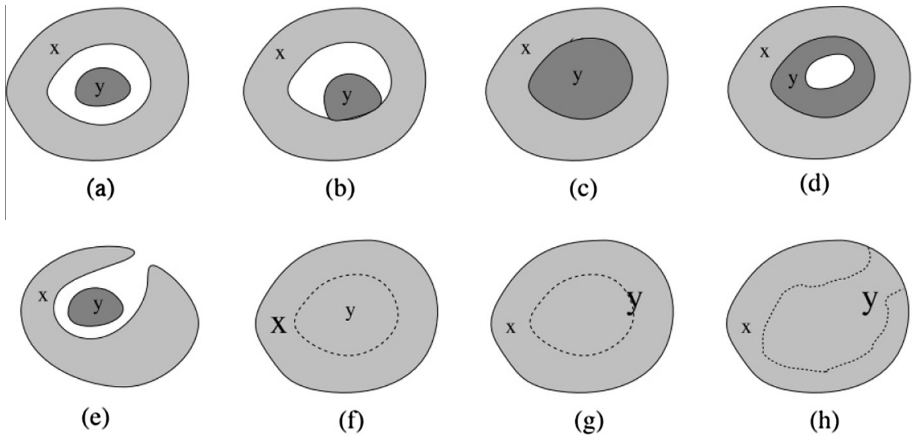

Fig. 1. The eight variations of Containment as discussed in Bennett and Cialone (2014).

transportation. With Kuhn's reasoning a vehicle can be described with a combination of the image schemas Support (alternatively Containment) and Path.

Acquired from natural language, Bennett and Cialone (2014) formally represented several different kinds of Containment schemas. They distinguish eight different spatial Containment relationships and their mappings to natural language constructs, illustrated in Fig. 1. Their work also demonstrates the non-trivial nature of formalising image schemas, and that there are many closely related variants of any given image schema.

3. Conceptual blending

3.1. A short introduction to conceptual blending

The theory of *Conceptual blending* was introduced during the 1990s as the cognitive machinery for novel concept generation (Fauconnier & Turner, 1998). The theory aims to explain the process behind creative thinking. It has strong support from research in cognitive psychology and linguistics (Gibbs, 2001; Grady, 2001; Yang, Bradley, Huq, Wu, & Krawczyk, 2012) as well as in more computational areas (Goguen & Harrell, 2010; Veale, 2012).

According to conceptual blending theory, generation of novel concepts occurs via the combination of already existing ideas and knowledge. It is suggested that such novel concepts are selective and 'compressed' combinations, or blends, of previously formed concepts, building on the notion that all novel generation builds from already existing knowledge. This cognitive process is thought to happen as two, or more, input domains, or information sources, are combined into a new domain, the blended domain, see Fig. 2. The blend inherits some of the attributes and relationships from the source domains and at the same time the unique mix allows the blends to have emergent properties that are unique to each particular blend.

Conceptual blending can be compared to the cognitive mechanisms behind analogical reasoning. In analogical reasoning information flows from a source domain to a target domain by using cognitive structure-mapping mechanisms. Conceptual blending is comparable insofar it employs a search for 'similar structure' in the two input domains, information then gathered in the generic space, the base space.[13] The abstracted structure found in the base ontology is later used to structure the blend as well.

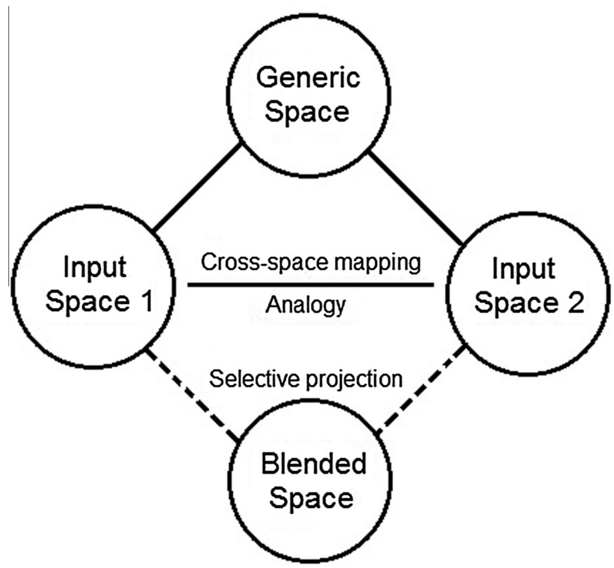

Fig. 2. The blending process as described by Fauconnier and Turner (1998).

Many monsters are examples for conceptual blends. For example, a griffin is a fictive creature with the body and the tail of a lion and with the head and the wings of an eagle. The blend of the two creatures does not just involve the physical attributes of the animals, but also the characteristics associated with them. The lion provides attributes such as strength and power, and the eagle precision and capacity

[13] Introduced by Fauconnier and Turner as *generic space*, the notion carries the name *base space* or *base ontology* in formal approaches.

for flight. Hence, the blended creature has the skills to master both land and sky.

The griffin exemplifies *one particular* blend of the two input spaces *lion* and *eagle*. There are other possibilities to blend a monster based on these two concepts. For example, one could consider an ‘inverted griffin’, which has the head of the lion and the body of the eagle but no wings. A third possible monster is a creature which has the shape and strength of a lion but cannot use its strength because of its fragile bird-like bone structure. The last example shows that not all blends are equally successful. In order for the blend to be considered creative, the blend needs to be “useful” (Boden, 2009). Given the task of blending a monster, a successful blend is required to produce a dangerous creature – a lion with brittle bones does not meet this requirement as well as a griffin.

The blended space preserves the information from the generic space. However, usually only some selected features of the input spaces are retained. In the griffin example, the generic space contains the head, the body, and two limbs of a vertebrate. In the blend, the head in the generic space is mapped to the head of the lion and the head of the eagle, respectively. The same holds for the body. In contrast, the two limbs are mapped to the forelimbs of the lion and the hindlimbs (legs) of the eagle. For this reason, the griffin has six limbs, namely two wings of the eagle, two hindlegs from the lion and two forelegs, which are inherited from both input spaces. Since the shape and features of lion legs and eagle legs are mutually exclusive (e.g., one has hair and the other has feathers), the forelegs of the griffin cannot inherit all properties from both input spaces. Thus, griffins forelegs are usually conceptualised as exemplifying either only the features of one animal or as inheriting a consistent subset of features from both input spaces.

For humans conceptual blending is effortless. We are able to create new blends spontaneously and have no difficulty to understand new conceptual blends when we encounter them. This include the selection of suitable input spaces, the identification of a relevant generic space, the identification of irrelevant features of the input spaces, the performance of the blend, and the evaluation of the usefulness of the blend. In contrast, for an automated system each of these steps provides a significant challenge. In the next section we discuss a formal, logic-based model for conceptual blending.

3.2. Formalising conceptual blending

We formalise conceptual blending following an approach based on Goguen’s (1999) work on *algebraic semiotics* in which certain structural aspects of semiotic systems are logically formalised in terms of algebraic theories, sign systems, and their mappings. In Goguen and Harrell (2010) algebraic semiotics has been applied to user interface design and conceptual blending. Algebraic semiotics does not claim to provide a comprehensive formal theory of blending – indeed, Goguen and Harrell admit that many aspects of blending, in particular concerning the meaning of the involved notions, as well as the optimality principles for blending, cannot be captured formally. However, the structural aspects *can* be formalised and provide insights into the space of possible blends. The formalisation of these blends can be formulated using languages from the area of algebraic specification, e.g. OBJ3 (Goguen & Malcolm, 1996).

In Hois, Kutz, Mossakowski, and Bateman (2010), Kutz, Mossakowski, Hois, Bhatt, and Bateman (2012), and Kutz, Neuhaus, Mossakowski, and Codescu (2014), an approach to computational conceptual blending was presented, which is in the tradition of Goguen’s proposal. In these earlier papers, it was suggested to represent the input spaces as ontologies (e.g., in the OWL Web Ontology Language[14]). The structure that is shared across the input spaces, i.e. the generic space, is also represented as an ontology, which is linked by mappings to the input spaces. As proposed by Goguen, the blending process is modelled by a colimit computation, a construction that abstracts the operation of disjoint unions modulo the identification of certain parts specified by the base and the interpretations, as discussed in detail in Goguen (2003), Kutz, Mossakowski, and Lücke (2010), and Kutz et al. (2012).

The inputs for a blending process (input concepts, generic space, mappings) can be formally specified in a *blending diagram* in the Distributed Ontology, Model, and Specification Language (DOL).[15]

DOL is a metalanguage that allows the specification of (1) new ontologies based on existing ontologies, (2) relations between ontologies, and (3) networks of ontologies, including networks that specify blending diagrams. These diagrams encode the relationships between the base ontology and the (two or more) input spaces. The blending diagrams can be executed by the *Heterogeneous Tool Set* Hets, a proof management system. Hets is integrated into Ontohub,[16] an ontology repository which allows users to manage and collaboratively work on ontologies. DOL, Hets, and Ontohub provide a powerful set of tools, which make it easy to specify and computationally execute conceptual

[14] With ‘OWL’ we refer to OWL 2 DL, see http://www.w3.org/TR/owl2-overview/.

[15] Regarding blending diagrams as displayed in Fig. 3, notice the following discrepancy in terminology and in the way the basic blending process is visualised. In the cognitive science literature following (Fauconnier & Turner, 1998), conceptual blending is visualised as shown in Fig. 2, with a *generic space* at the top identifying commonalities. In the technically oriented literature following (Goguen & Harrell, 2010), the formalisation of this process is represented as a diagram as shown in Fig. 3. This kind of diagram is on the one hand an upside-down version of the first illustration, following traditions of category theory to put the ‘simpler’ objects at the bottom of a diagram. On the other hand, it replaces the term ‘generic space’ with ‘base space’, partly because of a clash with mathematical terminology. In our work on formalisation of blending, we will make no technical difference between ‘generic space’ and ‘base space’ and treat them as synonymous.

[16] www.ontohub.org.

blends, as seen in Neuhaus et al. (2014). An extensive introduction to the features and the formal semantics of DOL can be found in Mossakowski et al. (2015).

As illustrated with the example in the previous section, a critical step in the blending process is the identification of the common structure of the generic space and its mapping to the input spaces. The structural similarity between conceptual blending and analogical thinking suggests to investigate and apply approaches to analogical reasoning as tools for computational conceptual blending.

One important theory in analogical research is the Structure Mapping Theory (Gentner, 1983). It claims that analogical reasoning is characterised by the relationships between objects rather than their attributes. Following this idea is the analogy engine Heuristic Driven Theory Projection, HDTP (Schmidt, Krumnack, Gust, & Kühnberger, 2014). HDTP computes a 'least general generalisation' B of two input spaces $O1$ and $O2$. This is done by anti-unification to find common structure in both input spaces $O1$ and $O2$. HDTP's algorithm for anti-unification is, analogously to unification, a purely syntactical approach that is based on finding matching substitutions.[17]

While this is an interesting approach, it has a major disadvantage. Typically, for any two input spaces there exists a large number of potential generalisations. Thus, the search space for potential base spaces and potential conceptual blends is vast. HDTP implements heuristics to identify interesting anti-unifiers; e.g., it prefers anti-unifiers that contain rich theories over anti-unifiers that contain weak theories. However, since anti-unification is a purely syntactical approach, there is no way to distinguish cognitively relevant from irrelevant information. As a result, an increase of the size of the two input ontologies leads to an explosion of possibilities for anti-unifications.

4. Blending with image schemas

Instead of relying on a purely syntactical approach to blending, the semantic content found in image schemas can be employed to help guiding the blending process. The basic idea here is that in order to identify common structure sufficient for defining a useful generic space for two (or more) given input spaces, we search for shared image-schematic information rather than *arbitrary* structure. As discussed above, a vast space of blends opens up if we work with more unconstrained resp. syntax-based shared structure in the generic space. Given the powerful role that image schemas generally seem to play in human conceptual (pre-linguistic) development, the working hypothesis is that the semantic content and cognitive relevance given by identifying shared image schemas will provide valuable information for constructing and selecting the more substantial or interesting possible blends.

This section therefore serves a twofold purpose. First, we will demonstrate in Section 4.1 that image schemas may enable similes based on a shared containment structure, and show how this extends to supporting the conceptual blending process. Second, in Sections 4.2 and 4.3, we will give formalised versions of blends where image schemas play a crucial role, showing that the gulf between the cognitive relevance of image schemas and formal, logic-based concept blending can be bridged.

4.1. Blending with image schemas in natural language

In this section we show that image schemas can provide the base structure for the blending of a wide variety of concepts.

Consider the concepts *Space Ship*, *North Korea*, *Spacetime*, *Marriage* and *Bank account*. Note that these concepts differ significantly. However, all of them can be construed as various kinds of *containers*. This is obvious in the case of space ships, which may contain passengers and cargo. Geopolitical entities like North Korea instantiate the CONTAINMENT schema, since they have boundaries and people may be inside and outside of countries. *Spacetime* conceived as a container is a particularly interesting case since it implies the notion of inertial frames of reference, which is arguably inconsistent with the Theory of Relativity (DiSalle, 2009). This does not prevent science fiction writers to construe spacetime as a container for planets, suns and other things; in many fictive stories it is possible to leave and return to the universe (e.g., by visiting a 'parallel universe'). While the first three examples are physical entities, *Marriage* is a social entity. Thus, in the literal physical sense marriage cannot be a container. Nevertheless, we use vocabulary that is associated with containers to describe marriage. E.g., one can *enter* and *leave* a marriage, some marriages are *open*, others are *closed*, and people may find happiness *in* their marriage. Similarly, a *bank account* may *contain* funds, and if it is *empty* we can put some additional funds *into* the account and take them *out* again later. These linguistic examples provide some evidence that we conceptualise *Marriage* and *Bank account* as kinds of *containers*.

The claim that these five concepts are indeed instantiating CONTAINMENT is supported by the behaviour of these concepts in *similes*. The first column ('target domain') of Table 1 contains our examples. The second column ('source domain') contains various concepts of physical containers which highlight some possible features of containers: e.g., a container may leak, be hard to get out of, or have a flexible boundary. Let us consider the similes *X is like a Y* that are the result of randomly choosing an element *X* from the first row and combining it with a random element *Y* from the second column. For example, 'The universe is like a treasure chest', 'Their marriage is like a prison', 'My bank account is like a leaky pot'. Note that all of the resulting similes are meaningful. Some of them

[17] There are several other methods for finding generalisations. One example is the Analogical Thesaurus (Veale, 2003) which uses WordNet to identify common categories for the source and target spaces.

Table 1
Containment similes: ⟨target⟩ is like a ⟨source⟩.

Target domain	Source domain
This space ship	Leaky pot
North Korea	Prison
The universe	Treasure chest
Their marriage	Bottomless pit
My bank account	Balloon

will intuitively have more appeal than others, which may only be meaningful within a particular context.[18]

The fact that Table 1 can be used to randomly produce similes is linguistically interesting, because the target concepts vary significantly. The concepts *space ship*, *marriage* and *North Korea* seem to have nothing in common. Therefore, the fact that they can all be compared meaningfully to the same concepts needs an explanation. The puzzle is solved if we assume all concepts in the first column share the underlying image schema Containment. For this reason they can be blended with the container concepts from the second column. In each simile we project some feature of the container in the source domain (second column) via analogical transfer onto the container aspect of the target domain (first column). Thus, Table 1 provides evidence that image schemas can help us to identify or (construe) shared structure between concepts.

The shared structure between concepts can be utilised in conceptual blending. For example, we can conceptually blend the concepts *universe* and *balloon* to a *balloon-universe*, that is a universe that continuously increases its size and expands. This concept is already lexicalised as *expanding universe* in English. Blending *space ship* with *prison* could lead to various interesting concepts: e.g., to a space ship that is used as a prison – a kind of space age version of the British prison hulks of the 19th century.

It is also possible to attempt to blend two different concepts from the first column from Table 1. However, since these concepts contain more prominent aspects than Containment, these blends may not involve the Containment as shared structure. E.g., a in a blend of *Space Ship* and *North Korea* probably other aspects of the concept of North Korea would be more dominant. E. g., a *North Korean Space Ship* may be, trivially, a space ship built in North Korea or a space ship with a dictatorial captain and a malnourished crew. Only by providing some additional context one can prime the Containment aspect of North Korea; e.g., in ‘People inside North Korea do not learn anything about the rest of the world, from their perspective they live in the space ship North Korea, which is surrounded by an empty void’.

Let us consider two different examples from our list. A blend of *marriage* and *bank account* may yield the concept of a *marriage account*. This new concept could be used in sentences like the following: ‘Marcus and Susie have just spent a long and happy holiday together, this was a big investment into their marriage account, it is now full of love’ or ‘Jim needs to watch the way he treats Jill, their marriage account is draining quickly and is nearly empty. She is probably going to leave him’. In this blend the *marriage account* is a container which contains feelings between the spouses instead of money. The blend inherits the domain from *marriage* (with the major difference that the spouses themselves are no longer inside the container). The main contribution of *bank account* to the blend is the ability to ‘invest’ and ‘check the balance’ of the content in the *marriage account*.

How something is conceptualised depends on the context. For example, surgeons may conceptualise people as containers of organs, blood, and various other anatomical entities, but in most contexts we do not conceptualise humans in this way. By choosing the appropriate context an image schema may be pushed from the background into the conceptual forefront. For example, in most contexts a *mother* is probably not conceptualised as a kind of container. However, in the appropriate contexts it is possible to generate similes for *mother* reusing the source domains from Table 1; e.g., ‘The mother is pregnant with twins, she looks like a balloon’ or ‘The mother is like a prison for the unborn child’.

The examples that we have discussed in this section show how the Containment image schema can be utilised as generic space in conceptual blending. In the next sections we present the formalisation of the blending of two of our examples, namely *space ship* and *mother*.

4.2. The mother ship example

Our thesis is that image schemas provide a useful heuristics for conceptual blending, because shared image schemas are good candidates for the generic space in the blending process.

The concepts *space ship* and *mother* share the Containment schema. As a first step towards the formalisation of the blending process, we need to represent Containment in some formal language.

For the sake of illustrating the basic ideas, we choose here a simplified representation in OWL (see Fig. 5). Containers are defined as material objects that have a cavity as a proper part. A container contains an object if and only if the object is located in the cavity that is part of the container.[19]

As mentioned in Section 4.1, many concepts contain a rich structure. We do not attempt here to provide a full axiomatisation of *mother* or *space ship*, but just focus on

[18] For example, ‘This space ship is like a bottomless pit’ may sound odd in isolation, but in the context of ‘I have already 20.000 containers in storage, and there is still empty cargo space’ the simile works.

[19] This is a simplified view on Containment. E.g., a more accurate formalisation of the Containment schema would need to cover notions like moving into or out of the container.

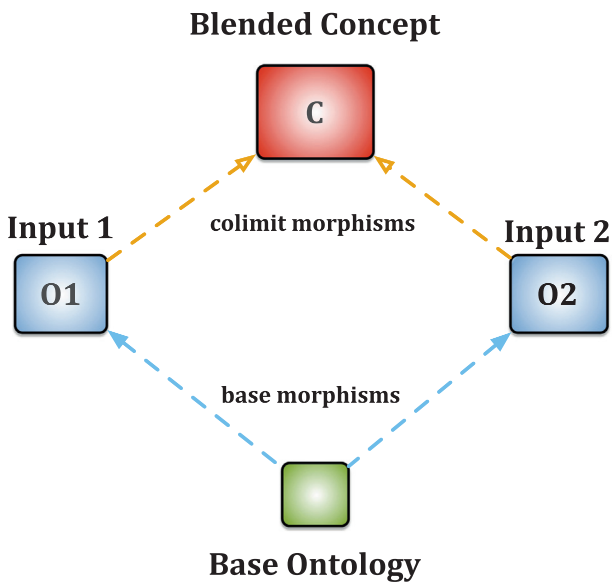

Fig. 3. The blending process as described by Goguen and Harrell (2010).

some salient points for the sake of illustrating the blending process.

As discussed in Section 4.1, *mother* realises the Containment schema, since mothers have a uterine cavity, which at some point in time contained some child. Further, *space ship* realises the Containment schema since space ships may be used to transport goods and passengers. Of course, in almost any other aspect mothers and space ships are completely different; in Fig. 6 we only represent that mothers are female humans with children and that space ships are capable of space travel.

During the blending of *mother* and *space ship* into *mother ship* the Containment schema structure of both input spaces is preserved (see Fig. 7). The uterine cavity and the cargo space are both mapped to the docking space. The *mother ship* inherits some features from both input spaces, while others are dropped. Obviously, a mother ship is a space travelling vessel. But like a mother, it is a 'parent' to some smaller entities of the same type. These smaller vessels can be contained within the mother ship, they may leave its hull (a process analogous to a birth) and are supported and under the authority of the larger vessel.[20]

To summarise, in our example we try to blend the input spaces of "Mother" and "Space ship". Instead of trying to utilise a syntactic approach like anti-unification to search for a base space, we recognise that both input spaces have cavities and, thus, are containers. Using the base space Containment in the blending process yields a blended concept of "Mother ship". Here, the precise mappings from the base space axiomatisation of Containment to the two input spaces regulate the various properties of the blended concept. Fig. 4 illustrates this blend by populating the generic blending schema shown in Fig. 3.

4.3. The satellite example

To further illustrate the role of image schemas in the construction of a newly blended concept, let us consider a second example. Assume we want to create a new concept by blending our mother ship with a moon. While this may not be astronomically completely correct, for the sake of this paper we consider a moon to be a celestial object that is part of some solar system, has a spheroidal shape, consists of rock, and orbits around a planet (see Fig. 8). Of course, many people would associate additional information with the concept *moon*, but even if we consider only these aspects, there are different possibilities how we could blend the two concepts. E.g., a structure mapping approach would probably first try to identify the parthood relationship between the docking station and the mother ship on one hand with the parthood relationship between the moon and the solar system. This may lead to the concept of a Moon/DockingStation that is part of a SolarSystem/MotherShip – While not being wrong, it might not be a useful concept.

In contrast, if one utilises shared image schemas as heuristics for conceptual blending, it is quite natural to look at a very different place for blending opportunities. Since the mother ship is a kind of vehicle it has the capability to move stuff or people, which involves movement from some place to another along a path. A moon also moves along a path, namely it's orbit. This commonality we can utilise in the blending process. However, in this case the situation is not as straightforward as in Section 4.2, because the movements of the mother ship and the moon are quite different and do not instantiate the same image schema.

When the mother ship (or any vessel) executes its capability in some movement process, the vehicle starts at some location of origin and moves along a path until it reaches its goal, where it stops. Thus, the image schema is Thing Moves On Path From Source To Goal. The orbital movement of the moon also follows some path. However, the orbital movement does not have a source or a goal; it is characterised by a focal point, which the orbiting object revolves around. Hence, the formal representation of both kinds of movement looks quite differently (see Fig. 9 for a representation in OWL).

As discussed in Section 2.2, spatial schemas, can be enriched by adding additional spatial primitives; the spatial schemas instantiated by the movement of a vessel and of a moon, respectively, are different (and mutually exclusive) refinements of Thing Moves On Path. For the purpose of blending the important lesson is that image schemas do not exist in isolation, but they are members of *families of image schemas*. The members of these image schema families are variants of some root conceptualisation (e.g.,

[20] To represent dynamic aspects like birth and vessels leaving a docking bay adequately, one needs a more expressive language than OWL.

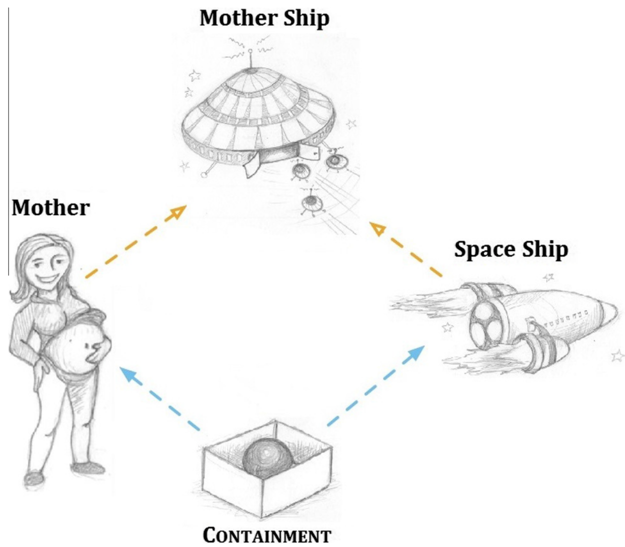

Fig. 4. The blending of mother ship.

movement) and can be partially ordered by their strength (see Section 5).

We can utilise this observation as a heuristic for conceptual blending: if two concepts involve two different image schemas, which are within the same image schema family, then a good candidate for the base space for blending both concepts is the least general member of the image schema family, which generalises the image schemas in the input spaces. In the case of our example, this is Movement on a Path. Further, the blended concept probably should include only one member of the image schema family. In our example, we can create a new concept that inherits the salient features of the mother ship, but replaces its ability to travel from one place to another by some orbital movement. The resulting theory describes a space station, which orbits around a planet or dwarf planet. Alternatively, we can think of a moon-like concept that is turned into a spacefaring vehicle. This is a kind of 'moon ship', that is a moon that has the capability to move from a location of origin along a path to a destination (see Fig. 10).

This example illustrates how the use of image schemas can provide heuristics for (i) identifying suitable base spaces and (ii) selecting features during running interesting blends – even if the input domains do not share exact structure. However, it raises the question how we should represent image schema families formally. This is a question we will discuss in the next section.

5. Image schema families as graphs of theories

In the previous section, we suggested that image schemas are members of families, which are partially ordered by generality. Formally, we can represent such a family as a graph of theories in DOL.[21] In this section, we discuss this approach in some more detail. On a technical level, our proposal for capturing image schemas as interrelated families of (heterogeneous) theories is quite similar to the ideas underlying the first-order ontology repository COLORE[22] (Grüninger, Hahmann, Hashemi, Ong, & Ozgovde, 2012).

In our blend of *mother ship* with *moon* we considered two variants of Movement on a Path. Fig. 11 shows a selection of some other members of the same image schema family.[23] One way Movement on a path can be specialised is as Movement on a looping path. Note that this change does not involve adding a new spatial primitive, but just an additional characteristic of the path. The resulting image schema can be further refined by adding the notion of a *focal point*, which the path revolves around – this leads to the notion of orbiting. Alternatively, we may change Movement on a Path by adding distinguished points; e.g., the source, the target, or both.

The latter image schema may be further specialised by identifying the source and the target. In this case the path is closed in the sense that any object which follows the path will end up at the location at where it started its movement (the source). The difference between a closed path and a looping path is that the closed path has a start and an end (e.g., a race on a circular track), while the looping path has neither (like an orbit). It is possible to further refine the schema by adding more designated points or other related spatial primitives.

The particular image schema family sketched in Fig. 11 is organised primarily via adding new spatial primitives to the participating image schemas and/or by refining an image schema's properties (extending the axiomatisation).[24] In general, different sets of criteria may be used depending, for example, on the context of usage, thereby putting particular image schemas (say, Revolve_Around) into a variety of families. Apart from a selection of spatial primitives, other dimensions might be deemed relevant for defining a particular family, such as their role in the developmental process.

In Bennett and Cialone (2014), eight closely related kinds of Containments were identified as being distinguishable within natural language corpora, illustrated in Fig. 1 and discussed above. Hence, the selection criteria for grouping together these particular forms of containment are not simply driven by a selection of spatial primitives. Although Bennett and Cialone (2014) do not explicitly formalise the structural relationships between

[21] These graphs are diagrams in the sense of category theory.

[22] See http://stl.mie.utoronto.ca/colore/.

[23] A disclaimer: in the following we will describe an approach to represent the connections between image schemas, belonging to the same family according to certain criteria. To illustrate some technical points, we will just postulate the existence of several image schemas and their connections. However, we here do not intend to make any claims regarding their empirical existence and/or their cognitive role in development.

[24] In Hedblom, Kutz, and Neuhaus (2015a,b) we present a more complete description of the image schema family of 'path following' and the corresponding formal methodology.

```
Class: Container
    EquivalentTo: MaterialObject and has_proper_part some Cavity

ObjectProperty: contains
    SubPropertyChain: has_proper_part o is_location_of
    DisjointWith: has_proper_part
    Domain: Container
    Range: MaterialObject
```

Fig. 5. A (partial) representation of CONTAINMENT in OWL.

```
Class: Mother
    EquivalentTo: Female and Human and parent_of some (Small and Human)
    SubClassOf: has_proper_part some UterineCavity

Class: SpaceShip
    EquivalentTo: Vehicle and has_capability some Spacefaring
    SubClassOf: has_proper_part some CargoSpace
```

Fig. 6. Mothers and space ships.

```
Class: MotherShip
    SubClassOf: Vehicle and has_capability some Spacefaring
    SubClassOf: has_proper_part DockingStation
    SubClassOf: parent_of some (Small and Vehicle)
```

Fig. 7. Mother ship.

```
Class: Moon
    EquivalentTo: CelestialObject and participates_in some
        (OrbitalMovement and revolves_around some (Planet or DwarfPlanet))
    SubClassOf: consists_of Rock
    SubClassOf: part_of some SolarSystem
    SubClassOf: has_shape some Spheroid
```

Fig. 8. Moon.

```
Class: Vehicle
    SubClassOf: has_capability some SourceToGoalMovementCapability
    SubClassOf: has_capability some TransportationCapability

Class: SourceToGoalMovement
    EquvialentTo: MovementProcess and
        (has_participant some MovingEntity) and
        (follows exactly 1 (Path and (has_source exactly 1 Location)
            and  (has_destination exactly 1 Location)))
    EquivalentTo: executes some SourceToGoalMovementCapability

Class: OrbitalMovement
    EquvialentTo: MovementProcess and
        (follows exactly 1 (LoopingPath and
                (revolves_around some owl:Thing)))

Class: LoopingPath
    EquvialentTo: Path and Looping
    SubClassOf: has_source only owl:Nothing
    SubClassOf: has_destination only owl:Nothing
```

Fig. 9. Movement.

```
Class: SpaceStation
    SubClassOf: has_capability some Spacefaring
    SubClassOf: has_proper_part DockingStation
    SubClassOf: parent_of some (Small and Vessel)
    SubClassOf: participates_in some
        (OrbitalMovement and revolves_around some (Planet or DwarfPlanet))

Class: MoonShip
    SubClassOf: Vehicle
    SubClassOf: consists_of Rock
    SubClassOf: has_shape some Spheroid
```

Fig. 10. Space station and moon ship.

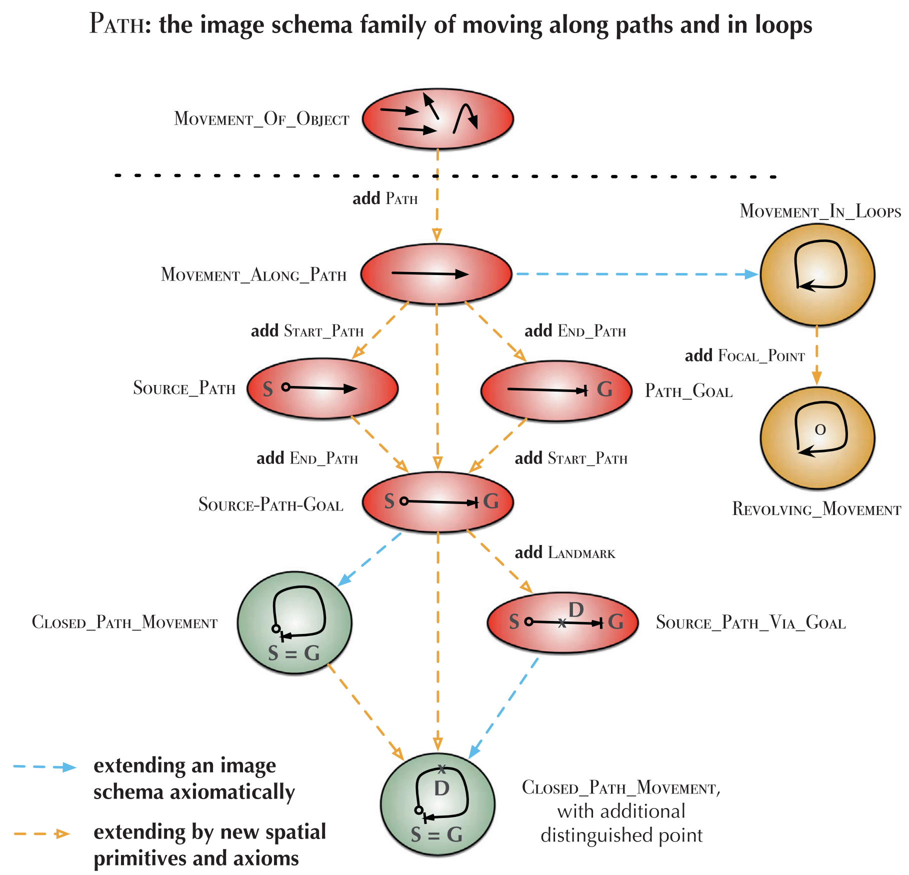

Fig. 11. A portion of the family of image schemas related to path following shown as DOL graph.

the different notions of containment, they are clearly present. Thus, their work provides an empirically well-motivated example of an image schema family.

To implement computationally the idea of using image schemas as generic spaces, two independent algorithmic problems have to be solved. Namely (1) the *Recognition Problem*: to identify an image-schematic theory within an input theory, and (2) the *Generalisation Problem*: to find the most specific image schema common to both inputs.

To address the recognition problem, suppose a lattice $\mathfrak{F}$ encoding an image schema family is fixed. We here assume for simplicity that elements of $\mathfrak{F}$ will be logical theories in a fixed formal logic, say first-order logic.[25] Given an input theory O_1 and $\mathfrak{F}$, solving the recognition problem means finding a member $f \in \mathfrak{F}$ that can be *interpreted* in O_1, i.e. such that we find a renaming σ of the symbols in f (called a signature morphism) and such that $O_1 \models \sigma(f)$ (also written $O_1 \models_\sigma f$).[26] Note that this is a more general statement than claiming the inclusion of the axioms of f (modulo renaming) in O_1 (the trivial inclusion interpretation) since establishing the entailment of the sentences in $\sigma(f)$ from O_1 might indeed be involved.

Computational support for automatic theory-interpretation search in first-order logic is investigated in Normann (2008), and a prototypical system was developed and tested as an add-on to the *Heterogeneous Tool Set*, Hets (Mossakowski, Maeder, & Lüttich, 2007). Experiments carried out in Normann and Kutz (2009) and Kutz and Normann (2009) showed that this works particularly well with more complex axiomatisations in first-order logic, rather than with simple taxonomies expressed in OWL, because in the latter case too little syntactic structure is available to control the combinatorial explosion of the search task. From the point of view of interpreting image schemas into non-trivial axiomatised concepts, we may see this as an encouraging fact, as image schemas are, despite their foundational nature, complex objects to axiomatise.

Once the recognition problem has been solved in principle, the given lattice structure of the image schema family $\mathfrak{F}$ gives us a very simple handle on the generalisation problem. Namely, given two input spaces O_1, O_2, and two image schemas f_1, f_2 from the same family $\mathfrak{F}$ (say, 'containment') such that $O_1 \models_{\sigma_1} f_1$ and $O_2 \models_{\sigma_2} f_2$, compute the most specific generalisation $G \in \mathfrak{F}$ of f_1 and f_2, i.e. their least upper bound in $\mathfrak{F}$. Since the signature of G will be included in

[25] Note that none of the ideas presented here depend on a particular, fixed logic. Indeed, heterogeneous logical specification is central to formal blending approaches, see Kutz, Bateman, et al. (2014).

[26] In more detail: a theory interpretation σ is a signature morphism renaming the symbols of the image schema theory f and induces a corresponding sentence translation map, also written σ, such that the translated sentences of f, written $\sigma(f)$, are logically entailed by O_1.

both signatures of f_1 and f_2, we obtain that $O_1 \models_{\sigma_1} G$ and $O_2 \models_{\sigma_2} G$. $G \in \mathfrak{F}$ is therefore an image schema common to both input spaces and can be used as generic space.

In order to implement this idea, a sufficiently comprehensive library of formalised image schema theories has to be made available for access by a blending engine. The first such library for the case of 'path following' is developed in Hedblom et al. (2015b).

6. Conclusion

In this paper we suggest that image schemas can provide useful heuristics for computational blending of concepts. They can serve as a driving force to identify or define the generic space and its mappings to the input spaces: because image schemas are building blocks of the concepts in the input spaces, we can generate generic spaces by identifying image schemas that are shared across both input spaces. Here, in particular, the idea of organising image schemas into lattice-like structures allows us to identify these shared structures even when the image schemas found in the input concepts do not precisely coincide, but are closely related within a common family.

Our hypothesis is that, compared to syntax-driven approaches (e.g., structure mapping), this approach allows us to identify more cognitively relevant generic spaces and, thus, for cognitively more interesting conceptual blends. Further, we conjecture that many interesting conceptual blends rely on generalisations of image schemas found in the input concepts and organised into well-motivated families.

To test this hypothesis, we intend to continue to develop and expand our image schema library, which formalises image schemas, their families, and interconnections between families, as DOL networks. This also includes more heterogeneous image schema families, where formal languages other than description logics are involved (Containment of Bennett & Cialone, 2014 is an example). While some formalisation of image schemas can be found in the literature on conceptual blending and common sense reasoning (Goguen & Harrell, 2010; Kuhn, 2002; Morgenstern, 2001), the area of systematically formalising and ontologically structuring image schemas is a largely unexplored ground.

The image schema library will allow us to use the tools for computational concept invention that are developed in the COINVENT Project,[27] improve the tool's heuristics, and at the same time test our hypotheses.

We are planning to compare the quality of the blended concepts that are generated based on our proposed heuristics with other blended concepts by using human judges. This way we will evaluate whether our hypotheses are correct and image schemas indeed provide a useful tool for computational concept invention.

[27] http://coinvent-project.eu/.

Acknowledgments

The project COINVENT acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET-Open Grant No.: 611553. We would also like to thank John Bateman, Tony Veale and Mihailo Antović for valuable feedback and discussions about image schemas.

References

Barsalou, Lawrence W. (2008). Grounded cognition. *Annual Review of Psychology, 59*, 617–645.

Bateman, John, Hois, Joana, Ross, Robert, & Tenbrink, Thora (2010). A linguistic ontology of space for natural language processing. *Artificial Intelligence, 174*(14), 1027–1071.

Bazjanac, Vladimir, Forester, James, Haves, Philip, Sucic, Darko, & Xu, Peng. (2002). HVAC component data modeling using industry foundation classes. In *Proceedings of the 5th international conference on system simulation in buildings*, Liege, Belgium.

Bennett, Brandon, & Cialone, Claudia (2014). Corpus guided sense cluster analysis: A methodology for ontology development (with examples from the spatial domain). In Pawel Garbacz & Oliver Kutz (Eds.), *Proceedings of the 8th international conference on formal ontology in information systems (FOIS). Frontiers in artificial intelligence and applications* (Vol. 267, pp. 213–226). Rio de Janeiro, Brazil, September: IOS Press.

Besold, Tarek R., Schorlemmer, Marco, & Smaill, Alan (Eds.). (2015). *Computational creativity research: Towards creative machines. Atlantis thinking machines* (Vol. 7). Atlantis Press.

Bhatt, Mehul, Hois, Joana, & Kutz, Oliver (2012). Ontological modelling of form and function in architectural design. *Applied Ontology, 7*(3), 233–267.

Boden, Margaret A. (2009). Computer models of creativity. *AI Magazine, 30*(3), 23–34.

DiSalle, Robert. (2009). Space and time: Inertial frames, In Edward N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*, winter 2009 edition.

Doerr, Martin (2003). The CIDOC conceptual reference module: An ontological approach to semantic interoperability of metadata. *AI Magazine, 24*(3), 75–92.

Egenhofer, Max J., & Mark, David M. (1995). Naive geography. In Andrew U. Frank & Werner Kuhn (Eds.), *Spatial information theory: A theoretical basis for GIS. Lecture notes in computer science* (Vol. 988, pp. 1–16). Berlin: Springer-Verlag.

Fauconnier, Gilles, & Turner, Mark (1998). Conceptual integration networks. *Cognitive Science, 22*(2), 133–187.

Gallese, Vittorio, & Lakoff, George (2005). The Brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology, 22*(3), 455–479.

Gärdenfors, Peter (2000). *Conceptual spaces – the geometry of thought*. Bradford Books, MIT Press.

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

Gibbs, Raymond W. Jr., (2001). Making good psychology out of blending theory. *Cognitive Linguistics, 11*(3–4), 347–358.

Goguen, Joseph A. (1999). An introduction to algebraic semiotics, with applications to user interface design. In Chrystopher L. Nehaniv (Ed.), *Computation for metaphors, analogy and agents. Lecture notes in computer science* (Vol. 1562, pp. 242–291). Springer.

Goguen, Joseph. A. (2003). Semiotic morphisms, representations and blending for interface design. In *Proceedings of the AMAST workshop on algebraic methods in language processing* (pp. 1–15). Verona, Italy: AMAST Press.

Goguen, Joseph A., & Harrell, D. Fox (2010). Style: A computational and conceptual blending-based approach. In Shlomo Argamon & Shlomo Dubnov (Eds.), *The structure of style: Algorithmic approaches to understanding manner and meaning* (pp. 147–170). Berlin: Springer.

Goguen, Joseph A., & Malcolm, Grant (1996). *Algebraic semantics of imperative programs*. MIT.

Grady, Joseph E. (2001). Cognitive mechanisms of conceptual integration. *Cognitive Linguistics, 11*(3–4), 335–345.

Grüninger, Michael, Hahmann, Torsten, Hashemi, Ali, Ong, Darren, & Ozgovde, Atalay (2012). Modular first-order ontologies via repositories. *Applied Ontology, 7*(2), 169–209.

Guhe, Markus, Pease, Alison, Smaill, Alan, Martínez, Maricarmen, Schmidt, Martin, Gust, Helmar, ... Krumnack, Ulf (2011). A computational account of conceptual blending in basic mathematics. *Cognitive Systems Research, 12*(3–4), 249–265.

Hampe, Beate (2005). Image schemas in cognitive linguistics: Introduction. In Beate Hampe & Joseph E. Grady (Eds.). *From perception to meaning: Image schemas in cognitive linguistics* (Vol. 29, pp. 1–14). Walter de Gruyter.

Hampe, Beate, & Grady, Joseph E. (Eds.). (2005). *From perception to meaning: Image schemas in cognitive linguistics. Cognitive linguistics research* (Vol. 29). Walter de Gruyter.

Hedblom, Maria M., Kutz, Oliver, & Neuhaus, Fabian. (2014). On the cognitive and logical role of image schemas in computational conceptual blending. In Antonio Lieto & Daniele Radicioni (Eds.), *Proceedings of the 2nd international workshop on artificial intelligence and cognition (AIC-2014)*. In *CEUR-WS* (Vol. 1315), Torino, Italy.

Hedblom, Maria M., Kutz, Oliver, & Neuhaus, Fabian. (2015a). Image schemas as families of theories. In Tarek R. Besold, Kai-Uwe Kühnberger, Marco Schorlemmer, & Alan Smaill (Eds.), *Proceedings of the workshop "computational creativity, concept invention, and general intelligence" (C3GI-15)*. In K.-U. Kühnberger, P. König, & S. Walter (Eds.), *Publications of the institute of cognitive science* (Vol. 02-2015), Osnabrück.

Hedblom, Maria M., Kutz, Oliver, & Neuhaus, Fabian (2015b). Choosing the right path: Image schema theory as a foundation for concept invention. *Journal of Artificial General Intelligence, 6*(1), 21–54.

Hois, Joana, Bhatt, Mehul, & Kutz, Oliver (2009). Modular ontologies for architectural design. In *Proceedings of FOMI-09. Frontiers in artificial intelligence and applications* (Vol. 198). Vicenza, Italy: IOS Press.

Hois, Joana, Kutz, Oliver, Mossakowski, Till, & Bateman, John. (2010). Towards ontological blending. In *Proceeing of the 14th international conference on artificial intelligence: Methodology, systems, applications (AIMSA-2010)*, Varna, Bulgaria.

Johanson, Megan, & Papafragou, Anna (2014). What does children's spatial language reveal about spatial concepts? Evidence from the use of containment expressions. *Cognitive Science, 38*(5), 881–910.

Johnson, Mark (1987). *The body in the mind: The bodily basis of meaning, imagination, and reason*. Chicago and London: The University of Chicago Press.

Krieg-Brückner, Bernd, Autexier, Serge, Rink, Martin, & Nokam, Sidoine Ghomsi (2015). Formal modelling for cooking assistance. In Rocco De Nicola & Rolf Hennicker (Eds.), *Software, services, and systems. Lecture notes in computer science* (Vol. 8950, pp. 355–376). Springer International Publishing.

Kuhn, Werner (2002). Modeling the semantics of geographic categories through conceptual integration. In *Proceedings of GIScience 2002* (pp. 108–118). Boulder, USA: Springer.

Kuhn, Werner (2007). An image-schematic account of spatial categories. In Stephan Winter, Matt Duckham, Lars Kulik, & Ben Kuipers (Eds.), *Spatial information theory. Lecture notes in computer science* (Vol. 4736, pp. 152–168). Springer.

Kutz, Oliver & Normann, Immanuel. (2009). Context discovery via theory interpretation. In *Proceedings of the IJCAI workshop on automated reasoning about context and ontology evolution, ARCOE-09*, Pasadena, California.

Kutz, Oliver, Bateman, John, Neuhaus, Fabian, Mossakowski, Till, & Bhatt, Mehul (2014). E pluribus unum: Formalisation, use-cases, and computational support for conceptual blending. In Tarek R. Besold, Marco Schorlemmer, & Alan Smaill (Eds.), *Computational creativity research: Towards creative machines. Thinking machines*. Atlantis/Springer.

Kutz, Oliver, Neuhaus, Fabian, Mossakowski, Till, & Codescu, Mihai. (2014). Blending in the hub—towards a collaborative concept invention platform. In *Proceedings of the 5th international conference on computational creativity*, Ljubljana, Slovenia.

Kutz, Oliver, Mossakowski, Till, Hois, Joana, Bhatt, Mehul, & Bateman, John. (2012.) Ontological blending in DOL. In Tarek R. Besold, Kai-Uwe Kühnberger, Marco Schorlemmer, & Alan Smaill (Eds.), *Procedigngs of the 1st international workshop C3GI@ECAI, computational creativity, concept invention, and general intelligence*. In *Publications of the institute of cognitive science* (Vol. 01), Osnabrück. Montpellier, France.

Kutz, Oliver, Mossakowski, Till, & Lücke, Dominik (2010). Carnap, goguen, and the hyperontologies: Logical pluralism and heterogeneous structuring in ontology design. *Logica Universalis, 4*(2), 255–333, Special Issue on 'Is Logic Universal?'.

Lakoff, George, & Núñez, Rafael E. (2000). *Where mathematics comes from*. Basic Books.

Lenat, Douglas B., Prakash, Mayank, & Shepherd, Mary (1985). CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI Magazine, 6*(4), 65–85.

Mandler, Jean M. (1992). How to build a baby: II. Conceptual primitives. *Psychological Review, 99*(4), 587–604.

Mandler, Jean M. (2004). *The foundations of mind: Origins of conceptual thought*. New York: Oxford University Press.

Mandler, Jean M. (2008). On the birth and growth of concepts. *Philosophical Psychology, 21*(2), 207–230.

Mandler, Jean M., & Pagán Cánovas, Cristóbal (2014). On defining image schemas. *Language and Cognition, 6*(4), 510–532.

Morgenstern, Leora (2001). Mid-sized axiomatizations of commonsense problems: A case study in egg cracking. *Studia Logica, 67*, 333–384.

Mossakowski, Till, Codescu, Mihai, Neuhaus, Fabian, & Kutz, Oliver. (2015). *The road to universal logic–Festschrift for 50th birthday of Jean-Yves Beziau*, chapter The distributed ontology, modelling and specification language – DOL. *Studies in universal logic* (Vol. II), Birkhäuser.

Mossakowski, Till, Maeder, Christian, & Lüttich, Klaus (2007). The heterogeneous tool set. In Orna Grumberg & Michael Huth (Eds.), *TACAS 2007. Lecture notes in computer science* (vol. 4424, pp. 519–522). Heidelberg: Springer-Verlag.

Neuhaus, Fabian, Kutz, Oliver, Codescu, Mihai, & Mossakowski, Till. (2014). Fabricating monsters is hard – towards the automation of conceptual blending. In *Proceedings of computational creativity, concept invention, and general intelligence (C3GI-14)*. In *Publications of the institute of cognitive science, Osnabrück* (Vol. 1-2014, pp. 2–5). Prague.

Normann, Immanuel. (2008). Automated theory interpretation (Ph.D. thesis). Bremen: Department of Computer Science, Jacobs University.

Normann, Immanuel & Kutz, Oliver. (2009). Ontology correspondence via theory interpretation. In *Workshop on matching and meaning (AISB-09)*, Edinburgh, UK.

Oakley, Todd (2007). Image schema. In Dirk Geeraerts & Hubert Cuyckens (Eds.), *The oxford handbook of cognitive linguistics* (pp. 214–235). Oxford University Press.

Randell, David A., Cui, Zhan, & Cohn, Anthony G. (1992). A spatial logic based on regions and connection. In *Proceedings of the 3rd international conference on the principles of knowledge representation and reasoning (KR'92)* (pp. 165–176). Los Altos: Morgan Kaufman.

Reed, Stephen L., Lenat, Douglas B., et al. (2002). Mapping ontologies into cyc. In *AAAI 2002 conference workshop on ontologies for the semantic web* (pp. 1–6).

Reed, Stephen K., & Pease, Adam (2015). A framework for constructing cognition ontologies using WordNet, FrameNet, and SUMO. *Cognitive Systems Research, 33*, 122–144.

Rohrer, Tim (2005). Image schemata in the brain. In Beate Hampe & Joseph E. Grady (Eds.), *From perception to meaning: Image schemas in cognitive linguistics*. *Cognitive linguistics research* (Vol. 29, pp. 165–196). Walter de Gruyter.

Runco, Mark A., & Jaeger, Garrett J. (2012). The standard definition of creativity. *Creativity Research Journal, 24*(1), 92–96.

Schmidt, Martin, Krumnack, Ulf, Gust, Helmar, & Kühnberger, Kai-Uwe (2014). *Computational approaches to analogical reasoning: Current trends*. *Studies in computational intelligence* (Vol. 548). Berlin, Heidelberg: Springer.

Schorlemmer, Marco, Smaill, Alan, Kühnberger, Kai-Uwe, Kutz, Oliver, Colton, Simon, Cambouropoulos, Emilios, & Pease, Alison. (2014). COINVENT: Towards a computational concept invention theory. In *Proceedings of the 5th international conference on computational creativity*, Ljubljana, Slovenia.

Smith, Barry, Ceusters, Werner, Klagges, Bert, Köhler, Jacob, Kumar, Anand, Lomax, Jane, ... Rosse, Cornelius (2005). Relations in biomedical ontologies. *Genome Biology, 6*(5), R46.

St. Amant, Robert, Morrison, Clayton T., Chang, Yu-Han, Mu, Wei, Cohen, Paul R., & Beal, Carole. (2006). An image schema language. In *Proceedings of the 7th international conference on cognitive modeling (ICCM)*, Trieste, Italy.

Turner, Mark (2014). *The origin of ideas: Blending, creativity, and the human spark*. Oxford University Press.

Veale, Tony (2003). The analogical thesaurus. In John Riedl & Randy Hill (Eds.), *Proceedings of the 15th innovative applications of artificial intelligence conference* (pp. 137–142). Acapulco, Mexico: AAAI Press.

Veale, Tony. (2012). From conceptual mash-ups to bad-ass blends: A robust computational model of conceptual blending. In Mary Lou Maher, Kristian Hammond, Alison Pease, Rafael Prez y Prez, Dan Ventura, & Geraint Wiggins (Eds.), *Proceedings of the 3rd international conference on computational creativity* (pp. 1–8), Dublin, Ireland.

Yang, Fan-Pei Gloria, Bradley, Kailyn, Huq, Madiha, Wu, Dai-Lin, & Krawczyk, Daniel C. (2012). Contextual effects on conceptual blending in metaphors: An event-related potential study. *Journal of Neurolinguistics, 26*, 312–326.