

# Choosing the Right Path: Image Schema Theory as a Foundation for Concept Invention\*

**Maria M. Hedblom**

HEDBLOM@IWS.CS.UNI-MAGDEBURG.DE

*Department for Knowledge and Language Engineering (IWS)  
Otto-von-Guericke University of Magdeburg, Germany*

**Oliver Kutz**

OLIVER.KUTZ@UNIBZ.IT

*Research Centre for Knowledge and Data (KRDB)  
Free University of Bozen-Bolzano, Italy*

**Fabian Neuhaus**

FNEUHAUS@IWS.CS.UNI-MAGDEBURG.DE

*Department for Knowledge and Language Engineering (IWS)  
Otto-von-Guericke University of Magdeburg, Germany*

**Editor:** Tarek R. Besold, Kai-Uwe Kühnberger, Tony Veale

## Abstract

Image schemas are recognised as a fundamental ingredient in human cognition and creative thought. They have been studied extensively in areas such as cognitive linguistics. With the goal of exploring their potential role in computational creative systems, we here study the viability of the idea to formalise image schemas as a set of interlinked theories. We discuss in particular a selection of image schemas related to the notion of ‘path’, and show how they can be mapped to a formalised family of microtheories reflecting the different aspects of path following. Finally, we illustrate the potential of this approach in the area of concept invention, namely by providing several examples illustrating in detail in what way formalised image schema families support the computational modelling of conceptual blending.

**Keywords:** image schemas, grounded cognition, computational creativity, concept invention, conceptual blending

## 1. Introduction

The cognitive processes underlying concept invention are still largely unexplored ground, although promising theories have been developed within the last few decades. Two of the most influential directions are the theory of grounded cognition (Barsalou, 2008) and the embodied mind theory (Lakoff and Johnson, 1999; Gallese and Lakoff, 2005). Both propose that human cognition is grounded in our bodily experience with the environment. In the case of grounded cognition, mental representations are thought to be derived from the embodied experience that is used to structure concepts, including even the most abstract ones. Embodied cognition takes this one step further by arguing against mental representations and by claiming that concepts are (or may be identified with) the neural activation of embodied experiences.

Following the reasoning of how embodied experiences shape our cognition, a more specific theory of concept formation and language understanding was introduced under the name of image schemas. It is a theory that focuses on basic spatial cognition. The theory was originally developed

---

\*. This paper is a revised and expanded version of (Hedblom, Kutz, and Neuhaus, 2015).

in the late 80s by Lakoff (1987) and Johnson (1987), and was quickly taken up by other researchers in the area (e.g. Mandler (1992)).

According to Johnson (1987) “*an image schema is a recurring dynamic pattern of our perceptual interaction and motor programs that gives coherence and structure to our experience.*”

The ‘image schema’ is thought to be the abstracted spatial pattern from repeated sensorimotor experience. These mental structures offer a foundation and a way to ground other cognitive phenomena, such as language capacity, understanding, and reasoning. They offer a connection between the bodily experienced relationships of physical objects in time and space with the internal conceptual world of an agent. In language, they can be seen as the conceptual building blocks for metaphoric and abstract thought. Some of the most commonly mentioned examples of image schemas are: CONTAINMENT, SUPPORT and MOVEMENT\_ALONG\_PATH<sup>1</sup>.

Another important approach towards an understanding of concept invention is the theory of *conceptual blending*, introduced by Fauconnier and Turner (1998), which developed further the idea of *bisociation* introduced by the psychologist Koestler (1964). Conceptual blending proposes that novel concepts arise from a selective combination of previously known information (see Section 5.1 for a more thorough introduction).

The primary goal of our research is to develop a computational system for concept invention by combining a formal representation of image schemas with the framework of conceptual blending. One of the major obstacles for implementing such a system is that image schemas are, typically, not crisply defined in the literature, but rather presented as mouldable concepts. Their adaptability is indeed part of the explanatory success of image schemas. Hence, to realise our goal of a computational system using image schemas in conceptual blending, we need to develop a formal representation of image schemas that captures their inherent complexity.

In this paper we suggest that image schemas should be considered as members of tightly connected image schema families, where the connecting relation is based on the notion of family resemblance. In particular, each of the image schemas covers a particular conceptual-cognitive scenario within the scope of the schema family. An image schema family may be formally represented as a set (i.e. a *family*) of interlinked theories.

To illustrate our approach, we will use the image schema of MOVEMENT\_ALONG\_PATH, analyse its use in natural language, and sketch a representation of MOVEMENT\_ALONG\_PATH in first-order logic, which focusses on the different branching points of microtheories involved in the family. Afterwards, we show how the image schemas may be used for conceptual blending.

The remainder of the paper is structured as follows: In Section 2, a more detailed account of image schemas is presented by discussing their internal structure and their role in language. In Section 3, the PATH-following is discussed in more detail. In Section 4, formal approaches to image schemas are discussed and we introduce our idea of how to gather image schemas in families of theories by using PATH-following as a proof of concept. This is done by representing PATH-following as a DOL-graph and as a FOL-axiomatisation. In Section 5, we investigate the role image schemas can play in concept invention within the framework of conceptual blending theory. Finally, in Sections 6 and 7, a discussion and a short conclusion are provided.

---

1. This image schema is also referred to as SOURCE\_PATH\_GOAL schema. For reasons that will become more obvious in Section 4.2, we use MOVEMENT\_ALONG\_PATH as the most generic term for this image schema.

## 2. Image schemas

### 2.1 Image schemas and embodiment

The theory of image schemas stems from the theory of embodied/grounded cognition. It is a theory that emphasises the role of bodily experiences as a source for cognitive capacities. It has become increasingly supported by findings in cognitive linguistics and neuroscience (e.g. Tettamanti et al. (2005); Gallese and Lakoff (2005); Feldman and Narayanan (2004); Wilson and Gibbs (2007); Louwerse and Jeuniaux (2010)).

The theory offers an interesting view of cognition for approaches to artificial intelligence as it provides a more direct route to computational cognition than traditional, more hard-coded approaches. I.e., artificial agents are encouraged to learn, and artificial cognitive structures are populated by their ‘experiences’, similar to the learning process observed in human children (see (Chrisley, 2003) for an introduction to embodied artificial intelligence and (Vernon, 2014) for an overview of such cognitive architectures).

Building on grounded cognition, image schemas are thought to be the mental representations extracted from bodily experiences, and more specifically, experiences that can be described using basic spatial relationships. Image schemas are therefore mental abstractions of learnt spatial relationships (e.g. CONTAINMENT or SUPPORT).

In the early stages of cognitive development, these image schemas are formed and ‘fine-tuned’ as the experiences with a particular spatial relationship are increased and extended to different situations (Mandler, 2008; Rohrer, 2005). Due to this fine-tuning, it appears prominent that image schemas consist of different ‘parts’ (Mandler, 2004). These ‘parts’ can either be removed or added while still capturing the same basic image schema, generating what can be described as an image schema ‘family’. Mandler and Pagán Cánovas (2014) refer to these ‘parts’ as *spatial primitives*; the fundamental spatial building blocks. As this view is essential to our approach for formalising image schemas we will regularly return to the notion of spatial primitives in later sections.

The explanatory value of image schemas lies in their function as abstracted spatial relationships. It is believed that they contain vital information for the understanding of concepts and their conceptual neighbourhoods. For example, to properly understand what a ‘cup’ is, an infant needs to learn that cup is kind of CONTAINER, and, thus, other objects may stand in CONTAINER-typical spatial relationships to a cup. (E.g., an object can be in the cup.) Respectively, the notion of a ‘table’ needs to be connected to a SUPPORT image schema since otherwise the child will not understand that objects remain on tables after being placed on them. This way, image schemas map affordances to objects (in the sense of Affordance Theory (Gibson, 1977)) and can be used to explain increasingly more complicated concepts. This is done through information transfer and can be observed in natural language, for example in conceptual metaphors.

### 2.2 Information transfer and conceptual metaphor

Image schemas are a abstracted spatial relationships that are associated with affordances. Their cognitive benefit is that they may map affordances to objects, which an agent has not encountered yet. For example, if the image schema of SUPPORT has been learnt through perceptual exposure of ‘plates on tables’, an infant can infer that table-like objects such as ‘desks’ also have the SUPPORT image schema and can SUPPORT objects such as ‘books’ as well. As the environment

becomes increasingly complex for the infant, this information transfer becomes a fundamental part of cognition.

Information transfer in language is often done through conceptual metaphors. Information is moved from one known source domain to an unknown target domain. Conceptual metaphors can be further specialised to image-schematic metaphors. These are the metaphors that do not transfer general conceptual knowledge from one domain to another but the skeletal structure of the image schema (Kövecses, 2010). Here the source domain in the analogical transfer is stripped down to the image schema skeleton which is mapped to the target domain and there fleshed out with local domain information. In sections 3 and 5 we will discuss several examples of image-schematic metaphors (and similes), which are the result of conceptual blending.

As language develops and the individual is exposed to increasingly more abstract concepts than those found in early infancy, image schemas can be used to ground the novel concepts in already comprehended concepts. Embodied experiences and image schemas are often used in natural language to explain abstract concepts. For example, in a social hierarchy people can be either ‘above’ or ‘below’ us, expressions learned from embodied experiences of the image schema ABOVE, which itself derives from experiencing the human body’s vertical axis. In natural language, metaphors such as ‘falling from grace’ or ‘the rise to power’ use the same image schema to represent status and success.

One of the most well-known examples of an attempt at grounding abstract concepts in image schemas is the work of Lakoff and Núñez (2000). In *Where Mathematics Comes From*, they defend the view that image schemas lay the foundation for abstract concepts in mathematics. They explain how the notions of addition and subtraction can be traced from back and forward MOVEMENT\_ALONG\_PATH, and extend the reasoning to more abstract constructions such as complex numbers. While being influential work, it has also received heavy criticism, in particular targeted at vague terminology and methodology (e.g. Goldin (2001); Schiralli and Sinclair (2003)) and mathematical errors (e.g. Voorhees (2004)).

The image schema CONTAINMENT is commonly described as the sum of the interrelationships of an inside, an outside and a boundary (Lakoff, 1987). An abstract example of the image schema CONTAINMENT is the conceptual metaphor “*to be in love*”. Obviously, there is no spatial region for the emotional state of love in the same sense as there is for a physical container such as a cup. Yet, we use the spatial language to talk about the phenomenon of love: e.g. we ‘fall in love’ or ‘fall out of love’.

There is a clear connection between CONTAINMENT and prepositions such as ‘in’, ‘into’ and ‘out of’. Bennett and Cialone (2014) investigated the CONTAINMENT relationship by searching text corpora for words similar to containment, e.g. ‘surrounding’ and ‘enclosing’. The authors’ method distinguished eight different kinds of CONTAINMENT.

Prepositions in combination with verbs often do appear to be the key words that help identify image schemas in language (Johanson and Papafragou, 2014). Below we will discuss natural language and conceptual metaphors of the PATH-following image schema family.

### 2.3 Image schemas and their structure

In this section we, look more closely at how image schemas are thought to be structured. The first pertinent distinction is that image schemas can be both static and dynamic. For example, CONTAINMENT can either describe the situation in which the cup already contains coffee, or

alternatively the situation in which the coffee is poured from a source: a kettle, to a goal: a cup, defined as an IN and OUT schema. The static image schemas can in turn be differentiated under three different categories: *orientational* (e.g. ABOVE), *topological* (e.g. CONTACT), and *force-dynamic* (e.g. SUPPORT) (Lakoff and Núñez, 2000).

A more dynamic way to understand the IN and OUT schema is to view them as combinations of the two image schemas CONTAINMENT and MOVEMENT\_ALONG\_PATH, building on the idea that image schemas can be combined with one another to generate more specific and complex image schemas (Kuhn, 2007; Walton and Worboys, 2009; Mandler and Pagán Cánovas, 2014; Oakley, 2010). Another example is how MOVEMENT\_ALONG\_PATH easily can be connected with the image schema LINK resulting in the higher level image schema LINKED\_PATH: The image schema concept that encompasses linked behaviour on two, or more, joint paths. This “Gestalt” grouping of image schemas means that there must be a distinction between the most perceptually primitive image schemas and the more complex image schemas.

In language, this corresponds to the observation that combinations of image schemas are suitable to describe more complex concepts. To illustrate this, Kuhn (2007) suggested that ‘transportation’ can be understood as a combination of the image schemas of SUPPORT and MOVEMENT\_ALONG\_PATH, and Mandler (2004) suggested that ‘marriage’ can be viewed as a LINKED\_PATH.

One proposal to hierarchically structure the range from simple to more complex and dynamic image schemas is the approach presented in (Mandler and Pagán Cánovas, 2014), which builds on empirical data from studies on cognitive development. In their work, the umbrella term ‘image schemas’ is divided into three different levels: *spatial primitives*<sup>2</sup> (the conceptual building blocks build from spatial information), *image schemas*<sup>3</sup> (simple spatial stories), and *conceptual integrations* (image schemas combined with a non-spatial element such as force or emotion).<sup>4</sup>

One major advantage of partitioning image schemas into spatial primitives, spatial schemas, and conceptual integration is that it enables a more fine-grained analysis of connections between image schemas. We believe that the change from one spatial schema to another can be accomplished by providing or detailing more spatial information, e.g., by adding additional spatial primitives. In Section 4.2, we will demonstrate this idea by using PATH-following as a proof of concept. In the next section PATH-following will be introduced in more detail.

### 3. The image schema family PATH-following

#### 3.1 Introducing PATH-following

In this section, we will explain how image schemas, like MOVEMENT\_ALONG\_PATH, are members of image schema families. For this purpose we introduce the PATH-following image schema family

---

2. The notion of spatial primitives is not novel. Research on such semantic building blocks can be found in the linguistic literature. E.g. the work on spatial semantics by Veale and Keane (1992) and the more general work on semantic primes that covers more than the spatial and temporal aspects found in spatial primitives (Wierzbicka, 1996).

3. When referring to this concept we will use the term *spatial schemas* to avoid ambiguity between the wider notion of ‘image schema’ and its narrower sense introduced by Mandler and Pagán Cánovas (2014).

4. For the purposes of this paper, only spatial primitives and spatial schemas will be further discussed. In principle, our approach is general enough to allow for heterogeneity, also on the logical level. Therefore one may also include conceptual integrations involving non-spatial elements in our image schema families, cf. the discussion in Sections 4.2 and 4.3.

and illustrate how the family is organised hierarchically from general to more specific versions of PATH-following by the addition of spatial information and primitives.

MOVEMENT\_ALONG\_PATH is one of the first image schemas to be acquired in early infancy as children are immediately exposed to movement from a range of objects. This, in combination with the neurological priority to process moving objects over static objects, suggests that the image schema is either innate or learnt at a very early stage in cognitive development (Rohrer, 2005). However, in order to understand how the PATH-following family is fine-tuned and in ‘more completion’ internally structured, experiments with children have provided some insights on distinguishing how the different spatial schemas may develop.

Firstly, already at an early age children pay more attention to moving objects than resting objects. Trivial as it may seem, it requires children to detect the spatial primitive OBJECT (or THING) and the spatial schema MOVEMENT\_OF\_OBJECT.<sup>5</sup> Secondly, children tend to remember the PATH of the movement of the object. The PATH is a spatial primitive, which is different from the movement and the moving object.<sup>6</sup>

In addition to these two basic spatial primitives and as the child becomes more and more familiar with PATH-following, image schemas that contain more spatial information are learned. This means that in more advanced stages, image schemas may include beyond MOVEMENT\_OF\_OBJECT and the spatial PATH itself also the spatial primitive END\_PATH, and later also a START\_PATH (Mandler and Pagán Cánovas, 2014). Already at five months infants can distinguish PATH-following that has an END\_PATH (the image schema PATH\_GOAL) from the initial PATH, while the START\_PATH is less interesting until the end of the first year of life. This is further supported by linguistic analyses in which an END\_PATH is initially more interesting than a START\_PATH (Johanson and Papafragou, 2014).

Table 1 summarises the spatial primitives that may be involved in image schemas of the PATH-following family.<sup>7</sup>

<b>Spatial primitive</b>	<b>Description</b>
OBJECT	an object
PATH	the path the object moves along
START_PATH	the initial location
END_PATH	the final location

Table 1: Spatial primitives of the PATH-following family

A more specified example of the PATH-following family is presented by Lakoff and Núñez (2000). In accordance with other linguistic literature on image schemas they are focussed on the SOURCE\_PATH\_GOAL schema, see Figure 1. Here, the object, called *trajector*, moves from a source to a goal. END\_PATH and START\_PATH are not identical to the SOURCE and GOAL found in the

5. OBJECT is understood here in a very wide sense that includes not only solid material objects but entities like waves on a pond or shadows. Mandler and Pagán Cánovas (2014) also discuss MOVE as a spatial primitive of its own. We consider MOVEMENT\_OF\_OBJECT to be a spatial schema, since movement necessarily involves a temporal dimension and, further, it always involves at least one spatial primitive, since any movement, necessarily, involves at least one OBJECT that moves.

6. This spatial primitive is not to be confused with the image schema family PATH-following.

7. Table 1 is based on Mandler and Pagán Cánovas (2014), but includes some changes. In particular, as mentioned above, we do not consider MOVE to be a spatial primitive.

SOURCE\_PATH\_GOAL schema. In SOURCE\_PATH\_GOAL, a direction and a purpose are implied in the image schema, which changes the conceptual nature of the movement. Lakoff and Núñez (2000) make the distinction of ‘elements’, or roles, that to some extent correspond to the spatial primitives discussed above, but additional distinctions are added. The elements<sup>8</sup> are listed in Table 2. Most importantly, they make the clear distinction between end location and goal, as they distinguish between ‘path’, the actual trajectory of a movement, and ‘route’, the expected movement.

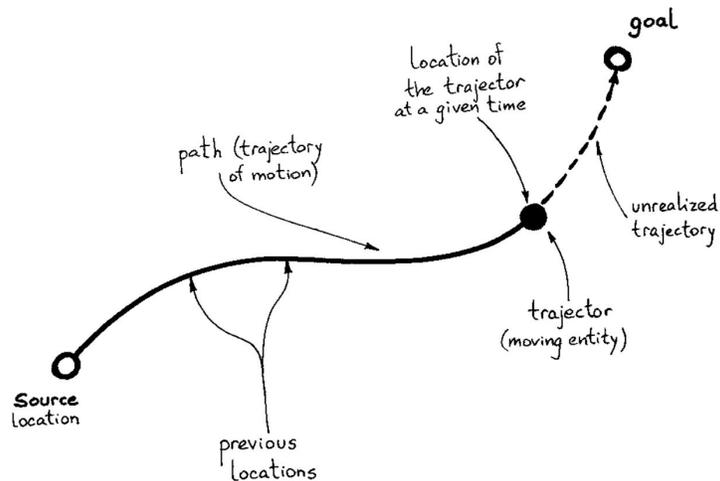


Figure 1: The SOURCE\_PATH\_GOAL schema as illustrated by Lakoff and Núñez (2000).

The distinction, made by Lakoff and Núñez (2000), between the *expected* movement and the *actual* movement is primarily interesting for a description of how new image schemas relate to actual events and how new image schemas are learned. Consider, for example, a situation where a child observes the movement of a billiard ball and is surprised that the ball stops because it is blocked by another billiard ball. In this case, a given instance of the MOVEMENT\_ALONG\_PATH spatial schema formed the expectations of the child, which were disappointed by the actual physical movement, because the expected END\_PATH (the goal) does not correspond to the actual END\_PATH (end location). Given a repeated exposure to similar events, the child may develop the new spatial schema BLOCKAGE. After learning BLOCKAGE, the child will no longer be surprised by blocked movement since the expected END\_PATH (the goal) will correspond to the actual END\_PATH (end location). While the terminological distinction between *expected trajectory* and *actual trajectory* is useful, these do not necessarily need to constitute two different spatial primitives. Indeed, spatial primitives are parts of image schemas and, thus, always parts of conceptualisations, and not parts of actual events.

While the notions of path-following of Mandler and Pagán Cánovas (2014) and Lakoff and Núñez (2000) coincide widely, there are differences in terminology and definitions. In this paper we follow primarily the former.

8. We altered the terms to better match the terminology in (Mandler and Pagán Cánovas, 2014), but no change in content was made.

Element	Description
trajector	The object
source	The initial location
goal	The intended end location
route	A pre-realised route from source to goal
path	The trajectory of motion
position	The position of the trajector at a given time
direction	The direction of the trajector at a given time
end location	End location, may not correspond to the goal location

Table 2: Elements of PATH according to Lakoff and Núñez (2000)

### 3.2 Concepts that involve PATH-following

As briefly demonstrated with CONTAINMENT and ABOVE, image schemas can be used as a source for grounding abstract concepts in already comprehended concrete concepts. In this section we consider examples for concepts, which involve members of the PATH-following family.

The most straightforward examples of concepts that involve PATH-following are concepts that are about the spatial relationship of movement between different points. Prepositions such as *from*, *to*, *across* and *through* all indicate a kind of PATH-following<sup>9</sup>. This also includes key verbs that describe movement, e.g. *coming* and *going*. Another example, here for the spatial schema SOURCE\_PATH\_GOAL, is *Going from Berlin to Prague*. Note that in many cases we do not provide information about START\_PATH and END\_PATH of a movement; e.g. *leaving Berlin* and *travelling to Berlin* are examples for the spatial schemas SOURCE\_PATH and PATH\_GOAL, respectively. *Meandering* is an example for a concept that realises MOVEMENT\_ALONG\_PATH, which involves a PATH but no START\_PATH or END\_PATH. In contrast, no discernable PATH is involved in *roaming the city*, which is an example for MOVEMENT\_OF\_OBJECT. – These examples illustrate that spatial schemas may be ordered hierarchically with respect to their content: SOURCE\_PATH\_GOAL contains more spatial primitives and more information than, for example, MOVEMENT\_ALONG\_PATH, which is the root of the PATH-following family. And MOVEMENT\_ALONG\_PATH is more specific than MOVEMENT\_OF\_OBJECT.

Beyond concepts that involve movement, PATH-following plays an important role in many abstract concepts and conceptual metaphors.

The concept of “*going for a joy ride*” realises the spatial schema SOURCE\_PATH, since it has a START\_PATH and a PATH but no END\_PATH. Similarly, the expression “*running for president*” describes the process of trying to get elected as president metaphorically as a PATH\_GOAL. In this metaphor the PATH consists of the various stages of the process (e.g. announcing a candidacy and being nominated by a party) with the inauguration as END\_PATH.

Another metaphor “*life is a journey*”, studied by Ahrens and Say (1999), makes an analogical mapping between the passing of time in life, to the passing of spatial regions on a journey. As in the example mentioned above, where the concept of “being in love” acquired information from the CONTAINMENT schema, this metaphor gains information from the spatial primitives connected to the image schema SOURCE\_PATH\_GOAL. Here, the most important spatial primitives are

9. Some prepositions include other image schemas at the same time. E.g. ‘through’ involves apart from PATH also some notion of CONTAINMENT.

START\_PATH and END\_PATH – in this metaphor they are mapped to the moments of birth and death, as well as the PATH itself, illustrating how “life goes on” in a successive motion without branching.

A different perspective on life and death is expressed in the metaphorical expression “*the circle of life*”. Implied is that life leads to death, but also that death gives rise to life, completing a cyclic movement – the image schema MOVEMENT\_IN\_LOOPS. This image schema can be considered as a version of PATH-following, in which START\_PATH and END\_PATH coincide at the same ‘location’.

These examples illustrate a general pattern, namely that many conceptual metaphors involving PATHS are about processes, and different events during such processes are treated metaphorically as locations on a path. This leads to a conceptualisation of the abstract concept of time, which we will further investigate in the next section.

	<b>Expression</b>	<b>Level in hierarchy</b>
Concrete:	Roaming the city	MOVEMENT_OF_OBJECT
	Meandering	MOVEMENT_ALONG_PATH
	Leaving Berlin	SOURCE_PATH
	Travelling to Berlin	PATH_GOAL
	Going from Prague to Berlin	SOURCE_PATH_GOAL
Abstract:	Going for a joy ride	SOURCE_PATH
	Running for president	PATH_GOAL
	Life is a journey	SOURCE_PATH_GOAL
	The circle of life	MOVEMENT_IN_LOOPS

Table 3: Summary of the mentioned expressions and their level in the PATH-following hierarchy

### 3.3 Time and processes as PATH

The conceptualisation of time has been investigated by Boroditsky (2000). Here we follow suit by looking at how members of the PATH-following image schema family are widely used as a conceptual metaphors for time. We consider several examples and discuss the role of PATH-following image schemas for the conceptualisation of processes in general.

One popular way to conceptualise time is as MOVEMENT\_ALONG\_PATH. Often, time is conceptualised as having a beginning, a START\_PATH; this may be the Big Bang or the moment of creation in a religious context. Depending on the cosmological preferences, time may also be conceptualised to have an end, an END\_PATH: the Big Rip or an apocalypse.

Other religious traditions embrace the notion of a ‘Wheel of Time’, that is time as a cyclic repetition of different aeons. The underlying image schema involves a MOVEMENT\_IN\_LOOPS. The same image schema is used in the conceptualisation of time within calendars: the seasons are a continuous cycle where any winter is followed by a new spring. Similarly, the hours of the day are represented on analogue clocks as 12 marks on a cycle, and the passing of time is visualised as MOVEMENT\_IN\_LOOPS of the handles of the clock.

The conceptualisation of time, in itself, is an interesting example for the usage of image schemas. However, the real significance is that these image schemas can be seen as providing the conceptual skeletal structure for our understanding of processes. Assume we want to understand a complex process, e.g. the demographic development of a country, the acceleration of a falling object,

or the economic situation of a country. In these situations we often use two-dimensional coordinate systems where the vertical axis represents the property in question (e.g. population, speed, GDP, respectively) and the horizontal axis represents time. These coordinate systems are so useful and so widely applicable because we can conceptualise arbitrary processes as MOVEMENT\_ALONG\_PATH, where the paths represent some important dimension or aspect of the process.

The importance of PATH-following image schemas for the conceptualisation of processes can be illustrated by considering *similes*. If we pick from Table 4 randomly a target domain  $X$  from the first column and a source domain  $Y$  from the second column, the resulting simile  $X$  is like  $Y$  will be sensible. (Of course, depending on the choice of  $X$  and  $Y$  the simile may be more or less witty.) Note that the target domains have little or nothing in common. Thus, at least on first glance, one would not expect that one can compare them meaningfully to one and the same source domain.

<b>Target Domain</b>	<b>Source Domain</b>
Watching the football game	the swinging of a pendulum
Their marriage	a marathon
The story	escaping a maze
This piece of music	a sail boat during a hurricane
Bob's career	a roller coaster ride
Her thoughts	a Prussian military parade
Democracy in Italy	stroll in the park

Table 4: PATH similes: <target> is like <source>.

The similes work because all of the concepts in the second column involve physical MOVEMENT\_ALONG\_PATH, which have some pertinent characteristics. These characteristics may concern the shape of the path itself (e.g. the path of a roller coaster involves many ups and downs and tight curves, the path out of a maze involves many turns, the path of a pendulum is regular and between two points), the way the movement is performed (e.g. the movement of a sail boat during a storm is erratic and involuntary, a stroll in the park is done leisurely), and the effects the movement may have (e.g. running a marathon is exhausting, a Prussian military parade may be perceived as threatening). In each of the similes we use some of the pertinent characteristics from the source domain to describe the process from the target domain. For example, in the simile 'Bob's career is like a Prussian military parade' we conceptualise the career as a path along time (with career-related events like promotions as the sites on the path) and transfer characteristics from the movement of a Prussian military parade on this path. Thus, one way to read the simile is that Bob moves through the stages of his career in an exceptionally predictable fashion. The example illustrates how the similes work: first, we conceptualise the process in the target domain as MOVEMENT\_ALONG\_PATH, where the events of the process are ordered by time, and then we transfer some pertinent characteristics of the MOVEMENT\_ALONG\_PATH of the source domain to the target domain. This pattern is not just applicable to the concepts in Table 4. As we discussed above, any process can be conceptualised as MOVEMENT\_ALONG\_PATH, thus, any process could be added as target domain in Table 4. Further, any concept that involves interesting physical movement along some path could be added as source domain. Hence, the use of the image schema MOVEMENT\_ALONG\_PATH enables the mechanical generation of similes for processes.

Similes are a particular form of concept generation in which two domains are combined. This phenomenon is strongly connected to conceptual blending that we will discuss further in Section 5.1.

To summarise, in this section we have introduced the image schema `MOVEMENT_ALONG_PATH`. We have seen that it is widely used in natural language and plays an important role in our understanding of time and processes. The examples show that the notion of `PATH`-following, at its core, is about movement along some trajectory. However, there are important differences both with respect to the spatial primitives that are involved and with respect to the shape of the `PATHS`. In the next section, we consider how image schemas can be represented in formal languages. One particular concern is to represent image schemas in a way that adequately captures the variety and flexibility of image schemas.

## 4. Formalising image schemas as graphs of theories

### 4.1 Previous work on formalising image schemas

Image schema research has had great impact in the cognitive sciences and in particular in cognitive linguistics. However, within computational cognitive systems, and artificial intelligence in general, it has not yet been explored to its full potential.

Looking at how image schemas can be computationally acquired, there are studies that attempt to model early cognitive development and learn from perceptual input. The connectionist model proposed by Regier (1996) learns to linguistically classify visual stimuli in accordance with the spatial terms of various natural languages. Similarly, the *Dev E-R* system by Aguilar and Pérez y Pérez (2015) is a computer model that simulates the first sensorimotor stages in cognitive development. Their system learns to distinguish and fine-tune visual clues such as nuances of colour, as well as different sizes of objects and directions of movement. Both approaches demonstrate how an artificial agent can develop cognitive abilities and language development from perceptual input.

Another study using perceptual input to simulate the development of image schemas was made by Nayak and Mukerjee (2012). They fed video material of `OBJECTS` moving `IN` and `OUT` of boxes into an unsupervised statistical model in order to capture the dynamic aspects of the `CONTAINMENT` schema. From this, the system learned how to categorise different `CONTAINMENT` contexts and could in combination with a linguistic corpus generate simple `CONTAINMENT`-related language constructions.

These are examples of systems that learn image schemas and visual relationships from perceptual input. More commonly, work on formalising image schemas is done when the image schemas are already identified. Prominent work in this field is the work by Kuhn (2002, 2007). He argues that image schemas capture abstractions in order to model affordances. Working top-down rather than bottom-up as above, he uses WordNet to define noun words and connects them to spatial categorisations related to image schemas based on affordance-related aspects of meaning.

Walton and Worboys (2009) build further on Kuhn's work by visualising and formalising the connections between different image schemas using bigraphs. By visually representing the topological and 'physical' image schemas relevant in built environments, they demonstrate how more complex dynamic image schemas such as `BLOCKAGE` could be generated using sequences of bigraph reaction rules on top of simpler static image schemas.

St. Amant et al. (2006) present what they call the *Image Schema Language*, ISL. In their paper, they provide a set of diagrams that illustrate how combinations of image schemas can lead to more complex image schemas, and provide some real life examples.

Brugman and Lakoff (1988) discuss how image schema transformations form networks that capture the relationships in polysemous words, in particular the preposition ‘over’ is investigated. This relates to our own approach of how to formalise and formally represent image schemas. Namely to use the hierarchical structure of image schemas demonstrated previously to represent image schemas as families of theories.

## 4.2 Image Schema Families as Graphs of Theories

In the previous sections, we argued for image schemas to be members of families, which are partially ordered by generality. In the following section, we will describe and visualise an approach to represent the connections between image schemas, belonging to the same family. In order to discuss the problem of how more complex image schemas can be constructed through a combination of different image schemas (e.g. LINKED\_PATH, MOVEMENT\_IN\_LOOPS), we will discuss the possible interconnection these families of theories allow. Formally, we can represent the idea as a graph<sup>10</sup> of theories in DOL, the *Distributed Ontology, Modeling and Specification Language* (Mossakowski et al., 2015).

This choice is motivated primarily by two general features of DOL: (1) the heterogeneous approach, which allows for a variety of image schematic formalisations without being limited to a single logic, and (2) the focus on linking and modularity. Therefore, DOL provides a rich toolkit to further formally develop the idea of *image schema families* in a variety of directions.

In more detail, DOL aims at providing a unified metalanguage for handling the diversity of ontology, modelling, and specification languages, for which it uses the umbrella term ‘OMS’. In particular, DOL includes syntactic constructs for:

1. “as-is” use of OMS formulated (as a logical theory) in a specific ontology, modelling or specification language,
2. defining new OMS by modifying and combining existing OMS (which are possibly written in different languages), and
3. mappings between OMS, resulting in networks of OMS.

DOL is equipped with an abstract model-theoretic semantics.<sup>11</sup> The theoretical underpinnings of the DOL language have been described in detail in (Kutz, Mossakowski, and Lücke, 2010) and (Mossakowski, Lange, and Kutz, 2012), whilst a full description of the language can be found in (Mossakowski et al., 2015) or (in a more condensed form) in (Mossakowski et al., 2013).

Building on similar ideas to those underlying the first-order ontology repository COLORE<sup>12</sup> (Grüninger et al., 2012), we propose to capture image schemas as interrelated families of (heterogeneous) theories. Similar ideas for structuring common sense notions have also been applied to various notions of time (Van Benthem, 1983; Allen and Hayes, 1985). This general approach also covers the introduction of non-spatial elements such as ‘force’ as a basic ingredient of image schemas, as for instance argued for by Gärdenfors (2007) and constitute the core of Mandler and Pagán Cánovas (2014)’s *conceptual integrations* mentioned above.

10. These graphs are diagrams in the sense of category theory.

11. The final DOL specification was submitted as a standard to the Object Management Group (OMG) in late 2015

12. See <http://stl.mie.utoronto.ca/colore/>

**PATH: the image schema family of moving along paths and in loops**

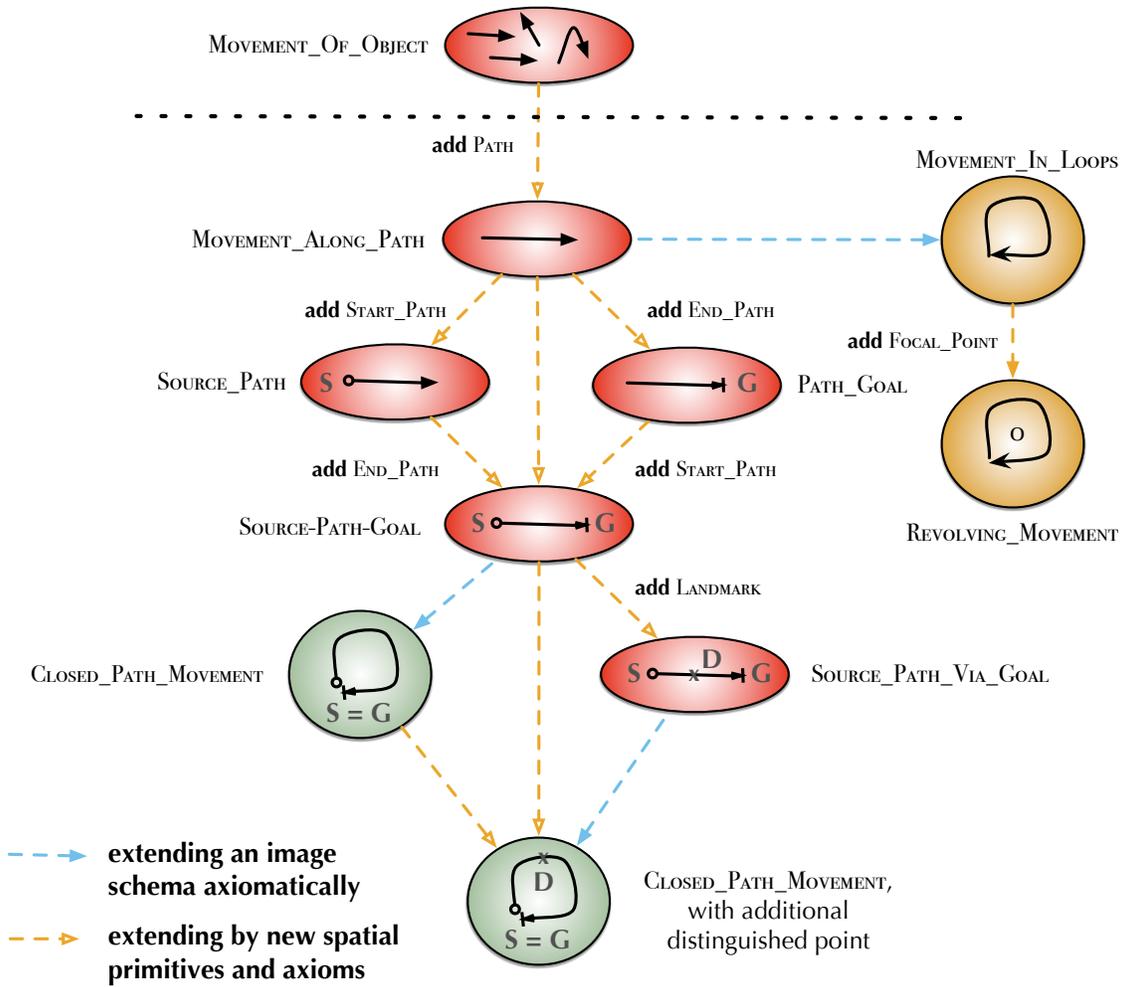


Figure 2: A portion of the family of image schemas related to path following shown as DOL graph.

In Figure 2, some of the first basic stages of the image schema family PATH-following are presented. Ranging from Mandler’s general definition presented above, of object movement in any trajectory, to more complex constructions.

The particular image schema family sketched 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). 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.

One way MOVEMENT\_ALONG\_PATH can be specialised is as the image schema of MOVEMENT\_IN\_LOOPS. 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 spatial information of a *focal point*, which the path revolves around – this leads to the notion of *orbiting*, or, by continuously moving the orbiting path away from the focal point, to create the concept of *spirals*. Alternatively, we may change MOVEMENT\_ALONG\_PATH by adding distinguished points; e.g. the START\_PATH, the target END\_PATH, or both.

The MOVEMENT\_IN\_LOOPS image schema may be further specialised by identifying (the location of) the START\_PATH and the END\_PATH. 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 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 (i.e. ‘landmarks’) or other related spatial primitives.

We will now show how the theories of image schemas and the various branching points in the graph can be characterised formally.

### 4.3 Axiomatisation of Path-Following

In this section, we present an axiomatisation of the image schemas represented in Figure 2. The focus of our axiomatisations is to capture the important differences of the branching points of the PATH-following family, not an exhaustive axiomatisation. For the sake of brevity, we will present only selected axioms in this section. A more complete axiomatisation is available at an Ontohub repository.<sup>13</sup>

Our axiomatisation approach is inspired by semantics in the neo-Davidsonian tradition (Davidson, 1967; Parson, 1990). We consider image schemas as a type of event (in generality quite similar to the view defended in (Clausner and Croft, 1999) to view image schemas as a kind of ‘domain’) and consider spatial primitives as thematic roles of these events. Thus, if a given image schema is enriched by adding a new spatial primitive, this is typically represented by adding a new entity (e.g. site) and a new relation (e.g. has\_start\_path) that determines the thematic role of the new entity in the event. As representation language we use ISO/IEC 24707 Common Logic. Common Logic is a standardised language for first-order logic knowledge representation, which supports some limited form of higher-order quantification and sequence variables (Menzel, 2011).

For the axiomatisation of the image schemas in the PATH-following family we assume an image schema MOVEMENT\_ALONG\_PATH as the root of the family. MOVEMENT\_ALONG\_PATH is

13. <https://ontohub.org/repositories/imageschemafamily/>

derived from a more general notion, namely MOVEMENT\_OF\_OBJECT. This is movement of some kind that involves only one spatial primitive, namely an OBJECT. This object plays the role of the *trajector* within the context of the MOVE. This can be formalised in Common Logic as follows:

```
(forall (m)
  (iff
    (MovementOfObject m)
    (exists (o)
      (and
        (Movement m)
        (Object o)
        (has_trajector m o))))))
```

No additional information about what kind of object is moving and how it is moving is assumed.<sup>14</sup>

The schema MOVEMENT\_ALONG\_PATH is the result of adding a new spatial primitive to MOVEMENT\_OF\_OBJECT, which plays the role of a PATH.

```
(forall (m)
  (iff
    (MovementAlongPath m)
    (exists (p)
      (and
        (MovementOfObject m)
        (Path p)
        (has_path m p))))))
```

Under a PATH we understand a collection of two or more sites, which are connected by successor relationships. Each of these sites have (relative to the path) at most one successor site. The transitive closure of the successor relation defines a *before* relationship (relative to the path); and for any two different sites  $x, y$  of a given path, either  $x$  is before  $y$  or  $y$  is before  $x$  (relative to the path).<sup>15</sup> This axiomatisation provides a representation of a quite abstract notion of MOVEMENT\_ALONG\_PATH. It needs to be sufficiently abstract, since it serves as the root node for the PATH-following family. All other image schemas in the family are derived from this root by adding additional spatial primitives and/or additional axioms.

Given this notion of PATH, we can axiomatise the relationship between the PATH and the OBJECT, which characterises a MOVEMENT\_ALONG\_PATH. During the movement, the moving object needs to pass through all sites of the path in a temporal order, which matches the before-relationship between the sites:

```
(forall (p o m s1 s2)
  (if
    (and
      (MovementAlongPath m)
      (has_path m p)
      (has_trajector m o)
      (before s1 s2 p))
    (exists (t1 t2)
      (and
```

---

14. From an ontological perspective, MOVEMENT\_OF\_OBJECT can be seen as a kind of process (or occurrent). Thus, any adequate axiomatisation of MOVEMENT\_OF\_OBJECT needs to represent change over time in some form. To keep things simple, we here just quantify over time points. We assume that time points are ordered by an *earlier* relationship. Further, we use two other relationships to connect time points to processes:  $(has\_start\ m\ t)$  means *The movement m starts at time point t* and  $(during\ t\ m)$  means *Time point t lies within the interval during which movement m happens*.

15. The before-relationship is not a total order, since antisymmetry is not postulated.

```
(Timepoint t1) (Timepoint t2)
(during t1 m) (during t2 m)
(located_at o s1 t1) (located_at o s2 t2)
(earlier t1 t2))))))
```

The image schema `SOURCE_PATH` is the result of adding the spatial primitive `START_PATH` to `MOVEMENT_ALONG_PATH`. We represent this with the *has\_starts\_path* relationship. The `START_PATH` of a `PATH` is a site on the path that is before any other site of the path:

```
(forall (m)
  (iff
    (SourcePathMovement m)
    (exists (s)
      (and
        (MovementAlongPath m)
        (has_start_path m s))))))

(forall (m s1 s2 p)
  (if
    (and
      (SourcePathMovement m)
      (Site s1)
      (Site s2)
      (not (= s1 s2))
      (has_path m p)
      (has_start_path m s1)
      (part_of s2 p))
    (before s1 s2 p)))
```

What distinguishes `SOURCE_PATH` from other movements is the following: at the start of a `SOURCE_PATH` movement the object that moves is located at the `START_PATH`:

```
(forall (m s t p o)
  (if
    (and
      (SourcePathMovement m)
      (has_start m t)
      (has_trajector m o)
      (has_start_path m s)
      (located_at o s t))))
```

Analogously, we can define `PATH_GOAL` as a `MOVEMENT_ALONG_PATH` with an `END_PATH`. A `SOURCE_PATH_GOAL` is a movement, which includes both landmarks of `START_PATH` and `END_PATH`. Thus, `SOURCE_PATH_GOAL` can be defined as the union (of the axioms) of `SOURCE_PATH` and `PATH_GOAL`.<sup>16</sup>

`CLOSED_PATH_MOVEMENT` is a special case of `SOURCE_PATH_GOAL`, where the location of the `START_PATH` and the `END_PATH` of the `PATH` coincide.

```
(forall (m s g)
  (if
    (and
      (has_start_path m s)
      (has_end_path m g))
    (iff
      (ClosedPathMovement m)
      (and
        (SourcePathGoalMovement m)
        (= (location_of s) (location_of g))))))
```

`SOURCE_PATH_VIA_GOAL` is a different way to refine `SOURCE_PATH_GOAL`. In this case an additional designated site is added, which lies between the `START_PATH` and the `END_PATH` of the `PATH`.

---

16. In DOL, this is done by using the keyword ‘and’, which amounts to taking the model-theoretic intersection of the model classes of the theories of `SOURCE_PATH` and `PATH_GOAL`.

```
( (forall (m)
  (iff
    (SourcePathViaGoalMovement m)
    (exists (s p)
      (and
        (SourcePathGoalMovement m)
        (has_path m p)
        (Site s)
        (part_of s p)
        (not (has_start_path m s))
        (not (has_end_path m s)))))))
```

Both CLOSED\_PATH\_MOVEMENT and SOURCE\_PATH\_VIA\_GOAL can be combined in the obvious way.

A completely different branch of the movement image schema family does not involve either START\_PATH or END\_PATH, but the PATH consists of a loop of sites. One way to represent this is by requiring that the before-relationship is reflexive (with respect to the path of the movement):

```
(forall (m)
  (iff
    (MovementInLoops m)
    (and
      (MovementAlongPath m)
      (forall (p s)
        (if
          (and
            (has_path m p)
            (Site s)
            (part_of s p))
          (before s s p))))))
```

The difference between MOVEMENT\_IN\_LOOPS and CLOSED\_PATH\_MOVEMENT is that in the latter case both START\_PATH and END\_PATH are present, they just spatially coincide. Hence, the movement is over when the object meets the target. In contrast, MOVEMENT\_IN\_LOOPS entails that the moving object is located at the same location more than once.

REVOLVING\_MOVEMENT is a subtype of MOVEMENT\_IN\_LOOPS. To define it, we need to consider two additional factors: the shape of the path is elliptical, and there is a focal point, which the movement revolves around. The focal point itself is a site, but it is typically the location of an object. A detailed axiomatisation of this image schema is beyond the scope of this paper, we just provide an initial sketch:

```
(forall (m)
  (iff
    (RevolvingMovement m)
    (and
      (MovementInLoops m)
      (exists (p s)
        (and
          (has_path m p)
          (Elliptical (shape p))
          (Site s)
          (has_focal_point p s))))))
```

## 5. Image Schemas in Computational Conceptual Blending

In this section, we will illustrate how formalised families of image schemas, as just sketched above, can help in the computational modelling of concept invention, an area at the heart of AGI. More

precisely, we will here focus on the highly influential framework of conceptual blending (Fauconnier and Turner, 2003; Turner, 2007; Lakoff and Núñez, 2000), and illustrate the foundational role that a formal theory of image schemas plays in its computational realisation.

### 5.1 A crash course on conceptual blending

Introduced by Fauconnier and Turner (1998), conceptual blending has been employed very successfully to understand the process of concept invention, studied particularly within cognitive psychology and linguistics. The theory argues that at the heart of novel concept creation lies a combination process involving already existing knowledge and understood concepts. By merging two, or more, conceptual spaces, a blended conceptual space results. This blend contains information from both input spaces and has emergent properties due to its own unique composition. The classic example is the blend of a ‘houseboat’, containing merged information from the input spaces ‘house’ and ‘boat’.<sup>17</sup>

One of the central aspects of blending is the way in which ‘common structure’ between the input concepts is understood to steer the creation of the new concept. The ‘merging’ of the input spaces is moderated by this common structure, represented as the generic space, or as it is called in formal approaches, the base ontology (see Figure 3).<sup>18</sup> The common structure of the input spaces is understood to play a vital role in rendering the newly constructed concept meaningful, as it ensures that the blended space also contains the structure found in the generic space.

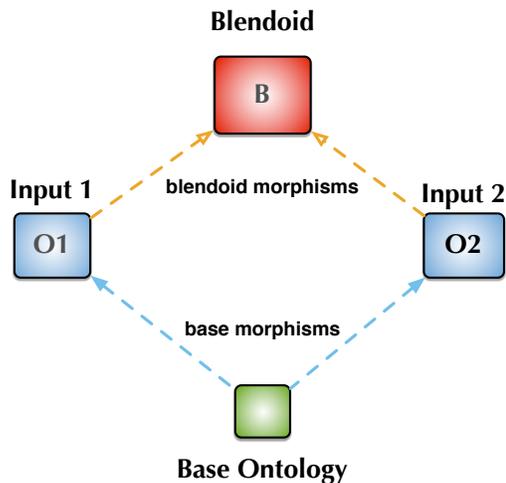


Figure 3: The basic integration network for blending: concepts in the base ontology are first refined to concepts in the input ontologies and then selectively blended into the blendoid.

However, despite this influential research, within computational creativity and AI in general, relatively little effort has been devoted to fully formalise these ideas and to make them amenable to computational techniques, but see (Schorlemmer et al., 2014; Kutz et al., 2014a) for overviews.

Unlike other combination techniques, *blending* aims at creatively generating (new) concepts on the basis of input theories whose domains are thematically distinct but whose specifications share some structural similarity.

17. This and the related blend of ‘boathouse’ were fully formalised in (Kutz et al., 2014b).

18. In the limit case, the shared structure might be trivial, and a concept such as ‘red pencil’ might be understood as a blend too, by simply imposing properties from one input space onto another.

Kutz et al. (2014a) describe in detail the basic formalisation of conceptual blending, as sketched by the late Joseph Goguen and discuss some of its variations (Goguen and Harrell, 2010). Moreover, it is illustrated how the Distributed Ontology Language DOL can be used to declaratively specify blending diagrams of various shapes, and how the workflow and creative act of generating and evaluating a new, blended concept can be managed and computationally supported within Ontohub, a DOL-enabled theory repository with support for a large number of languages and formal linking constructs, see (Kutz et al., 2014b; Mossakowski, Kutz, and Codescu, 2014). The reasoning engine managing heterogeneous theories and computationally supporting the Ontohub repository is the *Heterogeneous Tool Set* Hets (Mossakowski, Maeder, and Lüttich, 2007). The graph for the structured theory illustrating the PATH-following family and automatically being generated by the Hets system from its formal specification is shown in Figure 4.

Figure 3 illustrates the basic, classical case of an ontological blending diagram. The lower part of the diagram shows the generic space (tertium), i.e. the common generalisation of the two input spaces, which is connected to these via total (theory) morphisms, the base morphisms. The newly invented concept is at the top of this diagram, and is computed from the base diagram via a colimit. More precisely, any consistent subset of the colimit of the base diagram may be seen as a newly invented concept, a *blendoid*.<sup>19</sup> Note that, in general, ontological blending can deal with more than one base and two input ontologies, and in particular, the sets of input and base nodes need not exhaust the nodes participating in a base diagram.

## 5.2 Using image schemas in computational blending

One problem for conceptual blending, and related work on analogy engines (e.g. structure mapping (Gentner, 1983; Forbus, Falkenhainer, and Gentner, 1989) and heuristic-driven theory projection (HDPT) (Schmidt et al., 2014)) is the generation of a ‘sensible’ blend. In a completely automatised system, there is currently no simple way to distinguish the blendoids that a human would consider meaningful from those that lack cognitive value. This problem grows exponentially in relation to the size of the input spaces. The larger the input spaces, the more combinations can be generated resulting in a multitude of possible blendoids, most of which will make little sense if evaluated by humans. In real life scenarios, the amount of information in the input spaces can be vast, complicating things for successful concept invention tremendously when looked at as a formal, combinatorial problem.

A proposal to explain the ease with which humans perform blending is given via the ideas of *packing and unpacking*, as well as *compression and expanding* of conceptual spaces, as outlined by Turner (2014). These terms aim to capture how we mentally carry around ideas as compressed ‘idea packages’ that we can ‘unpack’ and utilise in different contexts on the fly. The process of packing and unpacking ideas is important for the contextualised usage of conceptual blends in various situations. Generally, the idea of *optimality principles* in blending theory is meant to account for an evaluation of the quality and appropriateness of the resulting blends (Fauconnier and Turner, 2003). However, there is currently no general formal proposal how such optimality principles could be implemented computationally, apart from some work on turning such principles into metrics for rather lightweight formal languages (Pereira and Cardoso, 2003).

---

19. A technically more precise definition of this notion is given in Kutz et al. (2012). Note also that our usage of the term ‘blendoid’ does not coincide with the (non-primary) blendoids defined in Goguen and Harrell (2010).

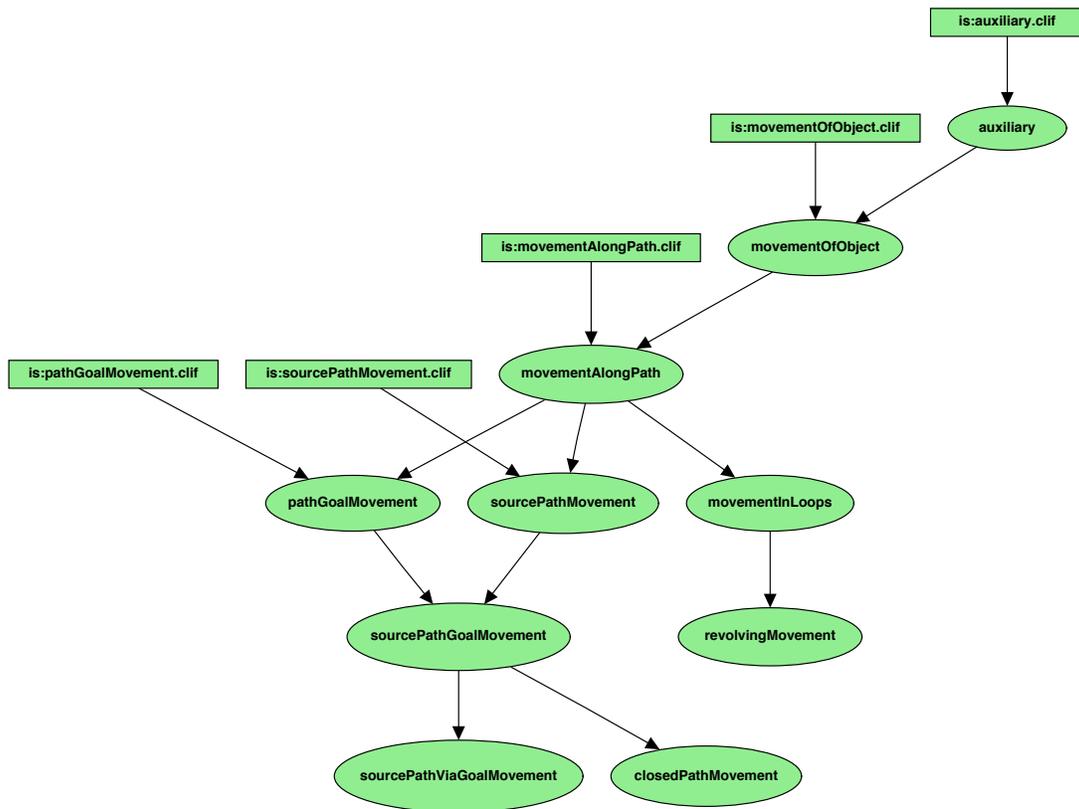


Figure 4: The PATH-following family displayed as a structured DOL theory in Ontohub/Hets.

Hedblom, Kutz, and Neuhaus (2014) suggested that instead of relying on purely syntactic approaches, image schemas in their role as conceptual building blocks could be used to guide the computational blending process. The principle idea here is that employing image schemas in the construction of generic spaces will not only result in a significant reduction of the number of generated blends, but will moreover filter out many of those blends human evaluators would deem meaningless. A related and complementary approach is (Veale, Feyaerts, and Forceville, 2013), where the problem of constraining the search space was addressed by suggesting that blending is performed in a task-specific context. Here, selecting a task-specific context in a blending scenario means to simultaneously work forward from the input spaces and backward from the *desired* elements of the blend space.

In this line of thinking, one way to use image schemas in blending is to identify them as the prime ingredient for the construction of a generic space. When performing the search for common structure in the different input spaces, the search could be guided by mapping (parts of) the content of the input spaces to nodes in a library of formally represented image schemas. As image schemas hold semantic value in the form of spatial relationships, the blendoids would be based on the same content. In theory, this is similar to classic structure mapping that preserves relationships, but as image schemas model e.g. affordances (Kuhn, 2007), a blendoid will inherit such information as well.

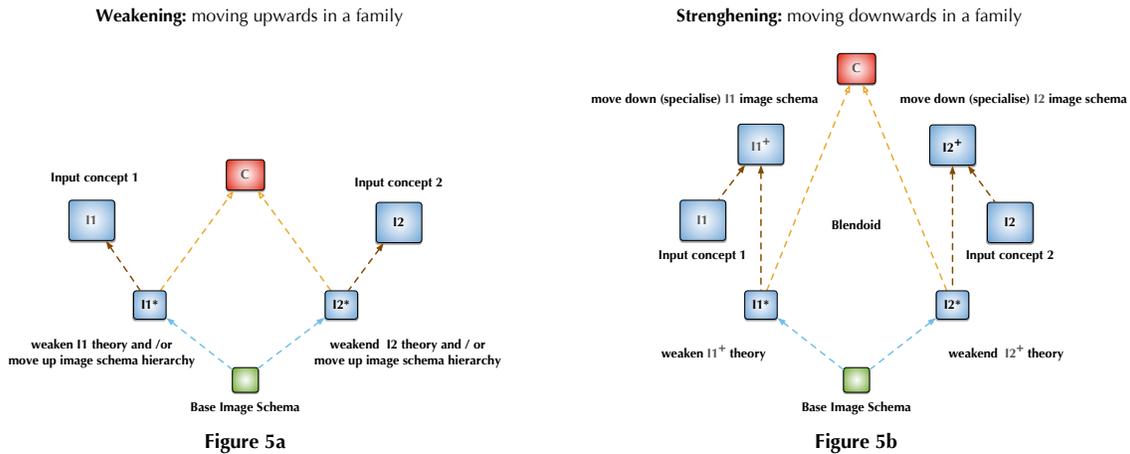


Figure 5: Blending using common image schemas: strengthening vs. weakening.

Figure 5 shows the two basic ways of using image schemas within the conceptual blending workflow. In both cases, the image schematic content takes priority over other information the input concepts might contain. On the left, following the core model of blending described above, we first identify different spatial structure within the same image schema family in the input concepts, and then generalise to the most specific common version within the image schema family to identify a generic space, using our pre-determined graph of spatial schemas (i.e. we compute the least upper bound in the lattice). The second case, shown on the right, illustrates the situation where we first want to specialise or complete the (description of the) spatial schemas found in the input concepts, before performing a generalisation step and to identify the generic space. This means moving down in the graph of the image schema family. Of course, also a mix of these two basic approaches is reasonable, i.e. where one input spatial schema is specialised within a family whilst the other is generalised in order to identify a generic space based on image-schematic content. Examples for both cases are described below in Sections 5.3 and 5.4.

### 5.3 The PATH-following family at work

To study how image schematic content can be used more concretely within conceptual blending, we will now look at a number of examples. In this section, we will illustrate how moving up and down within the image schema family of path-following opens up a space of blending possibilities, infused with the respective semantics of the (versions of) the image schema. In the next section, we will then discuss in more formal detail how these ideas work on a logic-based level.

As outlined in Section 3.3, processes in general can be easily combined with a variety of more specific PATH-following schemas. More specifically, we can explore the basic idea how to combine the input space of ‘thinking process’, which involves only an underspecified kind of ‘movement of thoughts’, with a second input space that carries a clearly defined path-following image schema. This leads intuitively to a number of more or less well known phrases that can be analysed as blends, including: ‘train of thought’, ‘line of reasoning’, ‘derailment’, ‘flow of arguments’, or ‘stream of consciousness’, amongst others. Indeed, a central point we want to make in this section is that these

blends work well and appear natural because of the effectiveness of the following heuristics:<sup>20</sup> (i) given two input spaces  $I_1$  and  $I_2$ , search for the strongest version  $G$  of some image schema that is *common to both*, according to the organisation of a particular image schema family  $\mathfrak{F}$ ; (ii) use  $G$  as generic space; and (iii) use again  $\mathfrak{F}$  to identify the stronger version of  $G$ , say  $G'$ , inherent in one of the two inputs, and use the semantic content of  $G'$  to steer the overall selection of axioms for the blended concept.

To illustrate this process informally, let us briefly consider the concepts of ‘stream of consciousness’, ‘train of thought’, and ‘line of reasoning’<sup>21</sup>.

On first inspection, the spatial schema of movement related to ‘thinking’ might be identified as MOVEMENT\_OF\_OBJECT, i.e. without necessarily identifying following a PATH at all. Indeed, in Figure 2, MOVEMENT\_OF\_OBJECT is marked as an ‘entry point’ to the path-following family.

The *stream of consciousness* may be seen as an unguided flow of thoughts, in which topics merge into each other without any defined steps, but rather in a continuous manner. It lacks a clear START\_PATH and has no guided movement towards a particular END\_PATH. It resembles the more basic forms of PATH-following that, according to Mandler and Pagán Cánovas (2014), is simply movement in any trajectory.

A *train of thought*<sup>22</sup> can be conceptualised in various ways. It differs from a stream of consciousness by having a more clear direction, often with an intended END\_PATH. It is possible to say that one “lost their train of thought”, or that “it was hijacked” or how “it reversed its course”. The ‘train’ may be understood as a chain-like spatial object (in which case ‘losing the train’ decodes to ‘disconnecting the chain’) or more plainly as a locomotive. In the Pixar film ‘Inside Out’ (2015), the ‘Train of Thought’ is an actual train that travels the mind of the fictional character Riley Anderson, and delivers daydreams, facts, opinions, and memories.

A *line of reasoning* might be seen as a strengthening of this blend, where the path imposed is linear. Although a ‘line’, mathematically speaking, has no beginning or end, the way this expression is normally understood is as a discrete succession of arguments (following logical rules) leading to an insight (or truth). This blend might therefore be analysed to correspond to SOURCE\_PATH\_GOAL in (Lakoff and Núñez, 2000), in which there is a clear direction and trajectory of the ‘thought’ (trajector).

In order to understand how blending can result in these concepts, and how image schemas are involved, let us have a closer look at the input spaces and their relationship to the PATH-following image schemas. Relevant input spaces include line (perhaps analysed as ‘discrete interval’), stream/river, train/locomotive, and, as secondary input space, ‘thinking process’.

‘Thinking’ as an input space is difficult to visualise. However, when ‘thinking’ is understood as a process it can be easily combined with various PATH-following notions (see Section 3.2 above). As thoughts (in the form of OBJECT) are moved around, the simplest form of thinking

20. By ‘heuristics’ we mean a method that imposes rules on how to select a base (i.e. introduces a preference order on possible generic spaces) and, moreover, rules to decide which axioms to push into the blend. I.e., without any heuristics we are left to perform a randomised axiom selection, followed by an evaluation of the resulting blended concept.

21. The examples presented here are chosen to illustrate the basic ideas how to employ families of image schemas in blending. It is not intended to capture fully the meaning of these terms as they are used in the psychological or linguistic literature, or indeed the subtle meaning they might carry in natural language.

22. The expression ‘train of thoughts’ appears to have been first used by Thomas Hobbes in his Leviathan (1651): “By ‘consequence of thoughts’ or ‘TRAIN of thoughts’ I mean the occurrence of thoughts, one at a time, in a sequence; we call this ‘mental discourse’, to distinguish it from discourse in words.”

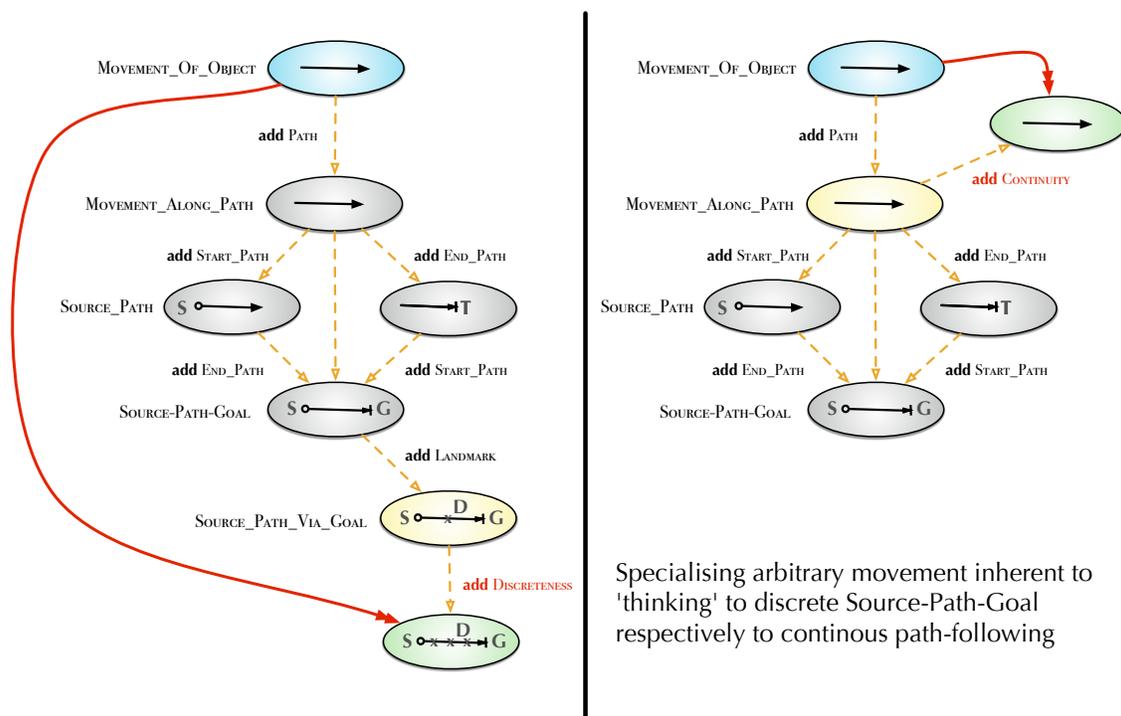
is MOVEMENT\_OF\_OBJECT. There is no START\_PATH nor an END\_PATH. Intuitively, it does not appear to have any particular PATH (in the sense of a spatial primitive).

A stream is characterised by a continuous flow along a PATH. Whilst a START\_PATH and END\_PATH can be part of a stream-like concept, like in the fleshed out concept of a river with a source and mouth, they do not constitute an essential part of the concept of stream.

For a train (understood as ‘locomotive’), the concepts of a START\_PATH and END\_PATH has a much higher significance. The affordances found in trains are primarily those concerning going from one place to another. A train ride can also be seen as a discrete movement in the sense that for most train rides, there are more stops than the final destination. This results in a discrete form of the spatial schema SOURCE\_PATH\_GOAL.

When blending such forms of movement with the thinking process, what happens is that the unspecified form of movement found in ‘thinking process’ is specialised to the PATH-following characteristics found in the second input space. The result is the conceptual metaphors for the different modes of thinking listed above, where the generic space contains just MOVEMENT\_OF\_OBJECT, and the blended concepts inherit the more complex PATH-following from ‘train’, ‘stream’, or ‘line’.

PATH: specialisation (and generalisation) of image schemas in the path family



Specialising arbitrary movement inherent to 'thinking' to discrete Source-Path-Goal respectively to continuous path-following

Figure 6: How ‘thinking’ transforms into ‘train of thought’ respectively ‘stream of consciousness’.

In more detail, Figure 6 shows two specialisations of the basic spatial schema of MOVEMENT\_OF\_OBJECT. The first, shown on the left, specialises to a discrete version of the schema SOURCE\_PATH\_GOAL with a designated element and discrete movement, supporting the ‘train of thought’ blend. The second, shown on the right, specialises to a continuous version of MOVEMENT\_ALONG\_PATH, where an axiom for gapless movement is added to the

```

ontology story =
%% A story is defined as a telling of a plot.
(forall (x)
  (iff
    (Story x)
    (exists (y) (and (Plot y) (tells x y)))))

%% A plot consists of some events
(forall (x)
  (if (Plot x)
    (exists (y) (and (Event y) (part_of y x)))))

%% Every event in a plot is causally connected to at least some other event in the plot.
(forall (x y)
  (if (and (Plot x) (Event y) (part_of y x))
    (exists (z) (and (Event z) (part_of z x) (or (causes y z x) (causes z y x))))))

%% Some stories have a protagonist
(exists (x y) (and (Story x) (has_protagonist x y) (Character)))
end

```

Figure 7: Ontology story in DOL and Common Logic.

MOVEMENT\_ALONG\_PATH spatial schema to support the ‘flow of consciousness’ blend. As a third possibility, in ‘line of reasoning’, we would impose additionally a linear (and perhaps discrete) path onto ‘thinking’.

#### 5.4 Blending in depth

In Section 3.3, we discussed how PATH spatial schemas support similes, where some process (e.g. a story) is compared to some other concept (e.g. a roller coaster ride). In the same way as the elements in the PATH schema family may be used for creating similes, they can be used for conceptual blending. In this case, the PATH spatial schemas play the role of the generic space and may also be used to strengthen the input spaces. In this subsection, we discuss the process in detail and show some of the relevant axioms. The example illustrates the blending pattern from Figure 5b.

The input spaces for our blending process are *Story* and *Roller Coaster Ride*. They are formally represented in Figures 7 and 8 as axiomatisations in Common Logic with a DOL wrapper.<sup>23</sup> Both axiomatisations are quite weak. Stories are defined as a telling of a plot, and they may involve a protagonist. A plot consists of some causally connected events. A roller coaster ride is a scary amusement ride that follows either some steel or wooden track. The track is fast-paced and consists at least of a start, a thrill element, and an end. Note that the track is an instantiation of a PATH image schema, more specifically a SOURCE\_PATH\_GOAL.<sup>24</sup>

Because SOURCE\_PATH\_GOAL is embedded in the roller coaster concept, it is natural to use it as the base space in our blend. (We reuse its axiomatisation from Section 4.3.) However, SOURCE\_PATH\_GOAL is not present in the space *Story*, thus we need to strengthen the concept by adding the image schematic content from SOURCE\_PATH\_GOAL to *Story*. This can be defined in DOL with the help of a signature map:

23. The first and the last line of each axiomatisation are DOL expressions. Their only purpose is to label these ontologies.

24. To save space we omitted many of the axioms that realise the SOURCE\_PATH\_GOAL schema.

## CHOOSING THE RIGHT PATH

```

ontology linearStory = { sourcePathGoalMovement with
  SourcePathGoalMovement |-> LinearStory ,
  Path |-> Plot ,
  Site |-> Event ,
  has_trajector |-> has_protagonist ,
  Object |-> Character ,
  has_path |-> has_plot ,
  has_start_path |-> starts_with ,
  has_end_path |-> ends_with ,
  successor_of |-> causes } and story

```

The resulting concept is a *Linear Story*, that is a story, where the protagonist participates in a linear succession of events with a clear start and ending.

While strengthening adds new information to an input space, weakening removes some information. This may be necessary, because the blend may otherwise be logically inconsistent. But even if logical consistency is not an issue, one of the input spaces may contain information that is not desirable or irrelevant for a specific blend. For example, the axiomatisation in Figure 8 provides information about the material of roller coasters, which may be removed completely. Further, the first axiom defines that roller coasters are scary amusement rides. This could be replaced by an axiom that keeps the information that roller coasters are scary, but omits the connection to amusement rides. All of these changes can be expressed in DOL with the help of filtering operations and extensions.

In this example it is not necessary to weaken the *Story* input space. On the other hand, since the SOURCE\_PATH\_GOAL image schema is already realised in the input space *Roller Coaster*, there is no strengthening necessary. Thus, in the case of this example  $I1^+$  and  $I1^*$  in Figure 5b are identical; so are  $I2^+$  and  $I2^*$ .

```

ontology rollerCoasterRide =
% A roller coaster ride is a kind of scary amusement ride.
(forall (x)
  (iff (RollerCoasterRide x)
    (and (AmusementRide x) (Scary x))))

% A roller coaster ride follows some track & involves some Person as participant.
(forall (x)
  (iff (RollerCoasterRide x)
    (exists (y z)
      (and (RollerCoasterTrack y) (follows x y) (Person z) (has_participant x z)))))

% A roller coaster track is fast-paced. It starts with the start, has at
% least one thrill element as part & ends with the end.
(forall (x)
  (if (RollerCoasterTrack x)
    (exists (y1 y2 y3)
      (and (fast_paced x) (starts_with x y1) (Start y1) (part_of y2 x)
        (ThrillElement y2) (ends_with x y3) (End y3)))))

% Roller coaster tracks are either made from wood or steel.
(forall (x)
  (if (RollerCoasterTrack x)
    (or (material_of x steel) (material_of x wood))))
[...]
```

Figure 8: Ontology rollerCoasterRide in DOL and Common Logic.

The blended concept is the result of merging both of the weakened *Linear Story*  $I1^*$  and weakened *Roller Coaster*  $I2^*$ . To achieve this in DOL we need to define the interpretations from the base image schema *SOURCE\_PATH\_GOAL* to  $I1^*$  and  $I2^*$ . The following is the definition for weakened *Roller Coaster*.

```

interpretation base2rollerCoaster: base to rollerCoasterRideWeakened =
  Path |-> RollerCoasterTrack ,
  Site |-> RollerCoasterElement ,
  SourcePathGoalMovement |-> RollerCoasterRide ,
  has_trajector |-> has_participant ,
  Object |-> Person ,
  has_path |-> follows ,
  has_start_path |-> starts_with ,
  has_end_path |-> ends_with
  
```

Analogously we define another interpretation `base2story`.<sup>25</sup>

By combining the two interpretations `base2story` and `base2rollerCoaster` we get a new concept: a *Thriller*. In DOL the blended concept can be defined as follows:<sup>26</sup>

```

ontology blend = combine base2rollerCoaster, base2plot with Story |-> Thriller
  
```

Figure 9 provides an overview over the whole blending diagram.

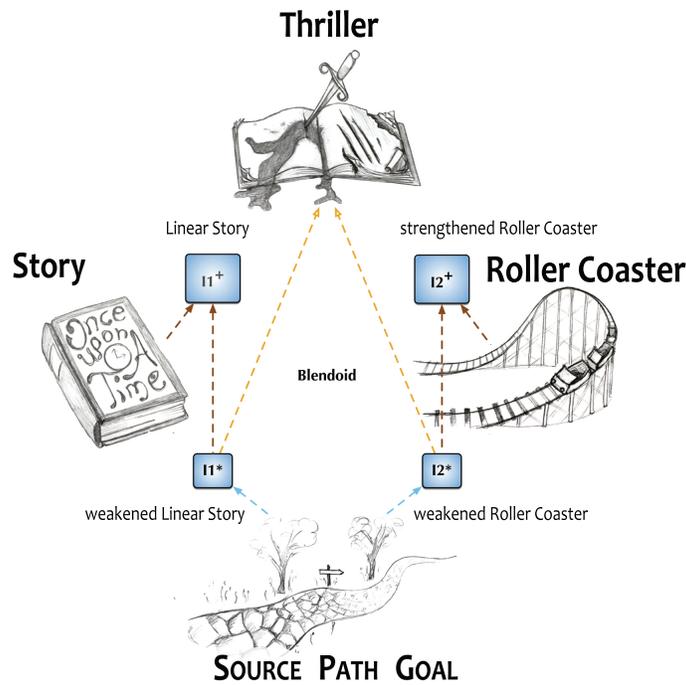


Figure 9: Blending *Thriller* with the input spaces *Story* and *Roller Coaster*.

The newly defined concept, *Thriller*, inherits aspects from the input spaces *Story* and *Roller Coaster* as well as from the *SOURCE\_PATH\_GOAL* spatial schema. In particular, thrillers are

25. The interpretation from the base space to  $I1^*$  reuses exactly the same signature map as in the strengthening process.

26. The `with Story |-> Thriller` part of the definition just renames “Story” into “Thriller” to make the axiomatisation easier to read. The content of the concept is not affected.

scary, they have a fast-paced linear plot, which involves thrill elements, and a protagonist which participates in the events of the plot. (Figure 10 shows some of the axioms of the blended theory.)

```

%% A thriller is scary.
(forall (x) (if (Thriller x) (Scary x)))

%% A thriller has a plot and a character, who is its protagonist
(forall (x)
  (iff
    (Thriller x)
    (exists (y z)
      (and (Plot y) (has_plot x y) (Character z) (has_protagonist x z)))))

%% The plot (of a thriller) is fast-paced and involves a start, at least one thrill
%% element and an end.
(forall (x)
  (if (Plot x)
    (exists (y1 y2 y3)
      (and (fast_paced x) (starts_with x y1) (Start y1) (part_of y2 x)
        (ThrillElement y2) (ends_with x y3) (End y3)))))

%% If x and y are two different events in a plot p,
%% then either x is before y (in p) or the other way round.
(forall (p x y)
  (if
    (and (Plot p) (Event y) (Event x) (part_of y p) (part_of x p) (not (= y x)))
    (or (before x y p) (before y x p))))

```

Figure 10: *Some axioms of the blended concept Thriller*

## 6. Discussion

One of the hardest problems yet to be solved in artificial general intelligence is the generation of concepts and their grounding in the environment, commonly known as the symbol grounding problem. The difficulty lies not only in establishing a relationship between objects in the real world and symbolic as well as mental representations, but also in the problem of defining ‘meaning’ itself.

This paper rests on the basic ideas of grounded and embodied cognition, in which physical experiences are thought of as the primary source that gives meaning to concepts. Indeed, some studies in linguistic neuroscience (e.g. Gallese and Lakoff (2005); Tettamanti et al. (2005)) indicate that the cortical regions of the sensorimotor cortex are activated also in word comprehension tasks. If bodily experiences are a primary source in constituting the meaning of concepts, the symbol grounding problem can be meaningfully approached within this general framework.

We proposed in this paper an approach to computational concept invention in which image schemas, understood as embodied conceptual building blocks, were utilised in conceptual blending, the suggested cognitive machinery behind concept invention. To successfully investigate and evaluate the fruitfulness of this idea, a more comprehensive formalisation of image schemas is needed. Formalising image schemas has been a rather recent undertaking in artificial intelligence research as a means to aid computational concept invention and common-sense reasoning (Kuhn, 2002; Walton and Worboys, 2009; Morgenstern, 2001; Goguen and Harrell, 2010; Kutz et al., 2014a).

So far, the research on image schemas has illuminated the inherent complexity of the formalisation problem due to their abstract cognitive nature. At the same time, the incoherent scholarly terminology and corresponding key definitions make it challenging to find stable ground for further research. In an attempt to bridge research from several research strands, the main part of this article was devoted to introducing the idea of how to formally structure image schemas as families of theories.

Our work differs from the approaches discussed in Section 4.1 by considering the structure of entire image schema families, using PATH as a proof of concept. While other approaches tend to look at the interconnections between particular image schemas, we have followed the psychological research of Mandler and Pagán Cánovas (2014) to analyse formally the PATH-following image schema family concentrating on the involved spatial primitives. It is our belief that this will allow for a more fine-tuned and specialised use of image schemas in computational systems.

The most basic image schemas develop early and become more specialised with experience (Rohrer, 2005). Currently there is no comprehensive and agreed upon list of these most basic image schemas, although the general consensus is that complex image schemas result from combining elements taken from various, more simple, image schemas (Oakley, 2010).

It is therefore likely that there are a limited number of core image schemas that, from their most basic form, can be formally fleshed out into a structure similar to the illustrated PATH image schema graph. Finding the intersections and combinations of these basic families, thus, would mean identifying the source of the cognitive machinery behind the development of complex image schemas. These would be the spatial schemas that contain similar and overlapping spatial information and share certain spatial primitives. In our graph, MOVEMENT\_IN\_LOOPS could be considered to be in the intersection of the image schema families PATH-following and CYCLE. MOVEMENT\_OF\_OBJECT marks the entry point to the PATH-family, yet lacking itself the spatial primitive of ‘path’. Here, the DOL language provides some of the tools to make such an interconnection of families formally feasible, and gives a handle on a formal rendering of the notion of *construal* (image schema transformation) discussed by Clausner and Croft (1999).

A second problem is the temporal nature of image schemas. Since image schemas are not only static but also capture change over time, any axiomatisation thereof needs to address the non-trivial problem of formally representing time. One motivation for the use of non-classical logics is the claim that these are cognitively and linguistically more adequate than classical logics involving variables and direct quantification over objects (Kamp, 1979; Blackburn, de Rijke, and Venema, 2001).

Moreover, the cognitive adequateness of particular formalisms has been studied in detail (e.g. Knauff, Rauh, and Renz (1997)). In this spirit, a large variety of temporal logics has been proposed to model various temporal aspects of natural language (Prior, 1967; Van Benthem, 1983). Similarly, qualitative spatial logics have been designed to capture more adequately the way humans conceptualise and reason about space (Cohn and Renz, 2007).

## 7. Conclusion

We have here presented an approach in which image schemas are treated as interconnected theories in a lattice (ordered by theory interpretation). This was motivated by image schematic structure found in language and the cognitive development of spatial primitives and image schemas. The main insights, we claim, support the hypotheses that the spatial primitives and their assumed properties

distinguish not only the different usages in natural language and various cognitive stages, but can be systematically seen as and mapped to branching points in the lattice of image schema theories.

The benefits of this approach lie not only in the provided structuring of image schemas, but also in how formal systems may use them. By using image schemas in conceptual blending, it is our belief that computational concept invention has taken a step in the right direction. Image schemas provide a cognitively very plausible foundation for the idea of a generic space found in the theory of conceptual blending. In analogy engines, or (formal) approaches to conceptual blending (Turner, 2014; Kutz et al., 2014a), the presented graph of image schemas can provide a method for theory weakening and strengthening based on the involved image schemas, employing basic ideas of amalgams (Ontañón and Plaza, 2010). This approach is therefore substantially different from the more syntactic-driven methods used by the Structure Mapping Engine (SME) (Gentner, 1983; Forbus, Falkenhainer, and Gentner, 1989) or Heuristic-Driven Theory Projection (HDTP) (Schwering et al., 2009; Schmidt et al., 2014).

Future work will focus on extending the presented formalisation approach to other basic image schema families. This will include studying their interconnections, formal methods for their combination to construct complex schemas, as well as algorithmic approaches for detecting image schemas within given input concepts. Conversely, we hope that the systematic study of the formal interconnections between image schema families will have unifying value also for image schema research within the cognitive sciences, and provide some of the still missing systematicity to the field. Finally, we believe that the semantic and cognitive grounding of the idea of a generic space in the notion of image schema has great potential for computational realisations of conceptual blending.

**Acknowledgments.** We thank the reviewers for constructive and valuable feedback. We would also like to thank John Bateman, Tarek R. Besold, Emiliós Cambouropoulos, Tony Veale, and Mihailo Antović for valuable input and interesting discussions on topics related to this paper.

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 number: 611553.

## References

- Aguilar, W., and Pérez y Pérez, R. 2015. Dev E-R: A computational model of early cognitive development as a creative process. *Cognitive Systems Research* 33:17–41.
- Ahrens, K., and Say, A. L. 1999. Mapping Image Schemas and Translating Metaphors. In *Proceedings of Pacific Asia Conference on Language, Information and Computation*, 1–8.
- Allen, J., and Hayes, P. 1985. A Common-Sense Theory of Time. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence (IJCAI-85)*, 528–531.
- Barsalou, L. W. 2008. Grounded cognition. *Annual review of psychology* 59:617–645.
- Bennett, B., and Cialone, C. 2014. Corpus Guided Sense Cluster Analysis: a methodology for ontology development (with examples from the spatial domain). In Garbacz, P., and Kutz, O., eds., *8th International Conference on Formal Ontology in Information Systems (FOIS)*, volume 267 of *Frontiers in Artificial Intelligence and Applications*, 213–226. IOS Press.

- Blackburn, P.; de Rijke, M.; and Venema, Y. 2001. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press.
- Boroditsky, L. 2000. Metaphoric structuring: Understanding time through spatial metaphors. *Cognition* 75(1):1–28.
- Brugman, C., and Lakoff, G. 1988. Cognitive Topology and Lexical Networks. In Stephen Small, Gary Cottrell, M. T., ed., *Lexical ambiguity resolution*. 477–508.
- Chrisley, R. 2003. Embodied artificial intelligence. *Artificial Intelligence* 149:131–150.
- Clausner, T. C., and Croft, W. 1999. Domains and image schemas. *Cognitive Linguistics* 10(1):1–31.
- Cohn, A. G., and Renz, J. 2007. Qualitative Spatial Representation and Reasoning. In van Harmelen et al., F., ed., *Handbook of Knowledge Representation*. Oxford: Elsevier. 551–596.
- Davidson, D. 1967. The logical form of action sentences. In Rescher, N., ed., *The logic of decision and action*. 81–94.
- Fauconnier, G., and Turner, M. 1998. Conceptual Integration Networks. *Cognitive Science* 22(2):133–187.
- Fauconnier, G., and Turner, M. 2003. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. Basic Books.
- Feldman, J., and Narayanan, S. 2004. Embodied meaning in a neural theory of language. *Brain and Language* 89(2):385–392.
- Forbus, K.; Falkenhainer, B.; and Gentner, D. 1989. The structure-mapping engine. *Artificial Intelligence* 41:1–63.
- Gallese, V., and Lakoff, G. 2005. The Brain's concepts: the role of the Sensory-motor system in conceptual knowledge. *Cognitive neuropsychology* 22(3):455–79.
- Gärdenfors, P. 2007. *Embodiment in Cognition and Culture*, volume 71 of *Advances in Consciousness Research*. John Benjamins Publishing Company. chapter Cognitive semantics and image schemas with embodied forces, 57–76.
- Gentner, D. 1983. Structure mapping: A theoretical framework for analogy. *Cognitive Science* 7(2):155–170.
- Gibson, J. J. 1977. The theory of affordances, in Perceiving, Acting, and Knowing. Towards an Ecological Psychology. In Shaw, R., and Bransford, J., eds., *Perceiving, Acting, and Knowing: Toward an Ecological Psychology*. Hillsdale: NJ: Lawrence Erlbaum. 67–82.
- Goguen, J. A., and Harrell, D. F. 2010. Style: A Computational and Conceptual Blending-Based Approach. In Argamon, S., and Dubnov, S., eds., *The Structure of Style: Algorithmic Approaches to Understanding Manner and Meaning*. Berlin: Springer. 147–170.
- Goldin, G. A. 2001. Counting on the metaphorical. *Nature* 413(6851):18–19.

- Grüninger, M.; Hahmann, T.; Hashemi, A.; Ong, D.; and Ozgovde, A. 2012. Modular First-Order Ontologies Via Repositories. *Applied Ontology* 7(2):169–209.
- Hedblom, M. M.; Kutz, O.; and Neuhaus, F. 2014. On the cognitive and logical role of image schemas in computational conceptual blending. In *Proceedings of the 2nd International Workshop on Artificial Intelligence and Cognition (AIC-2014), Torino, Italy, November 26th–27th*, volume Volume 1315 of *CEUR-WS*.
- Hedblom, M. M.; Kutz, O.; and Neuhaus, F. 2015. Image Schemas as Families of Theories. In Besold, T. R.; Kühnberger, K.-U.; Schorlemmer, M.; and Smaill, A., eds., *Proceedings of the Workshop “Computational Creativity, Concept Invention, and General Intelligence” 2015*, volume 2 of *Publications of the Institute of Cognitive Science*, 19–33. Institute of Cognitive Science.
- Johanson, M., and Papafragou, A. 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, M. 1987. *The Body in the Mind. The Bodily Basis of Meaning, Imagination, and Reasoning*. The University of Chicago Press.
- Kamp, H. 1979. Instants, events and temporal discourse. In Bäuerle, R.; Schwarze, C.; and von Stechow, A., eds., *Semantics from Different Points of View*. Springer, Berlin. 376–417.
- Knauff, M.; Rauh, R.; and Renz, J. 1997. A cognitive assessment of topological spatial relations: Results from an empirical investigation. In Hirtle, S. C., and Frank, A. U., eds., *Spatial Information Theory: A Theoretical Basis for GIS*, volume 1329 of *Lecture Notes in Computer Science*. Springer. 193–206.
- Koestler, A. 1964. *The Act of Creation*. Macmillan.
- Kövecses, Z. 2010. *Metaphor: A Practical Introduction*. Oxford University Press, USA.
- Kuhn, W. 2002. Modeling the Semantics of Geographic Categories through Conceptual Integration. In *Proceedings of GIScience 2002*, 108–118. Springer.
- Kuhn, W. 2007. An Image-Schematic Account of Spatial Categories. In Winter, S.; Duckham, M.; Kulik, L.; and Kuipers, B., eds., *Spatial Information Theory*, volume 4736 of *Lecture Notes in Computer Science*. Springer. 152–168.
- Kutz, O.; Mossakowski, T.; Hois, J.; Bhatt, M.; and Bateman, J. 2012. Ontological Blending in DOL. In Besold, T.; Kühnberger, K.-U.; Schorlemmer, M.; and Smaill, A., eds., *Computational Creativity, Concept Invention, and General Intelligence, Proceedings of the 1st International Workshop C3GI@ECAI*, volume 01-2012. Montpellier, France: Publications of the Institute of Cognitive Science, Osnabrück.
- Kutz, O.; Bateman, J.; Neuhaus, F.; Mossakowski, T.; and Bhatt, M. 2014a. E pluribus unum: Formalisation, Use-Cases, and Computational Support for Conceptual Blending. In Besold, T. R.; Schorlemmer, M.; and Smaill, A., eds., *Computational Creativity Research: Towards Creative Machines*, Thinking Machines. Atlantis/Springer.

- Kutz, O.; Neuhaus, F.; Mossakowski, T.; and Codescu, M. 2014b. Blending in the Hub—Towards a collaborative concept invention platform. In *Proceedings of the 5th International Conference on Computational Creativity*.
- Kutz, O.; Mossakowski, T.; and Lücke, D. 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, G., and Johnson, M. 1999. *Philosophy in the Flesh*. Basic Books.
- Lakoff, G., and Núñez, R. 2000. *Where Mathematics Comes from: How the Embodied Mind Brings Mathematics Into Being*. New York: Basic Books.
- Lakoff, G. 1987. *Women, Fire, and Dangerous Things. What Categories Reveal about the Mind*. The University of Chicago Press.
- Louwerse, M. M., and Jeuniaux, P. 2010. The linguistic and embodied nature of conceptual processing. *Cognition* 114(1):96–104.
- Mandler, J. M., and Pagán Cánovas, C. 2014. On defining image schemas. *Language and Cognition* 0:1–23.
- Mandler, J. M. 1992. How to build a baby: II. Conceptual primitives. *Psychological review* 99(4):587–604.
- Mandler, J. M. 2004. *The Foundations of Mind : Origins of Conceptual Thought: Origins of Conceptual Thought*. New York: Oxford University Press.
- Mandler, J. M. 2008. On the Birth and Growth of Concepts. *Philosophical Psychology* 21(2):207–230.
- Menzel, C. 2011. Knowledge representation, the World Wide Web, and the evolution of logic. *Synthese* 182:269–295.
- Morgenstern, L. 2001. Mid-Sized Axiomatizations of Commonsense Problems: A Case Study in Egg Cracking. *Studia Logica* 67:333–384.
- Mossakowski, T.; Kutz, O.; Codescu, M.; and Lange, C. 2013. The Distributed Ontology, Modeling and Specification Language. In Del Vescovo, C.; Hahmann, T.; Pearce, D.; and Walther, D., eds., *Proceedings of the 7th International Workshop on Modular Ontologies (WoMO-13)*, volume 1081. CEUR-WS.
- Mossakowski, T.; Codescu, M.; Neuhaus, F.; and Kutz, O. 2015. *The Road to Universal Logic—Festschrift for 50th birthday of Jean-Yves Beziau, Volume II*. Studies in Universal Logic. Birkhäuser. chapter The distributed ontology, modelling and specification language - DOL.
- Mossakowski, T.; Kutz, O.; and Codescu, M. 2014. Ontohub: A semantic repository for heterogeneous ontologies. In *Proc. of the Theory Day in Computer Science (DACS-2014)*. Satellite workshop of ICTAC-2014.

- Mossakowski, T.; Lange, C.; and Kutz, O. 2012. Three Semantics for the Core of the Distributed Ontology Language. In Grüninger, M., ed., *7th International Conference on Formal Ontology in Information Systems (FOIS)*, Frontiers in Artificial Intelligence and Applications. IOS Press.
- Mossakowski, T.; Maeder, C.; and Lüttich, K. 2007. The Heterogeneous Tool Set. In Grumberg, O., and Huth, M., eds., *Tools and Algorithms for the Construction and Analysis of Systems. 13th International Conference, TACAS 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007 Braga, Portugal, March 24 - April 1, 2007. Proceedings*, volume 4424 of *Lecture Notes in Computer Science*, 519–522. Springer.
- Nayak, S., and Mukerjee, A. 2012. Concretizing the image schema: How semantics guides the bootstrapping of syntax. In *2012 IEEE International Conference on Development and Learning and Epigenetic Robotics, ICDL 2012*.
- Oakley, T. 2010. Image Schema. In Geeraerts, D., and Cuyckens, H., eds., *The Oxford Handbook of Cognitive Linguistics*. Oxford: Oxford University Press. 214–235.
- Ontañón, S., and Plaza, E. 2010. Amalgams: A formal approach for combining multiple case solutions. In *Case-Based Reasoning. Research and Development*. Springer. 257–271.
- Parson, T. 1990. *Events in the Semantics of English: A Study in Subatomic Semantics*. MIT Press.
- Pereira, F. C., and Cardoso, A. 2003. Optimality Principles for Conceptual Blending: A First Computational Approach. *AISB Journal* 1(4).
- Prior, A. N. 1967. *Past, Present and Future*. Oxford: Oxford University Press.
- Regier, T. 1996. *The Human Semantic Potential: Spatial Language and Constrained Connectionism*. The MIT Press.
- Rohrer, T. 2005. Image Schemata in the Brain. In Hampe, B., and Grady, J. E., eds., *From perception to meaning: Image schemas in cognitive linguistics*, volume 29 of *Cognitive Linguistics Research*. Walter de Gruyter. 165–196.
- Schiralli, M., and Sinclair, N. 2003. A Constructive Response to ‘Where Mathematics Comes From’. *Educational Studies in Mathematics* 52:79–91.
- Schmidt, M.; Krumnack, U.; Gust, H.; and Kuhnberger, K.-U. 2014. Heuristic-Driven Theory Projection: An Overview. In Prade, H., and Richard, G., eds., *Computational Approaches to Analogical Reasoning: Current Trends*, Computational Intelligence 548. Springer-Verlag.
- Schorlemmer, M.; Smaill, A.; Kühnberger, K.-U.; Kutz, O.; Colton, S.; Cambouropoulos, E.; and Pease, A. 2014. COINVENT: Towards a Computational Concept Invention Theory. In *Proceedings of the 5th International Conference on Computational Creativity*.
- Schwering, A.; Krumnack, U.; Kühnberger, K.-U.; and Gust, H. 2009. Syntactic Principles of Heuristic-Driven Theory Projection. *Cognitive Systems Research* 10(3):251–269.
- St. Amant, R.; Morrison, C. T.; Chang, Y.-H.; Cohen, P. R.; and Beal, C. 2006. An image schema language. In *International Conference on Cognitive Modeling (ICCM)*, 292–297.

- Tettamanti, M.; Buccino, G.; Saccuman, M. C.; Gallese, V.; Danna, M.; Scifo, P.; Fazio, F.; Rizzolatti, G.; and Perani, D. 2005. Listening to action-related sentences activates fronto-parietal motor circuits. *Journal of Cognitive Neuroscience* 273–281.
- Turner, M. 2007. The Way We Imagine. In Roth, I., ed., *Imaginative Minds - Proceedings of the British Academy*. Oxford: OUP. 213–236.
- Turner, M. 2014. *The Origin of Ideas: Blending, Creativity, and the Human Spark*. Oxford University Press.
- Van Benthem, J. F. A. K. 1983. *The Logic of Time*. Dordrecht, Holland: D. Reidel Publishing Company.
- Veale, T., and Keane, M. T. 1992. Conceptual Scaffolding: a Spatially Founded Meaning Representation for Metaphor Comprehension. *Computational Intelligence* 8(3):494–519.
- Veale, T.; Feyaerts, K.; and Forceville, C. 2013. E Unis Pluribus: Using Mental Agility to Achieve Creative Duality in Word, Image and Sound. In *Creativity and the Agile Mind: A Multi-Disciplinary Study of a Multi-Faceted Phenomenon (Applications of Cognitive Linguistics)*. 37–57.
- Vernon, D. 2014. *Artificial Cognitive Systems: A Primer*. MIT Press.
- Voorhees, B. 2004. Embodied Mathematics: Comments on Lakoff and Núñez. *Journal of Consciousness Studies* 11(9):83–88.
- Walton, L., and Worboys, M. 2009. An Algebraic Approach to Image Schemas for Geographic Space. In *Proceedings of the 9th International Conference on Spatial Information Theory (COSIT)*, 357–370.
- Wierzbicka, A. 1996. *Semantics : Primes and Universals: Primes and Universals*. Oxford University Press, UK.
- Wilson, N. L., and Gibbs, R. W. 2007. Real and imagined body movement primes metaphor comprehension. *Cognitive science* 31(4):721–731.