

E pluribus unum*

Formalisation, Use-Cases, and Computational Support for Conceptual Blending

Oliver Kutz, John Bateman, Fabian Neuhaus, Till Mossakowski, and Mehul Bhatt

Abstract Conceptual blending has been employed very successfully to understand the process of concept invention, studied particularly within cognitive psychology and linguistics. However, despite this influential research, within computational creativity little effort has been devoted to fully formalise these ideas and to make them amenable to computational techniques. 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 structural similarity based on a relation of analogy, identified in a generic space, the base ontology. We here introduce the basic formalisation of conceptual blending, as sketched by the late Joseph Goguen, and discuss some of its variations. We illustrate the vast array of conceptual blends that may be covered by this approach and discuss the theoretical and conceptual challenges that ensue. Moreover, we show how the Distributed Ontology Language DOL can be used to declaratively specify blending diagrams of various shapes, and discuss in detail 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 logical languages and formal linking constructs.

Oliver Kutz, Fabian Neuhaus, Till Mossakowski
Institute of Knowledge and Language Engineering, Otto-von-Guericke University of Magdeburg,
Germany. e-mail: {okutz, mossakow, fneuhaus}@iws.cs.ovgu.de

John Bateman
Faculty of Linguistics and Literary Sciences / Research Center on Spatial Cognition (SFB/TR 8),
University of Bremen, Germany. e-mail: bateman@uni-bremen.de

Mehul Bhatt
Research Center on Spatial Cognition (SFB/TR 8), University of Bremen, Germany. e-mail:
bhatt@informatik.uni-bremen.de

* This paper summarises our work on conceptual blending, to create a new concept out of many, with material drawn from Hois et al. (2010); Kutz et al. (2012, 2014) and Neuhaus et al. (2014).

1 Concept Invention via Blending

One broad area of phenomena that is often brought into connection with issues of creativity and the emergence of new ideas concerns notions such as metaphor, blending, category mistakes, similes, analogies and the like. In each of these, seemingly inconsistent material is combined in a manner that results in a productive growth of information instead of simple logical contradiction. Approaches to treat this phenomenon are varied but commonly come to the conclusion that more or less well developed notions of ‘structure’ are crucial for bringing the growth of information about — e.g., ‘implication complexes’ for metaphor (Black 1979), ‘conceptual spaces’² for blending (Fauconnier and Turner 2003), ‘structure mapping’ in analogy (Gentner 1983), and so on. On the one hand, the less structure that is available, the less productive the combinations appear to be; on the other, the presence of structure raises the challenge of how such formal commitments can be productively ‘overridden’ or rearranged in order to avoid contradiction.

In our ongoing work on ontology and its formal underpinnings, we have been led to a very similar set of questions. By ‘ontology’ we here refer to the now rather standard notion of a formal specification of a shared understanding of the entities, relations and general properties holding in some domain of interest (cf. Guarino 1994; Gruber 1995). Achieving adequate treatments in various domains has demonstrated to us the need for heterogeneous ontological specifications that are capable of capturing distinct perspectives on the phenomena being modelled. In an architectural context, for example, it is beneficial to maintain distinct perspectives on structural integrity, spatial distribution, movement patterns by the occupants of a building (‘flow’), navigation networks (possibly varying according to ‘normal’ and ‘emergency’ conditions), ‘visibility’ patterns (both for users and for sensors in the case of security) and many more (Bhatt et al. 2012) — each of these perspectives can be modelled well by employing ontological engineering techniques but there is no guarantee that they are simply compatible. Our work on natural language dialogue systems involving spatial language comes to the same conclusion (Bateman 2010), while similar concerns are already well known in Geographic Information Science (Frank and Kuhn 1999; Kuhn 2003). To support this fundamental ‘multi-perspectivalism’ we have therefore been developing an entire toolset of more sophisticated combination methods (Kutz et al. 2008a), leading to the formal definition of the notion of a ‘hyperontology’ in Kutz et al. (2010).

The similarities apparent between the goals of heterogeneous ontology ‘alignment’ and the creative combination of thematically distinct information spaces can be built on quite concretely by treating such information spaces explicitly in terms of ontological specifications. This allows us to link directly with previous work by exploring the application of techniques for combining distinct perspectives that are now becoming available. For example, much work on creativity has been pursued in the context of Fauconnier and Turner’s (2003) account of conceptual blending,

² The usage of the term ‘conceptual space’ in blending theory is not to be confused with the usage established by Gärdenfors (2000).

in which the blending of two thematically rather different *conceptual spaces* yields a new conceptual space with emergent structure, selectively combining parts of the given spaces whilst respecting common structural properties. The ‘imaginative’ aspect of blending is summarised as follows in Turner (2007):

[...] the two inputs have different (and often clashing) organising frames, and the blend has an organising frame that receives projections from each of those organising frames. The blend also has emergent structure on its own that cannot be found in any of the inputs. Sharp differences between the organising frames of the inputs offer the possibility of rich clashes. Far from blocking the construction of the network, such clashes offer challenges to the imagination. The resulting blends can turn out to be highly imaginative.

We see the almost unlimited space of possibilities supported by ‘ontological blending’ for combining existing ontologies to create new ontologies with emergent structure as offering substantial benefits not only for ontological engineering — where conceptual blending can be built on to provide a structural and logic-based approach to ‘creative’ ontological engineering — but also for conceptual blending and related frameworks themselves — by providing a far more general and nevertheless computational, formalised foundation. Re-considering some of the classic problems in conceptual blending in terms of ontological modelling and ontological blending opens up an exciting direction for future research.

This endeavour primarily raises the following two challenges: (1) when combining the terminologies of two ontologies, the shared semantic structure is of particular importance to steer possible combinations — this shared semantic structure leads to the notion of a base ontology, which is closely related not only to the notion of ‘tertium comparationis’ found in classical rhetoric and poetics, but also to more recent cognitive theories of metaphor (see, e.g., Jaszczolt (2003)); (2) having established a shared semantic structure, there typically remains a considerable number of possibilities that can capitalise on this information in the combination process — here, structural optimality principles as well as ontology evaluation techniques can take on a central role in selecting ‘interesting’ blends.

There is still much to explore concerning the relationships between the principles governing ontological blending and the principles explored to date for blending phenomena in language or poetry or, indeed, the rather strict principles ruling blending in mathematics, in particular in the way formal inconsistencies are dealt with. For instance, whilst blending in poetry might be particularly inventive or imaginative when the structure of the basic categories found in the input spaces is almost completely ignored, in areas such as mathematics a rather strict adherence to sort structure is important in order to generate meaningful blends.³ The use that we might typically make of ontological blending is situated somewhere in the middle: re-arrangement and new combination of basic categories can be quite interesting, but has to be finely controlled through corresponding interfaces, often regulated by or related to choices found in foundational or upper ontologies so that basic categorial relationships are maintained.

³ For instance when creating the theory of transfinite cardinals by blending the perfective aspect of counting up to any fixed finite number with the imperfective aspect of ‘endless counting’ Núñez (2005).

For all such cases, however, we can consider the formal mechanisms that support specific blends that we explore with respect to their potential relevance and value for understanding ‘blending’ phenomena in general. This will be the main purpose of the current chapter. We will summarise some of the progress that has been made in recent years towards adopting the fruitful idea of conceptual blending in a theoretically well-understood and computationally supported formal model for concept invention, focusing in particular on ontology languages. Here we elaborate on ideas first introduced in Hois et al. (2010), with detailed technical definitions given in Kutz et al. (2012). More specifically, we:

- briefly characterise the kinds of creativity that have been considered hitherto in the areas of blends, metaphors and related operations where structured mappings or analogies are relied upon;
- sketch the logical analysis of conceptual blending in terms of blending diagrams and colimits, as originally proposed by Joseph Goguen, and give an abstract definition of ontological blendoids capturing the basic intuitions of conceptual blending in the ontological setting;
- sketch a formal meta-language, namely the distributed ontology language DOL, that is capable of declaratively specifying blending diagrams in a variety of ontology languages. This provides a structured approach to ontology languages and blending and combines the simplicity and good tool support for languages such as OWL⁴ with the more complex blending facilities of OBJ3 (Goguen and Malcolm 1996) or Haskell (Kuhn 2002); DOL also facilitates the specification of a range of variations of the basic blending technique;
- discuss the capabilities of the Ontohub/HETS ecosystem with regard to collaboratively managing, creating, and displaying blended concepts, ontological theories, and entire blending diagrams; this includes an investigation of the evaluation problem in blending, together with a discussion of structural optimality principles and current automated reasoning support.

We close with a discussion of open problems and future work.

2 An Ocean of Blends

In this section, we briefly characterise the rather diverse phenomena that may be subject to beneficial formalisations in terms of ontological blending. The starting point is the obvious one of conceptual blending, which we use as a prototypical case of emergent organisation throughout this chapter. As noted above, conceptual blending in the spirit of Fauconnier and Turner (2003) operates by combining two input ‘conceptual spaces’, construed as rather minimal descriptions of some thematic domains, in a manner that creates new ‘imaginative’ configurations. A classic example for this is the blending of the concepts *house* and *boat*, yielding as most

⁴ With ‘OWL’ we refer to OWL 2 DL, see <http://www.w3.org/TR/owl2-overview/>

straightforward blends the concepts of a *houseboat* and a *boathouse*, but also an *amphibious vehicle* (Goguen and Harrell 2010); we return to this example below. This case shows well how it is necessary to maintain aspects of the structural semantics of the spaces that are blended in order to do justice to the meanings of the created terms: the houseboat stops neither being a vehicle on water nor being a place of residence, for example.

Very similar processes appear to be operating in cases of metaphor (Black 1979; Kövecses 2010). Here a semantically structured ‘source’ is used so that facets of the semantics of the source are selected for appropriate take up by a semantically structured ‘target’. This can operate on a small scale, analogously to the *house* and the *boat*, as for example in metaphors such as that evident in the 1940s film title “Wolf of New York” or the recent “The Wolf of Wall Street” (2013), where certain conventionalised properties of the wolf as animal (the source) are transferred to the people referred to by the titles (the target). Structure is essential here since the transfer is very specific: a reading of the metaphor in which ‘four-leggedness’ or ‘furry’ is transferred is in the given contexts most unlikely. Only particular relations and relational values are effected. Metaphors can also operate on much broader scales, as in considerations of metaphors as contributions to creative scientific theory construction, as in the well known transfer of a ‘sun-and-planet’ conceptual model to models of the atom (Miller 2000; Guhe et al. 2011; Gust et al. 2003) (see Sec. 3.1 below). Structural transfer of this kind has consequently been suggested to play a substantial role for persuasive text creation as such. Hart, for example, discusses the use of phrases such as ‘limitless flow of immigration’, ‘flood of asylum seekers’ and so on as ideologically-loaded constructions that need to be unpacked during critical discourse analysis (Hart 2008).

Metaphors also bring with them some particular formal features of their own — for example, they are typically seen as *directed* in contrast to blends and have been related to models of embodiment via accounts of *image schemas* (Johnson 1987). Image schemas suggest how multimodal patterns of experience can be linked to increasingly abstract conceptualisations: abstract thought is then seen as a metaphorical construction on top of concrete experience. The use of the word ‘flood’ in the above example can then be expected to bring about a physical component in its reception where feelings of force, damage and lack of control are activated; this makes it clear that much more than ‘flowery language’ might be involved in such phrasings and their selection.

A related consideration is the proposal for internalised spatial representations for supporting reasoning and more abstract conceptualisations (such as time) as well as externalised spatial representations for diagrammatic reasoning. In the former case, it is common to work within blended spaces where time and spatial extent appear to have ‘collapsed’, giving rise to language use such as “keep going straight until the church” or “turn left before the tower” and so on. Blends of this kind are so familiar that they may be considered to be *entrenched* in the cognitive linguistic sense of having become part of the semantics of the respective terms and shared by the language community (Fauconnier and Turner 2003: 49).

Blends may also be multiple in that once established, for example in a text, further conceptual spaces might be added as an argument progresses. These may progressively add details to a developing emergent space (or, alternatively, lead to a space which strains the credibility of a reader or hearer too far resulting in a charge of ‘mixing metaphors’). In the right-wing immigration example above from Hart, the texts do in fact continue with phrases such as ‘Britain is full up’, ‘no matter how open or closed its immigration policy’, and ‘our first step will be to shut the door’. This builds on the previous blend of immigration-as-flood by (i) combining ‘Britain’ with a ‘container’ (which can then be full) that is itself (ii) combined with a ‘building’ or ‘room’ that has ‘doors’ that can be closed, and (iii) those doors can in turn also be ‘policies’ (which can be open or closed) (Hart 2008: 102). There need in principle be no end to this creative extension and combination of concepts. This aspect of iteration of blending is also explored in the area of conceptual mathematics as explored in Lakoff and Núñez (2000), where it is argued that abstract mathematical concepts such as modern number systems, algebra, or set theory, are created through a succession of conceptual metaphors and blends, grounded in embodied concepts and image schemas. The structure of such blends and blending patterns in general are discussed more formally in Section 4.2.

There is also now increasing discussion of the potential role of blending or similar mechanisms when considering the creative use of combinations of information from different semiotic modes, e.g., drawing relations between verbal information, visual and gestural information (Forceville and Urios-Aparisi 2009).




Fig. 1 Visual blending of a car and an anatomical representation used for advertising discussed by van Mulken et al. (2010).

In such cases representations or entities in one mode of presentation are made to take on properties or behaviours in another. The general applicability of an ontological approach to semiotic blending of this kind is argued in Bateman (2011). Again, there are many examples of such creativity in action. Consider for example the extract from an advertisement discussed by van Mulken et al. (2010) and shown in Fig. 1. Here an open-ended set of potential further inferences, all supporting the general intention of the advertisement, is opened up by virtue of the blend. There are also commonly discussed combinations such as the

use of space for time in comics and visual narrative — moving across the space of a comic’s panel, typically in Western comics therefore from left to right, often correlates with a progression in time (McCloud 1994: 95) — as well as blends for dramatic or emotional effect, such as when typography is shaped visually for affective purposes (Eisner 1992: 12).

A particularly creative and novel example of semiotic blending across media can be seen in the following example. In this case the film director Ang Lee works with the *dynamic* possibilities of the film medium to enlist graphic resources for ex-


Wonderwoman with motion lines connecting panels



Pérez, G., Potter, G. and Wein, L. (2008), *Biblioteca DC – Mulher-Maravilha/Wonder Woman*, Brasil, São Paulo: Panini. p. 140.

Hulk with motion panels connecting actions

1:37:26.7



camera tracking left,
following helicopter

camera tracking left,
following helicopter.
insert panel moves from
left to right, picks up Hulk
as he enters frame right

second insert panel moves
from left to right, picks up
Hulk just as he leaves first
insert panel moving left

complete right-to-left
trajectory of Hulk over three
insert panels

Lee, Ang (2003), *Hulk*, based on Marvel characters by Stan Lee and Jack Kirby, Screenplay: John Turman, Michael France and James Schamus, Story: James Schamus, USA: Universal Pictures, Marvel Enterprises, Valhalla Motion Pictures, Good Machine.

Fig. 2 Blending expressive resources from comics and film.

pressing movement developed within the *static* medium of comics. The result is an interesting and highly explorative expansion of the creative potential of what can be done with film. An illustration is shown in Fig. 2. On the left for comparison is a now quite traditional static rendition of movement from a comic — in this particular case showing ‘continuity’ of movement across panels. In contrast, on the right we see a short sequence of stills taken from a chase scene in Lee’s film *Hulk* (2003), where the main character is trying to escape from pursuers in a helicopter. In this case, the escape trajectory is shown in a sequence of dynamically inserted ‘panels’ that move across the screen to the point where they can pick up the character’s movement. This blending of properties in Lee’s film does much more than ‘re-create’ a visual effect analogous to comics as sometimes suggested in analyses of this film. Lee’s appropriation of framing and movement techniques within an already dynamic medium appears instead to provide a resource that considerably heightens continuity for narrative effect. A more detailed discussion of the consequences of this appropriation for interpretation and reception is given in Bateman and Veloso (2013).

We are just beginning to be able to explore extensions of meaning-making potential of these kinds. Indeed, although there are now many examples in the literature of such creative meaning growth in action, deep questions remain concerning how precisely this may be modelled. In particular, following simpler operations of ‘alignment’ of structures across spaces (e.g., by graph matching (Forbus et al. 1989; Veale and O’Donoghue 2001)), it is by no means clear how the results that are achieved can function as productively as they evidently do. This relates also to

Fauconnier’s suggestion that it is actually what is done with the result of blending, termed *elaboration* (or ‘running the blend’), that is the most significant stage of the entire blending process. Elaboration “consists in cognitive work performed within the blend, according to its own emergent logic” (Fauconnier 1997: 151). This makes it evident that something more is required in the formalisation than a straightforward recording or noting of a structural alignment: a new blended theory should also be ‘logically productive’, with new and surprising entailments which may well be quite specific to the blend. This is therefore another motivation for the rather more formal and ontologically-driven approach to this kind of creative meaning creation that we now present.

3 Blending Computationalised

There have now been several approaches moving towards effective computational treatments of blending, metaphor and related constructs such as analogy (cf. e.g. Veale and O’Donoghue 2001; Pereira 2007; Schwering et al. 2009b; Li et al. 2012; Veale 2012; Mamakos et al. 2014). Here we follow the research direction of *algebraic semiotics* established by Goguen. In this approach certain structural aspects of semiotic systems are logically formalised in terms of algebraic theories, sign systems, and their mappings (Goguen 1999). Sign systems are theories ‘with extra structure’ connected by a particular class of mappings, which Goguen terms ‘semiotic morphisms’, which preserve that extra structure to a greater or lesser degree. In Goguen and Harrell (2010), algebraic semiotics has been applied to user interface design and blending. Algebraic semiotics does not claim to provide a comprehensive formal theory of blending—indeed, Goguen and Harrell admit that many aspects of blending, in particular concerning the meaning of the involved notions, as well as the optimality principles for blending, cannot be captured formally. However, the structural aspects *can* be formalised and provide insights into the space of possible blends.

Goguen defines sign systems as algebraic theories that can be formulated by using the algebraic specification language OBJ3 (Goguen and Malcolm 1996). One special case of such a sign system is a *conceptual space*: it consists only of constants and relations, one sort, and axioms that define that certain relations hold on certain instances.

We now relate such spaces to a general formalisation of ontologies as we understand them and as introduced above. Since we will focus on standard ontology languages, namely OWL and first-order logic, we use these to replace the logical language OBJ3 used by Goguen and Malcolm. However, as some structural aspects are necessary in the ontology language to support blending, we augment these standard ontology languages with structuring mechanisms known from algebraic specification theory (Kutz et al. 2008b). Such mechanisms are now included in the DOL language specification discussed below in Section 4. This allows us to translate most parts of Goguen’s theory to these augmented ontology languages. Goguen’s main

insight has been that sign systems and conceptual spaces can be related via *morphisms*, and that blending is comparable to *colimit* construction. In particular, the blending of two concepts is often a *pushout* (also called a *blendoid* in this context). Some basic definitions we then need are the following.⁵

Non-logical symbols are grouped into **signatures**, which for our purposes can be regarded as collections of typed symbols (e.g. concept names, relation names). **Signature morphisms** are maps between signatures that preserve (at least) types of symbols (i.e. map concept names to concept names, relations to relations, etc.). A **theory** or **ontology** pairs a signature with a set of sentences over that signature, and a **theory morphism** (or **interpretation**) between two theories is just a signature morphism between the underlying signatures that preserves logical consequence, that is, $\rho : T_1 \rightarrow T_2$ is a theory morphism if $T_2 \models \rho(T_1)$, i.e. all the translations of sentences of T_1 along ρ follow from T_2 . This construction is completely logic independent. Signature and theory morphisms are an essential ingredient for describing conceptual blending in a logical way.

We can now give a general definition of ontological blending capturing the basic intuition that a blend of input ontologies shall partially preserve the structure imposed by base ontologies, but otherwise be an almost arbitrary extension or fragment of the disjoint union of the input ontologies with appropriately identified base space terms.

For the following definition, a variant of which we first introduced in Kutz et al. (2012), a diagram consists of a set of ontologies (the nodes of the diagram) and a set of morphisms between them (the arrows of the diagram). The **colimit** of a diagram is similar to a disjoint union of its ontologies, with some identifications of shared parts as specified by the morphisms in the diagram. We refrain from presenting the category-theoretic definition here (which can be found in Adámek et al. (1990)), but will explain (the action of) the colimit operation in the examples in Section 4.1. In the following definition, we use $|D|$ to denote the set of all nodes in a diagram.

Definition 1 (Ontological Base Diagram). An **ontological base diagram** is a diagram D for which a distinguished set $\mathcal{B} = \{B_i \mid i \in I\} \subset |D|$ of nodes are called **base ontologies**, and where a second distinguished set of nodes $\mathcal{I} = \{I_j \mid j \in J\} \subset |D|$ are called **input ontologies**, and where the theory morphisms $\mu_{ij} : B_i \rightarrow I_j$ from base ontologies to input ontologies are called the **base morphisms**.

If there are exactly two inputs I_1, I_2 , and precisely one base $B \in \mathcal{B}$ and two base morphisms $\mu_k : B \rightarrow I_k, k = 1, 2$, the diagram D is called **classical** and has the shape of a ‘V’. In this case, \mathcal{B} is also called the **tertium comparationis**.

Fig. 3 illustrates the basic, classical case of an ontological blending diagram. The lower part of the diagram shows the base 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

⁵ Note that these definitions apply not only to OWL, but also to many other logics. Indeed, they apply to *any* logic formalised as an *institution* (Goguen and Burstall 1992).

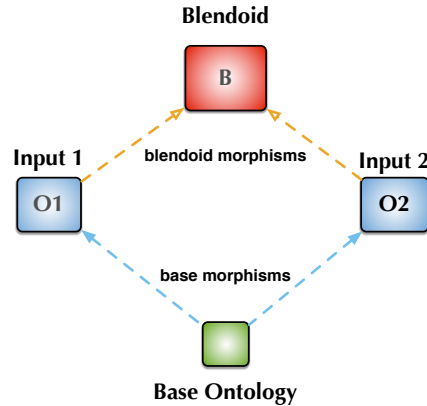


Fig. 3 The basic integration network for blending: concepts in the base ontology are first refined to concepts in the input ontologies and then selectively blended into the blendoid.

consistent subset of the colimit of the base diagram may be seen as a newly invented concept, a **blendoid**.⁶ 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. We will further discuss this and give some examples in Section 4.2.

3.1 Computing the Tertium Comparationis

To find candidates for base ontologies that could serve for the generation of ontological blendoids, much more shared semantic structure is required than the surface similarities that statistical term alignment approaches rely on (Euzenat and Shvaiko 2007). The common structural properties of the input ontologies that are encoded in the base ontology are typically of a more abstract nature. The standard example here relies on *image schemata*, such as the notion of a *container* (see e.g. Kuhn (2002)). Thus, in particular, foundational ontologies can support such selections. In analogical reasoning, ‘structure’ is (partially) mapped from a source domain to a target domain (Forbus et al. 1989; Schwering et al. 2009a). Therefore, intuitively the operation of computing a base ontology can thus be seen as a bi-directional search for analogy or generalisation into a base ontology together with the corresponding mappings. Providing efficient means for finding a number of suitable such candidate generalisations is essential to making the entire blending process computationally feasible. Consider the example of blending ‘house’ with ‘boat’ discussed in detail in

⁶ 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).

Section 4.1: even after fixing the base ontology itself, guessing the right mappings into the input ontologies means guessing within a space of approximately 1.4 Billion signature morphisms. Three promising candidates for finding generalisations are:

(1) **Ontology intersection:** Normann (2009) has studied the automatization of theory interpretation search for formalised mathematics, implemented as part of the Heterogeneous Tool Set (HETS, see below). Kutz and Normann (2009) applied these ideas to ontologies by using the ontologies' axiomatisations for finding their shared structure. Accidental naming of concept and role names is deliberately ignored and such names are treated as arbitrary symbols (i.e., any concept may be matched with any other). By computing mutual theory interpretations between the inputs, the method allows the computation of a base ontology as an *intersection* of the input ontologies together with corresponding theory morphisms. While this approach can be efficiently applied to ontologies with non-trivial axiomatisations, lightweight ontologies are less applicable, e.g., 'intersecting' a smaller taxonomy with a larger one clearly results in a huge number of possible taxonomy matches (Kutz and Normann 2009). In this case, the following techniques are more appropriate.

(2) **Structure-based ontology matching:** matching and alignment approaches are often restricted to find simple correspondences between atomic entities of the ontology vocabulary. In contrast, work such as Ritze et al. (2009); Walshe (2012) focuses on defining a number of *complex correspondence patterns* that can be used together with standard alignments in order to relate complex expressions between two input ontologies. For instance, the 'Class by Attribute Type Pattern' may be employed to claim the equivalence of the atomic concept `PositiveReviewedPaper` in ontology O_1 with the complex concept $\exists \text{hasEvaluation.Positive}$ of O_2 . Such an equivalence can be taken as an axiom of the base ontology; note, however, that it could typically not be found by intersecting the input ontologies. Giving such a library of design patterns may be seen as a variation of the idea of using image schemata.

(3) **Analogical Reasoning:** *Heuristic-driven theory projection* is a logic-based technique for analogical reasoning that can be employed for the task of computing a common generalisation of input theories. Schwering et al. (2009a) establish an analogical relation between a source theory and a target theory (both first-order) by computing a common generalisation (called 'structural description'). They implement this by using anti-unification (Plotkin 1970). A typical example is to find a generalisation (base ontology) formalising the structural commonalities between the Rutherford atomic model and a model of the solar system. This process may be assisted by a background knowledge base (in the ontological setting, a related domain or foundational ontology). Indeed, this idea has been further developed in Martinez et al. (2011).

3.2 *Selecting the Blendoids: Optimality Principles*

Having a common base ontology (computed or given) with appropriate base morphism, there is typically still a large number of possible blendoids whenever some kind of partiality is allowed. For example, even in the rather simple case of combining House and Boat, allowing for blendoids which only partially maintain structure (called *non-primary* blendoids in Goguen and Harrell (2010)), i.e., where any subset of the axioms may be propagated to the resulting blendoid, the number of possible blendoids is in the magnitude of 1000. Clearly, from an ontological viewpoint, the overwhelming majority of these candidates will be rather meaningless. A ranking therefore needs to be applied on the basis of specific ontological principles. In conceptual blending theory, a number of **optimality principles** are given in an informal and heuristic style (Fauconnier and Turner 2003). While they provide useful guidelines for evaluating natural language blends, they do not suggest a direct algorithmic implementation, as also analysed in Goguen and Harrell (2010) who in their prototypical implementation only covered certain structural, logical criteria. However, the importance of designing computational versions of optimality principles has been realised early on, and one such attempt may be found in the work of Pereira and Cardoso (2003), who proposed an implementation of the eight optimality principles presented in Fauconnier and Turner (1998) based on quantitative metrics for their more lightweight logical formalisation of blending. Such metrics, though, are not directly applicable to more expressive languages such as OWL or first-order logic.

Moreover, the standard blending theory of Fauconnier and Turner (2003) does not assign types, which might make sense in the case of linguistic blends where type information is often ignored. A typical example of a type mismatch in language is the operation of *personification*, e.g., turning a boat into an ‘inhabitant’ of the ‘boathouse’. However, in the case of blending in mathematics or ontology, this loss of information is often rather unacceptable: on the contrary, a fine-grained control of type or sort information may be of the utmost importance.

Optimality principles for ontological blending will be of two kinds:

(1) purely **structural/logical principles**: these will extend and refine the criteria as given in Goguen and Harrell (2010), namely *degree of commutativity* of the blend diagram, *type casting* (preservation of taxonomical structure), *degree of partiality* (of signature morphisms), and *degree of axiom preservation*. In the context of OWL, typing needs to be replaced with preservation of specific axioms encoding the taxonomy.

(2) **heuristic principles**: these include introducing preference orders on morphisms (an idea that Goguen (1999) labelled $3/2$ pushouts) reflecting their ‘quality’, e.g. measured in terms of degree of type violation; specific ontological principles, e.g. adherence to the OntoClean methodology (Guarino and Welty 2002) and ontological modelling principles, or general ontology evaluation techniques such as competency questions and fidelity requirements, as further discussed in Section 5.2.

4 Blending with the Distributed Ontology Language DOL

The distributed ontology language DOL is a formal language for specifying both ontologies, base diagrams, and their blends. DOL is a metalanguage in the sense that it enables the reuse of existing ontologies (written in some ontology language like OWL or Common Logic) as building blocks for new ontologies and, further, allows the specification of intended relationships between ontologies. One important feature of DOL is the ability to combine ontologies that are written in different languages without changing their semantics. DOL is going to be submitted as response to the Object Management Group's (OMG) Ontology, Model and Specification Integration and Interoperability (OntoIOp) Request For Proposal.⁷

In this section, we introduce DOL only informally. A formal specification of the language and its model-theoretic semantics can be found in Mossakowski et al. (2013, 2014).

For the purpose of ontology blending the following features of DOL are relevant:

- a **distributed ontology** consists of basic and structured ontologies and ontology interpretations. A basic ontology is an ontology written in some ontology language (e.g., OWL or Common Logic). A structured ontology builds on basic ontologies with the help of ontology translations, ontology unions, and symbol hiding.
- a **basic ontology** $\langle \Sigma, \Gamma \rangle$ written in some ontology language; Σ is a signature, Γ a set of Σ -sentences, with $\mathbf{Mod}(\langle \Sigma, \Gamma \rangle)$ containing all Σ -models satisfying Γ ;
- **ontology translation** (written O_1 **with** σ). A translation takes an ontology O_1 and a renaming function (technically, signature morphism) σ . The result of a translation is an ontology O_2 , which differs from the ontology O_1 only by substituting the symbols as specified by the renaming function.
- **ontology union** (written O_1 **and** O_2). The union of two ontologies O_1 and O_2 is a new ontology O_3 , which combines the axioms of both ontologies.
- **symbol hiding** (written O_1 **hide** s_1, \dots, s_n). A symbol hiding takes an ontology O_1 and a set of symbols s_1, \dots, s_n . The result of the hiding is a new ontology O_2 , which is the result of 'removing' the symbols s_1, \dots, s_n from the signature of ontology O_1 . Nevertheless, O_2 keeps all semantic constraints from O_1 .⁸
- **ontology interpretation** (written **interpretation** $INT_NAME : O_1$ to $O_2 = \sigma$). An ontology interpretation is a claim about the relationship between two ontologies O_1 and O_2 , giving some renaming function σ . It states that all the constraints that are the result of translating O_1 with σ logically follow from O_2 .
- **ontology alignment** (written **alignment** $ALIGN_NAME : O_1$ to $O_2 = c_1, \dots, c_n$), where the correspondences c_i relate a symbol in O_1 with one in O_2 , e.g. $s_1 = s_2$ or $s_1 < s_2$. Alignments can be seen as a relational variant of interpretations, with the major difference that no logical consequence is involved.

⁷ <http://www.omg.org/cgi-bin/doc?ad/2013-12-02>

⁸ By approximation, one could consider O_2 as the ontology that is the result of existentially quantifying s_1, \dots, s_n in O_1 .

An essential novelty introduced in DOL is that a user can specify the ontological base diagram as a DOL theory, from which the colimit and other blendoids can then be computed.⁹ This is a crucial task, as the computed colimit ontology depends on the dependencies between symbols that are stored in the diagram. Ontohub, a DOL-enabled repository discussed further in Section 5, is able to use the specification of a base diagram to automatically generate the colimit ontology. In the next section, we illustrate the specification of base diagrams in DOL and the computation of the resulting blendoids by blending house and boat to houseboat and boathouse.

4.1 The classic House+Boat Blend

The main inputs for the blendings consist of two ontologies, one for HOUSE and the other for BOAT. We adapt them from Goguen and Harrell (2010) but give a stronger axiomatisation to make them more realistic and ontologically sound. Fig. 4 shows the ontology for HOUSE in OWL Manchester Syntax. The ontology is a fragment introducing several concepts necessary for understanding the basic meaning of the term ‘house’, including that it is an artefact that has the capability of serving as a residence for people and is generally located on a plot of land. The precise formalisation is not criterial at this point; any adequate ontological description of ‘house’ would, however, needs to provide similar distinctions.¹⁰

As discussed above, finding candidate base ontologies and base morphisms is a non-trivial task. For the purpose of this example, we create them manually. The purpose of the example is to show how the DOL specifications naturally allow us to express these kinds of ‘re-mappings’ of relations and entities that are required when considering blends in general. The base ontologies used for the two blends discussed here are both quite simple, they mostly introduce shared concepts and contain only weak axiomatisations. The second base ontology only differs from the first by replacing the class `Agent` by `Person` and two additional classes, namely `Object` and `Site`.

```
ontology base1 =
  Class: Artifact [...] Class: Agent
end

ontology base2 =
  Class: Artifact [...] Class: Person
  Class: Object      Class: Site
end
```

⁹ While OBJ3 already provides the possibility to write down theory morphisms, only DOL provides means to collect them into a formally defined diagram; see the **graph** construct below.

¹⁰ In the examples, note that concepts such as ‘ArtifactThatExecutesResidenceFunction’ are auxiliary symbols that are needed because of limitation of the Manchester Syntax being used, which does not allow the use of complex concepts on the left-hand side of subsumption statements. The ontology for BOAT is axiomatised similarly, it can be found at <http://www.ontohub.org/conceptportal>.

```

Class: Artifact
Class: Capability
ObjectProperty: has_function
    Range: Capability
ObjectProperty: executes
    Range: Capability
ObjectProperty: is_located_on
Class: Person
Class: Plot
ObjectProperty: is_inhabited_by
    Domain: House
    Range: Person
Class: ServeAsResidence
    SubClassOf: Capability
Class: ArtifactThatExecutesResidenceFunction
    EquivalentTo: Artifact that executes
        some ServeAsResidence
    SubClassOf: is_inhabited_by some Person
Class: House
    SubClassOf: Artifact
        that is_located_on some Plot
        and has_function some
            ServeAsResidence

```

Fig. 4 Ontology House

The blending of boat and house to houseboat is achieved by turning the boat into a habitat and moving the house from a plot of land to a body of water. This can be represented by two interpretations `boat_habitable` and `house_floating`.

```

interpretation boat_habitable : base2 to Boat =
    Object  $\mapsto$  Boat,
    Site  $\mapsto$  BodyOfWater

```

```

interpretation house_floating : base2 to House =
    Object  $\mapsto$  House,
    Site  $\mapsto$  Plot

```

The base ontologies and the interpretations above provide the necessary ingredients for a blending of BOAT and HOUSE to HOUSEBOAT. The syntax of diagrams is

```

graph  $D = O_1, \dots, O_m, M_1, \dots, M_n, A_1, \dots, A_k, D_1, \dots, D_l$ 

```

where the O_i are ontologies, the M_i are morphisms, the A_i are alignments and the D_i are existing diagrams. The syntax of combinations is

```

combine  $O_1, \dots, O_m, M_1, \dots, M_n, A_1, \dots, A_k, D_1, \dots, D_l$ 

```

with the ingredients as above. The simplest (and still fully general) form is just

```

combine  $D$ 

```

where D is a diagram. The semantics of combinations is the colimit of the generated diagram. A colimit involves both pasting together (technically: disjoint union) and identification of shared parts (technically: a quotient).

In our example, houseboat can be defined by the colimit based on the interpretations. To make the result easier to read, some of the classes are renamed:

```
ontology house_boat =
  combine boat_habitable, house_floating
  with Object ↦ HouseBoat, Site ↦ BodyOfWater
```

This captures formally the informal description of the house+boat blend as often given in examples of blending diagrams. Our specification then allows us to go further and derive both consequences of this and other blends. Here Ontohub is able to compute the colimit, which combines both the BOAT and HOUSE ontologies along the morphism. The colimit inherits the axioms of the input ontologies and the base with appropriate identifications of symbols. Here we just show the declaration of the blended class Houseboat:

```
Class: HouseBoat
SubClassOf: Artifact
  and has_function some MeansOfTransportation
  and has_function some Floating
  and is_navigated_by some Agent
SubClassOf: Artifact
  and is_located_on some BodyOfWater
  and has_function some ServeAsResidence
```

In the case of blending BOAT and HOUSE to BOATHOUSE, the crucial part in this blend is to view a boat as a kind of “person” that lives in a house. The two ontologies House and Boat presented above can be blended by selecting a base, which here provides (among others) a class Agent, and two interpretations, mapping Agent to Boat and Person, respectively. Therefore, the second base ontology only differs from the first by replacing the class Agent by Person and two additional classes, namely Object and Site.

```
ontology base1 =
  Class: Artifact [...] Class: Agent
end
```

In this way, we let a boat play the role of a person (that inhabits a house).¹¹

```
interpretation boat_personification :
  base1 to Boat =
  Agent ↦ Boat
```

```
interpretation house_import :
  base1 to House =
  Agent ↦ Person
```

¹¹ Compared to Goguen and Harrell (2010), the advantage of our formulation is that no projections (“retracts”) from a superset to a subset are needed. Instead, we can carefully select which parts of the theory of houses and their inhabitants are instantiated with boats.


```

ontology boat_house =
  combine boat_personification, house_import
  with Agent  $\mapsto$  Boat, House  $\mapsto$  BoatHouse

```

As before, Ontohub is able to compute the colimit. As above, we present here only the relevant declarations of the blended concept.

```

Class: BoatHouse
  SubClassOf: Artifact
    and is_located_on some Plot
    and has_function some ServeAsResidence
Class: ArtifactThatExecutesResidenceFunction
  EquivalentTo: Artifact
    and executes some ServeAsResidence
  SubClassOf: is_inhabited_by some Boat

```

Figure 5 shows the representation of the ontologies and their relations in Ontohub.

Of course, the possibilities for blending the two concepts do not stop here. For example, we could map the agent in the base ontology to person in the boat ontology. This can be achieved by first defining an additional interpretation and by blending all three interpretations.

```

interpretation boat_import :
  base1 to Boat =
    Agent  $\mapsto$  Person

ontology boat_house =
  combine boat_personification, house_import, boat_import
  with Agent  $\mapsto$  Boat, House  $\mapsto$  BoatHouse

```

The resulting blendoid is consistent, but it contains some strange consequences. For example, in the blendoid boats are driven by boats. However, if we are interested both in hosting boats and a hub for autonomous vehicles, this would count as an interesting result. In general, whether such more creative aspects of blendoids are desirable or not will depend on the context of the blending. We will address this issue in the section on evaluation below. It should be noted, however, that an ontologically cleaner axiomatisation of the input spaces makes blending in fact easier — this is because it reveals more clearly the type structure of the inputs, whose modification can then be more elegantly controlled via the base morphisms.

4.2 Variations: Blends of Blends and Partiality

We have discussed a more sophisticated version of the classic HOUSE + BOAT blending in order to illustrate some of the fine detail in the workings of formalised blending in the Goguen tradition, here based on the DOL language. However, the basic blending diagram only covers the most basic situation, that of an ‘atomic blend’ using basic concepts and one base space. The real power of blending, however, is only unleashed when blends are iterated and when partiality is allowed.

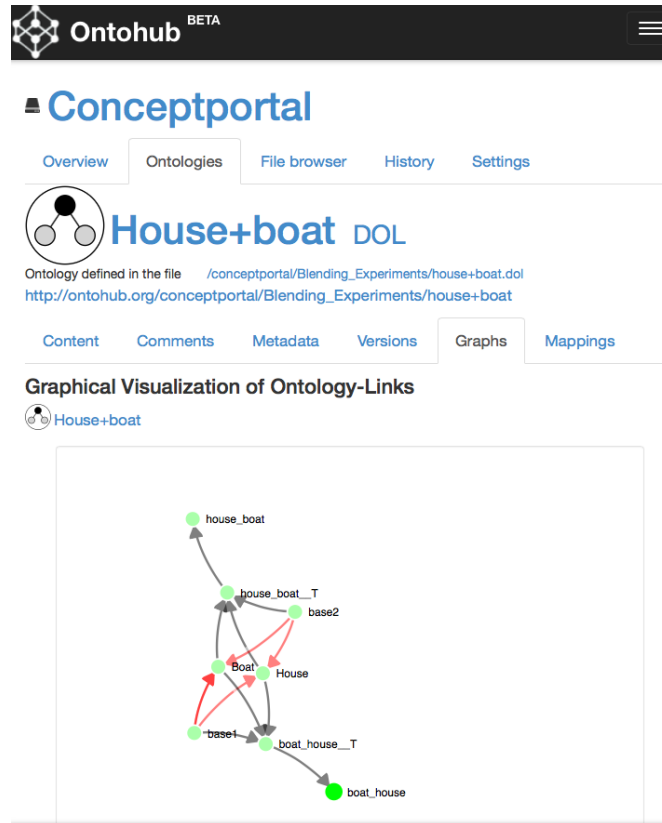


Fig. 5 Blendoid representation and colimit computation via HETS/Ontohub: the screenshot of Ontohub shows Conceptportal, which contains the double-blend of house and boat. In the graph the dots represent the ontologies: the input spaces (House, Boat), the two bases, and the computed blendoids (boat_house, house_boat). (The ontologies boat_house_T and house_boat_T are generated by Ontohub as an intermediate step before the terms in the signature are renamed.) The arrows denote the relationships between the ontologies (interpretations, blending, and renaming).

Lakoff and Núñez (2000) give a detailed and powerful analysis of this in the field of conceptual mathematics. A basic claim they make is that the most sophisticated mathematical concepts have been created, over time, through a tower of blended concepts, generating more and more abstract notions. A basic case is that of arithmetic, where several metaphors, image schemas, and analogies are successively blended into modern number systems such as rationals, reals, or complex numbers, including ‘arithmetic as object collection’, ‘object construction’, the ‘measuring stick metaphor’ and ‘arithmetic as motion along a path’ (see Lakoff and Núñez (2000) and Guhe et al. (2011) for further details and Fleuriot et al. (2014) for a conceptual blend of the complex numbers along these lines). A detailed formal reconstruction of such iterated blends is a challenging task, both conceptually and on

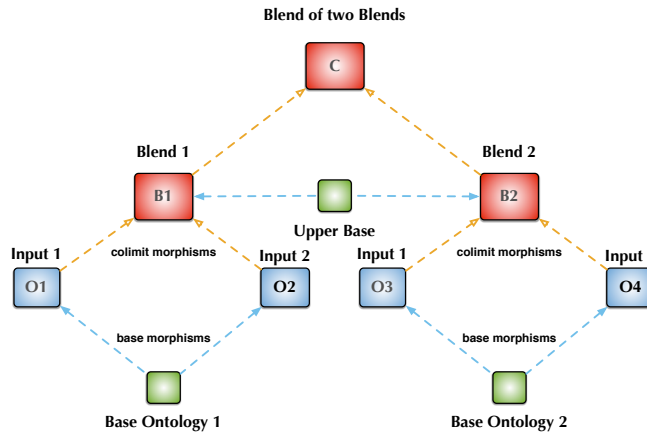


Fig. 6 Blending two basic blends into a third.

a technical level. Figure 6 shows the basic diagrammatical structure of such iterated blends.

Iteration of blends, however, is not the only variation of the basic blendoid structure. Figure 7 shows two triple blends, both have three input spaces, but the one on the left has one base, the one on the right has two base spaces. For instance, we might have 3 inputs that are simultaneously aligned with a basic image schema in the base (left), or we have three ontologies that pairwise interpret different metaphors, e.g. ‘arithmetic as object collection’ and ‘arithmetic as motion along a path’.




Fig. 7 Blending three input spaces using one respectively two base ontologies.

Note that on a purely technical level, such complex diagrams can always be reduced to a succession of squares, possibly by duplicating some nodes or adding trivial ones¹²—however, such a reduction loses the direct connection between the diagrammatic representation and the cognitive-conceptual processes that are being formalised here. In a similar vein, Def. 1 introducing the notion of an ontological

¹² A well-known theorem of category theory states that every finite colimit can be expressed by pushouts and initial objects.

base diagram in Section 3 easily generalises to the case of partial base morphisms, i.e. where only parts of the signature of an ontology are mapped. Such partial morphisms can be coded as spans of two (total) theory morphisms $B_i \leftarrow \text{dom}(\mu_{ij}) \rightarrow I_j$, where the first morphism is the embedding of the domain (actually, the larger $\text{dom}(\mu_{ij})$ is, the more defined is the partial morphism), and the second action represents the action of the partial morphism.¹³ Similarly, arbitrary relations can be coded as spans $B_i \leftarrow R \rightarrow I_j$. Here, $R \subseteq B_i \times I_j$ is a relation, and the arrows are the projections to the first and second component. However, such complexities can be hidden from a user by allowing partial morphisms to be used directly in the specification of a blending diagram, and by letting a tool handle the simulation through total morphisms as discussed above.

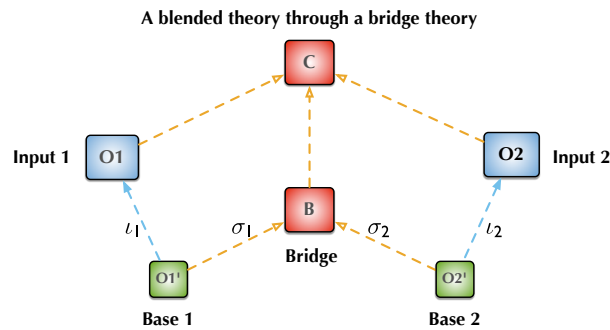


Fig. 8 Blending two input spaces through two bases and a bridge theory, deviating from the Goguen construction.

Finally, a more severe deviation from the basic blending diagram is shown in Fig. 8. Here, we interpret Base1 into Input1, Base2 into Input2, and connect the two bases by a bridge theory. For example, the bridge might introduce a higher-level connection between two image schemas which is then used to create the blended theory. An example of this might be where we have image schemas involved with time and with space and combine these first for the definition of a theory in which time and space are linked (as in our navigation examples above or even in the linking between time and space discussed for comics and visual narrative).

Technically, such diagrams are closely related to alignments (Zimmermann et al. 2006; Codescu et al. 2014), and to distributed modelling languages (Kutz et al. 2004). Concerning the formalisation of conceptual blending, these techniques and diagram patterns will be of particular importance to tackle a computational treatment and formal representation of so-called *generic integration templates* (GIT), i.e. the idea of studying general blending templates, first introduced and discussed in detail by Pagán Cánovas (2010), with more examples to be found in Turner (2014).

¹³ In this case, the base diagram becomes a bit more complex; in particular, there are minimal nodes $\text{dom}(\mu_{ij})$ which have only auxiliary purposes and do not belong to the base.

5 Blending in the Hub

In this section we will discuss the computational and representational support for formalised blending via the Ontohub.org platform as well as the evaluation problem.

5.1 Representation and Computation

To begin, combinations (or, alternatively, the underlying colimits) can be computed directly by the web platform Ontohub. Ontohub is a repository engine for managing distributed heterogeneous ontologies. Ontohub supports a wide range of formal logical and ontology languages and allows for complex inter-theory (concept) mappings and relationships with formal semantics, as well as ontology alignments and blending. Ontohub understands various input languages, among them OWL and DOL.

We describe the basic design and features of Ontohub in general, and outline the extended feature-set that we pursue to add to Ontohub for conceptportal.org — a specialised repository for blending experiments within the distributed Ontohub architecture.

The back-end of Ontohub is the Heterogeneous Tool Set HETS, which is used by Ontohub for parsing, static analysis and proof management of ontologies. HETS can also compute colimits of both OWL and first-order logic diagrams and even approximations of colimits in the case where the input ontologies live in different ontology languages (Codescu and Mossakowski 2008).

Computation of colimits in HETS is based on HETS' general colimit algorithm for diagrams of sets and functions (note that signatures in most cases are structured sets, and signature morphisms structure preserving functions.) Such a colimit of sets and functions is computed by taking the disjoint union of all sets, and quotienting it by the equivalence relation generated by the diagram, which more precisely is obtained by the rule that given any element x of an involved set, any images of x under the involved functions are identified. The quotient is computed by selecting a representative of each equivalence class.

A difficulty that arises is that we have to make a choice of these representatives, and therefore of names for the symbols in the colimit, since a symbol is often not identically mapped in the base diagram of the blendoid. The convention in HETS is that, in case of ambiguity, from among all symbols of the equivalence class, that name of the symbol is chosen which is the most frequently occurring one. In any case, the user has control over the namespace because the symbols in the colimit can later be renamed. We can see this for our boathouse example above, where `Agent` appears most often in the diagram and therefore the symbol has been correspondingly renamed.

5.2 Evaluating the Blending Space

Optimality principles (see Section 3.2), in particular structural ones, can be used to rank candidate blendoids on-the-fly during the ontology blending process. However, even if they improve on existing logical and heuristic methods, optimality principles will only narrow down the potential candidates and not tell us whether the result is a ‘successful’ blend of the ontologies. For example, assume that we had optimality principles that would show that from the roughly 1000 candidate blendoids of *House* and *Boat* that Goguen computed, only two candidates \mathfrak{B}_{hb} and \mathfrak{B}_{bh} are optimal. Is either \mathfrak{B}_{hb} or \mathfrak{B}_{bh} any good? And, if so, which of them should we use? To answer these questions, it seems natural to apply ontology evaluation techniques.

Ontologies are human-intelligible and machine-interpretable representations of some portions and aspects of a domain that are used as part of information systems. To be more specific, an ontology is a logical theory written in some knowledge representation language, which is associated with some intended interpretation. The intended interpretation is partially captured in the choice of symbols and natural language text (often in the form of annotations or comments). The evaluation of an ontology covers both the logical theory and the intended interpretation, their relationship to each other, and how they relate to the requirements that are derived from the intended use within a given information system. Therefore, ontology evaluation is concerned not only with formal properties of logical theories (e.g., logical consistency), but, among other aspects, with the *fidelity* of an ontology; that is whether the formal theory accurately represents the intended domain (Neuhaus et al. 2013). For example, if \mathfrak{B}_{hb} is an excellent representation of the concept *houseboat*, then \mathfrak{B}_{hb} provides a poor representation of the concept *boathouses*. Thus, any evaluation of the blend \mathfrak{B}_{hb} depends on what domain \mathfrak{B}_{hb} is intended to represent.

Given these considerations, \mathfrak{B}_{hb} and \mathfrak{B}_{bh} are not ontologies, they are logical theories that are the result of the blending of two logical theories that are part of ontologies. This is illustrated by the following thought-experiment: let’s assume the theory \mathfrak{B}_{hb} captures the concept *houseboat* very well, and that \mathfrak{B}_{hb} is not the result of some automatic blending process, but was intentionally developed by an ontology engineer. In case that the ontology engineer intended to develop an ontology of houseboats \mathfrak{B}_{hb} , he would have done very well. However, if the engineer intended to develop an ontology of boathouses, then \mathfrak{B}_{hb} would be a poor outcome. In other words, the ontology consisting of \mathfrak{B}_{hb} and the intention *houseboat* would have high fidelity, but the ontology consisting of \mathfrak{B}_{hb} and the intention *boathouse* would have low fidelity. Thus, the evaluation of the theory \mathfrak{B}_{hb} is dependent on the domain it is supposed to represent.

The lesson from this thought experiment is that the evaluation of the results of ontology blending is dependent on the intended goal and, more generally, on the requirements that one expects the outcome of the blending process to meet. One way to capture these requirements is similar to competency questions, which are widely used in ontology engineering (Grüninger and Fox 1995). Competency questions are usually initially captured in natural language; they specify examples for questions that an ontology needs to be able to answer in a given scenario. By formalising the

competency questions one can use automatic theorem provers to evaluate whether the ontology meets the intended interpretation.

The requirements that are used to select between the different blends fall, roughly, into two categories: *ontological constraints* and *consequence requirements*. Ontological constraints prevent the blends from becoming ‘too creative’ by narrowing the space for conceptual blending. E.g., it may be desirable to ensure that the `is_inhabited_by` relationship is asymmetric and that `is_navigated_by` is irreflexive. To achieve that any blendoid can be checked for logical consistency with the following ontology:

```
ontology OntologicalConstraints =
  ObjectProperty: is_inhabited_by
    Characteristics: Asymmetric
  ObjectProperty: is_navigated_by
    Characteristics: Irreflexive
```

Given these requirements, any blendoid that involves a house that lives in itself, or any boat navigated by itself (see the blendoid `boat_house1` above) would be discarded.

Consequence requirements specify the kind of characteristics the blendoid is supposed to have. E.g., assume the purpose of the conceptual blending is to find alternative housing arrangements, because high land prices make newly built houses unaffordable. In this case, the requirement could be ‘a residence that is not located on a plot of land’, which can be expressed in OWL as follows:

```
ontology ConsequenceRequirements =
[...]
  Class PlotFreeResidence
    EquivalentTo: Residence
      and (is_located_on only (not (Plot)))
```

For the evaluation of a blendoid against requirements (both ontological constraints and consequence requirements) it is often not sufficient to just consider the information that is contained in the blendoid itself. Some background knowledge usually needs to be added in order to evaluate a blendoid.

Background knowledge plays another crucial role in the blending process, which we have not addressed in this paper so far. The basic blending diagram in Figure 3 presents a static view, which describes how two input spaces, a base, and two interpretations give rise to a blendoid. However, any system that attempts to automate conceptual blending will need to perform not one but many blends in order to get a decent result. In this process, the background knowledge and the evaluation of previous blending results can be utilised in the selection of candidate bases and interpretations. Further, the violation of ontological constraints may be symptom of an attempt to blend input spaces that are too rich. In these cases, the result of the evaluation can be used to guide heuristics, which remove information from the original input spaces that may have caused the violation of the ontological constraints. The result is a new, weakened input space, which may be easier to blend. In short, following a proposal by Marco Schorlemmer discussed in detail in Neuhaus et al. (2014), we envision an approach where background knowledge and evaluation are driving an iterative blending process, as illustrated in Figure 9.

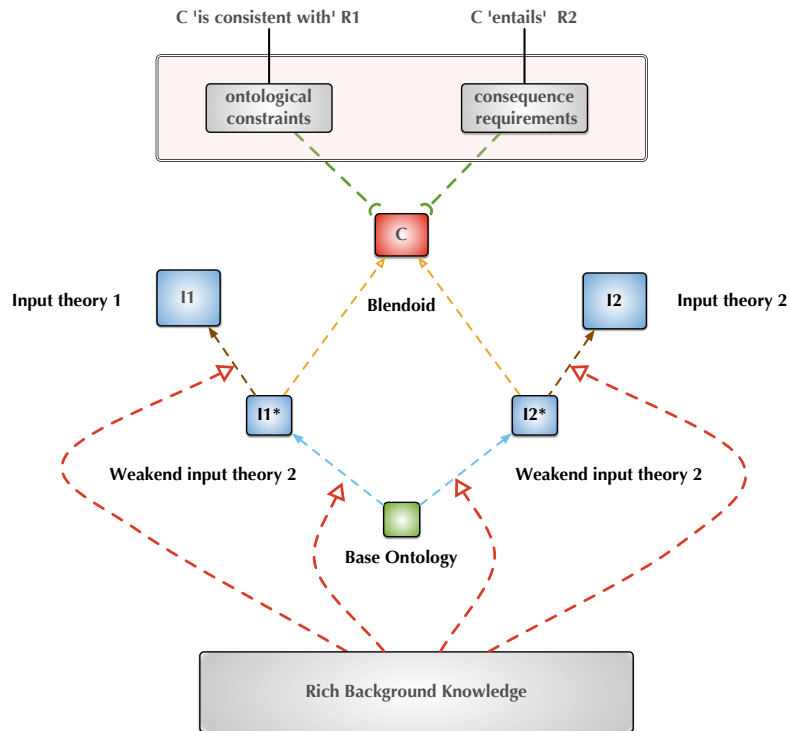


Fig. 9 The core Schorlemmer model for computational blending enriched with evaluation and background layers

Ontohub allows to use ontological constraints and consequence requirements to evaluate blended concepts automatically. The requirements are managed as DOL files, which allow to express that a given blendoid (together with some background knowledge) is logically consistent with a set of ontological constraints or that it entails some consequence requirements. The requirements themselves may be stored as regular ontology files (e.g., in OWL Manchester syntax). Ontohub executes the DOL files with the help of integrated automatic theorem provers, and is able to detect whether a blendoid meets the specified requirements. At this time, the evaluation of blendoids for ontological constraints and consequence requirements depends on the use of DOL files. We are planning to integrate this functionality into the GUI of Ontohub to make it more convenient for the user.

6 E pluribus unum: The future of the melting pot

Our work in this paper follows a research line in which blending processes are primarily controlled through mappings and their properties (Gentner 1983; Forbus et al. 1989; Veale 1997; Pereira 2007). By introducing blending techniques to ontology languages, we have provided a method which allows us to combine two or more thematically different ontologies into a newly created ontology, the blendoid, describing a novel concept or domain. The blendoid creatively mixes information from both input ontologies on the basis of structural commonalities of the inputs and selective combination of their axiomatisations.

We have moreover illustrated that the Ontohub/HETS tool ecosystem and the DOL language provide an excellent starting point for developing the theory and practice of ontology blending further (Mossakowski et al. 2013). They (1) support various ontology languages and their heterogeneous integration (Kutz et al. 2008b); (2) allow the specification of theory interpretations and other morphisms between ontologies (Kutz et al. 2010); (3) support the computation of colimits as well as the approximation of colimits in the heterogeneous case (Codescu and Mossakowski 2008); (4) provide (first) solutions for automatically computing a base ontology through ontology intersection (Kutz and Normann 2009) and blendoid evaluation using requirements (Kutz et al. 2014; Neuhaus et al. 2014).

In particular, we have shown that the blending of ontologies can be declaratively encoded in a DOL theory representing the respective blending diagram—here, employing the homogeneous fragment of DOL just using OWL ontologies. Blendoid ontologies, as well as their components, i.e. input and base ontologies, can be stored, formally related, and checked for consistency within Conceptportal, a repository node within Ontohub dedicated to blending experiments carried out in the European FP7 Project COINVENT (Schorlemmer et al. 2014). Ontohub moreover gives access to thousands of ontologies from a large number of different scientific and common sense domains, searchable via rich metadata annotation, logics used, formality level, and other dimensions, to provide not only a rich pool of ontologies for blending experiments, but also for the evaluation of newly created concepts.

Of course, constructing a homogeneous blendoid from a basic blending diagram is one of the simplest cases of conceptual blending. As discussed in Section 4.2, on a technical level a blendoid is just like an alignment diagram, except that instead of dealing with synonymy and homonymy relations, and just signature in the base, in the blendoid case we are dealing with selectively merging axioms. Following this intuition, a whole range of more complex alignment and theory combination techniques can be combined with the basic blending ideas of Goguen: this includes constructions such as W-alignments (Zimmermann et al. 2006; Kutz et al. 2008c; Codescu et al. 2014), and connections of theories following the \mathcal{L} -connection/DDL paradigm (Kutz et al. 2004; Borgida and Serafini 2003; Nalon and Kutz 2014).

The next important milestone for computational conceptual blending will be to make the step from a *reconstructive* approach, where conceptual blending is illustrated by blending one concept (e.g., houseboat) with the help of some carefully selected input spaces (e.g., a house and a boat) and a hand-crafted base ontology,

to a system that *autonomously* selects two (or more) ontologies from a repository in Ontohub and attempts to blend them in a way that meets some given requirements. In Neuhaus et al. (2014), we have described the first steps towards designing a computational architecture that performs conceptual blending autonomously and self-evaluates its own creations.

Within the extensive literature on conceptual blending, only few attempts have been made at a (more or less) complete automation of the blending process; notable exceptions include Goguen and Harrell (2010), Pereira (2007), Li et al. (2012), and Veale and O’Donoghue (2001); Veale (2012). To make concept invention via ontological blending more feasible in practice from within Ontohub, a number of further plugins into the architecture and refinements are planned covering in particular the automatic creation of base ontologies together with their mappings, the implementation of filtering blendoids by structural optimality principles and preference orders on morphisms, as well as the addition of more ontologically motivated evaluation techniques as discussed above.

Acknowledgements. 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. Work on this paper was moreover supported by the DFG-funded collaborative research centre SFB/TR 8 ‘Spatial Cognition’ of the Universities of Bremen and Freiburg.

We thank the anonymous referees as well as Mihai Codescu for detailed feedback on this chapter.

References

- Adámek, J., Herrlich, H., and Strecker, G. (1990). *Abstract and Concrete Categories*. Wiley, New York.
- Bateman, J. A. (2010). Language and Space: a two-level semantic approach based on principles of ontological engineering. *International Journal of Speech Technology*, 13(1):29–48.
- Bateman, J. A. (2011). The decomposability of semiotic modes. In O’Halloran, K. L. and Smith, B. A., editors, *Multimodal Studies: Multiple Approaches and Domains*, Routledge Studies in Multimodality, pages 17–38. Routledge, London.
- Bateman, J. A. and Veloso, F. O. D. (2013). The Semiotic Resources of Comics in Movie Adaptation: Ang Lee’s *Hulk* (2003) as a case study. *Studies in Comics*, 4(1):137–159.
- Bhatt, M., Hois, J., and Kutz, O. (2012). Ontological modelling of form and function in architectural design. *Applied Ontology*, 7(3):233–267.
- Black, M. (1979). More about metaphor. In Ortony, A., editor, *Metaphor and Thought*, pages 19–43. Cambridge University Press, Cambridge.

- Borgida, A. and Serafini, L. (2003). Distributed Description Logics: Assimilating Information from Peer Sources. *Journal of Data Semantics*, 1:153–184.
- Codescu, M. and Mossakowski, T. (2008). Heterogeneous colimits. In Boulanger, F., Gaston, C., and Schobbens, P.-Y., editors, *MoVaH'08 Workshop on Modeling, Validation and Heterogeneity*, pages 131–140. IEEE press.
- Codescu, M., Mossakowski, T., and Kutz, O. (2014). A Categorical Approach to Ontology Alignment. In *Proc. of the 9th International Workshop on Ontology Matching (OM-2014)*, ISWC-2014, Riva del Garda, Trentino, Italy. CEUR-WS.
- Eisner, W. (1992). *Comics and Sequential Art*. Kitchen Sink Press Inc., Princeton, WI.
- Euzenat, J. and Shvaiko, P. (2007). *Ontology Matching*. Springer.
- Fauconnier, G. (1997). *Mappings in Thought and Language*. Cambridge University Press, Cambridge.
- 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.
- Fleuriot, J., Maclean, E., Smaill, A., and Winterstein, D. (2014). Reinventing the Complex Numbers. In Besold, T., Kühnberger, K.-U., Schorlemmer, M., and Smaill, A., editors, *Computational Creativity, Concept Invention, and General Intelligence, Proc. of the 3rd Int. Workshop C3GI@ECAI-14*, volume 1-2014, Prague, Czech Republic. Publications of the Institute of Cognitive Science, Osnabrück.
- Forbus, K., Falkenhainer, B., and Gentner, D. (1989). The structure-mapping engine. *Artificial Intelligence*, 41:1–63.
- Forceville, C. J. and Urios-Aparisi, E., editors (2009). *Multimodal Metaphor*. Mouton de Gruyter, Berlin/New York.
- Frank, A. and Kuhn, W. (1999). A specification language for interoperable GIS. In Goodchild, M., Egenhofer, M., Fegeas, R., and Kottmann, C., editors, *Interoperating geographic information systems*, pages 123–132. Kluwer, Norwell, MA.
- Gärdenfors, P. (2000). *Conceptual Spaces - The Geometry of Thought*. Bradford Books. MIT Press.
- Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170.
- Goguen, J. A. (1999). An Introduction to Algebraic Semiotics, with Applications to User Interface Design. In *Computation for Metaphors, Analogy and Agents*, number 1562 in LNCS, pages 242–291. Springer.
- Goguen, J. A. and Burstall, R. M. (1992). Institutions: Abstract Model Theory for Specification and Programming. *Journal of the Association for Computing Machinery*, 39(1):95–146. Predecessor in: LNCS 164, 221–256, 1984.
- Goguen, J. A. and Harrell, D. F. (2010). Style: A Computational and Conceptual Blending-Based Approach. In Argamon, S. and Dubnov, S., editors, *The Structure of Style: Algorithmic Approaches to Understanding Manner and Meaning*, pages 147–170. Springer, Berlin.

- Goguen, J. A. and Malcolm, G. (1996). *Algebraic Semantics of Imperative Programs*. MIT.
- Gruber, T. (1995). Toward principles for the design of ontologies used for knowledge sharing. *International Journal of Human-Computer Studies*, 43(5/6):907–928.
- Grüninger, M. and Fox, M. S. (1995). The role of competency questions in enterprise engineering. In *Benchmarking—Theory and Practice*, pages 22–31. Springer.
- Guarino, N. (1994). The Ontological Level. In Casati, R., Smith, B., and White, G., editors, *Philosophy and the Cognitive Sciences*, pages 443–456. Hölder-Pichler-Tempsky, Vienna.
- Guarino, N. and Welty, C. (2002). Evaluating ontological decisions with OntoClean. *Commun. ACM*, 45(2):61–65.
- Guhe, M., Pease, A., Smaill, A., Martínez, M., Schmidt, M., Gust, H., Kühnberger, K.-U., and Krumnack, U. (2011). A computational account of conceptual blending in basic mathematics. *Cognitive Systems Research*, 12(3–4):249–265.
- Gust, H., Kühnberger, K.-U., and Schmid, U. (2003). Metaphors and anti-unification. In *Proc. Twenty-First Workshop on Language Technology: Algebraic Methods in Language Processing, Verona, Italy*, pages 111–123.
- Hart, C. (2008). Critical discourse analysis and metaphor: toward a theoretical framework. *Critical Discourse Studies*, 5(2):91–106.
- Hois, J., Kutz, O., Mossakowski, T., and Bateman, J. (2010). Towards Ontological Blending. In *Proc. of the The 14th International Conference on Artificial Intelligence: Methodology, Systems, Applications (AIMSA-2010)*, Varna, Bulgaria, September 8th–10th.
- Jaszczolt, K. M. (2003). On Translating ‘What Is Said’: *Tertium Comparationis* in Contrastive Semantics and Pragmatics. In *Meaning Through Language Contrast Vol. 2*, pages 441–462. J. Benjamins.
- Johnson, M. (1987). *The Body in the Mind. The Bodily Basis of Meaning, Imagination, and Reasoning*. The University of Chicago Press.
- Kövecses, Z. (2010). *Metaphor: A Practical Introduction*. Oxford University Press, Oxford, 2 edition.
- Kuhn, W. (2002). Modeling the Semantics of Geographic Categories through Conceptual Integration. In *Proc. of GIScience 2002*, pages 108–118. Springer.
- Kuhn, W. (2003). Semantic Reference Systems. *International Journal of Geographical Information Science*, 17(5):405–409.
- Kutz, O., Lücke, D., and Mossakowski, T. (2008a). Heterogeneously Structured Ontologies—Integration, Connection, and Refinement. In *Proc. KROW 2008*, volume 90 of *CRPIT*, pages 41–50. ACS.
- Kutz, O., Lücke, D., Mossakowski, T., and Normann, I. (2008b). The OWL in the CASL—Designing Ontologies Across Logics. In *Proc. of OWLED-08*, volume 432. CEUR.
- Kutz, O., Lutz, C., Wolter, F., and Zakharyashev, M. (2004). \mathcal{E} -Connections of Abstract Description Systems. *Artificial Intelligence*, 156(1):1–73.

- Kutz, O., Mossakowski, T., and Codescu, M. (2008c). Shapes of Alignments: Construction, Combination, and Computation. In Sattler, U. and Tamilin, A., editors, *Proc of the 1st Workshop on Ontologies: Reasoning and Modularity (WORM-08)*, ESWC, Tenerife, Spain. CEUR-WS, Vol-348.
- 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., editors, *Computational Creativity, Concept Invention, and General Intelligence, Proc. of the 1st Int. Workshop C3GI@ECAI*, volume 01-2012, Montpellier, France. Publications of the Institute of Cognitive Science, Osnabrück.
- 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?’.
- Kutz, O., Neuhaus, F., Mossakowski, T., and Codescu, M. (2014). Blending in the Hub—Towards a collaborative concept invention platform. In *Proc. of the 5th International Conference on Computational Creativity*, Ljubljana, Slovenia.
- Kutz, O. and Normann, I. (2009). Context Discovery via Theory Interpretation. In *Workshop on Automated Reasoning about Context and Ontology Evolution, ARCOE-09 (IJCAI-09)*.
- Lakoff, G. and Núñez, R. E. (2000). *Where Mathematics Comes From*. Basic Books.
- Li, B., Zook, A., Davis, N., and Riedl, M. O. (2012). Goal-Driven Conceptual Blending: A Computational Approach for Creativity. In *Proc. of the 2012 International Conference on Computational Creativity*, Dublin, Ireland.
- Mamakos, C., Stefaneas, P., Dimarogkona, M., and Ireson-Paine, J. (2014). Polytopos Project: Experiments in Blending. In Besold, T., Kühnberger, K.-U., Schorlemmer, M., and Smaill, A., editors, *Computational Creativity, Concept Invention, and General Intelligence, Proc. of the 3rd Int. Workshop C3GI@ECAI-14*, volume 1-2014, Prague, Czech Republic. Publications of the Institute of Cognitive Science, Osnabrück.
- Martinez, M., Besold, T. R., Abdel-Fattah, A., Kühnberger, K.-U., Gust, H., Schmidt, M., and Krumnack, U. (2011). Towards a Domain-Independent Computational Framework for Theory Blending. In *Proc. of the AAAI Fall 2011 Symposium on Advances in Cognitive Systems*.
- McCloud, S. (1994). *Understanding comics: the invisible art*. HarperPerennial, New York.
- Miller, A. I. (2000). Metaphor and scientific creativity. In Hallyn, F., editor, *Metaphor and Analogy in the Sciences*, pages 147–164. Kluwer Academic Publishers, Dordrecht.
- 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., editors, *Proceedings of the 7th International Workshop on Modular Ontologies (WoMO-13)*, volume 1081. CEUR-WS.
- Mossakowski, T., Kutz, O., Neuhaus, F., Codescu, M., Lange, C., Gruninger, M., and Keet, M. (2014). The distributed ontology, modeling and specification language. Draft answer to the OMG RFP “OntoIOP”.

- Nalon, C. and Kutz, O. (2014). Towards Resolution-based Reasoning for Connected Logics. *Electronic Notes in Theoretical Computer Science*, 305:85–102. Post-proceedings of the 8th Workshop on Logical and Semantic Frameworks (LSFA).
- Neuhaus, F., Kutz, O., Codescu, M., and Mossakowski, T. (2014). Fabricating Monsters is Hard: Towards the Automation of Conceptual Blending. In Besold, T., Kühnberger, K.-U., Schorlemmer, M., and Smaill, A., editors, *Computational Creativity, Concept Invention, and General Intelligence, Proc. of the 3rd Int. Workshop C3GI@ECAI-14*, volume 1-2014, Prague, Czech Republic. Publications of the Institute of Cognitive Science, Osnabrück.
- Neuhaus, F., Vizedom, A., Baclawski, K., Bennett, M., Dean, M., Denny, M., Grüninger, M., Hashemi, A., Longstreth, T., Obrst, L., et al. (2013). Towards ontology evaluation across the life cycle: The Ontology Summit 2013. *Applied Ontology*, 8(3):179–194.
- Normann, I. (2009). *Automated Theory Interpretation*. PhD thesis, Jacobs University Bremen.
- Núñez, R. E. (2005). Creating mathematical infinities: Metaphor, blending, and the beauty of transfinite cardinals. *Journal of Pragmatics*, 37:1717–1741.
- Pagán Cánovas, C. (2010). Erotic Emissions in Greek Poetry: A Generic Integration Network. *Cognitive Semiotics*, 6:7–32.
- Pereira, F. C. (2007). *Creativity and Artificial Intelligence: A Conceptual Blending Approach*, volume 4 of *Applications of Cognitive Linguistics (ACL)*. Mouton de Gruyter, Berlin.
- Pereira, F. C. and Cardoso, A. (2003). Optimality Principles for Conceptual Blending: A First Computational Approach. *AISB Journal*, 1(4).
- Plotkin, G. D. (1970). A note on inductive generalization. *Machine Intelligence*, 5:153–163.
- Ritze, D., Meilicke, C., Šváb Zamazal, O., and Stuckenschmidt, H. (2009). A Pattern-based Ontology Matching Approach for Detecting Complex Correspondences. In *OM-09*, volume 551 of *CEUR*.
- 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 *Proc. of the 5th International Conference on Computational Creativity*, Ljubljana, Slovenia.
- Schwering, A., Krumnack, U., Kühnberger, K.-U., and Gust, H. (2009a). Syntactic Principles of Heuristic-Driven Theory Projection. *Cognitive Systems Research*, 10(3):251–269.
- Schwering, A., Kühnberger, K.-U., Krumnack, U., Gust, H., Wandmacher, T., Indurkha, B., and Ojha, A. (2009b). A Computational Model for Visual Metaphors: Interpreting Creative Visual Advertisements. In *International Conference on Agents and Artificial Intelligence (ICAART-09)*, pages 339–344.
- Turner, M. (2007). The Way We Imagine. In Roth, I., editor, *Imaginative Minds - Proc. of the British Academy*, pages 213–236. OUP, Oxford.
- Turner, M. (2014). *The Origin of Ideas: Blending, Creativity, and the Human Spark*. Oxford University Press.

- van Mulken, M., le Pair, R., and Forceville, C. (2010). The impact of perceived complexity, deviation and comprehension on the appreciation of visual metaphor in advertising across three European countries. *Journal of Pragmatics*, 42:3418–3430.
- Veale, T. (1997). Creativity as pastiche: A computational treatment of metaphoric blends, with special reference to cinematic “borrowing”. In *Proc. of Mind II: Computational Models of Creative Cognition*.
- Veale, T. (2012). From Conceptual Mash-ups to “Bad-Ass” Blends: A Robust Computational Model of Conceptual Blending. In *Proc. of the 2012 International Conference on Computational Creativity*, Dublin, Ireland.
- Veale, T. and O’Donoghue, D. (2001). Computation and Blending. *Cognitive Linguistics*, 11(3/4):253–281.
- Walshe, B. (2012). Identifying complex semantic matches. In *The Semantic Web: Research and Applications*, pages 849–853. Springer.
- Zimmermann, A., Krötzsch, M., Euzenat, J., and Hitzler, P. (2006). Formalizing Ontology Alignment and its Operations with Category Theory. In *Proc. of FOIS-06*, pages 277–288.