

Generalize and Blend: Concept Blending Based on Generalization, Analogy, and Amalgams

Tarek R. Besold

Institute of Cognitive Science
University of Osnabrück
D-49069 Osnabrück, Germany
tarek.besold@uni-osnabrueck.de

Enric Plaza

IIIA, Artificial Intelligence Research Institute
CSIC, Spanish Council for Scientific Research
Campus U.A.B., 08193 Bellaterra, Catalonia (Spain)
enric@iiia.csic.es

Abstract

Concept blending, a cognitive process which allows for the combination of certain elements (and their relations) from originally distinct conceptual spaces into a new unified space combining these previously separate elements and allowing the performance of reasoning and inference over the combination, is taken as a key element of creative thought and combinatorial creativity. In this paper, we provide an intermediate report on work towards the development of a computational-level and algorithmic-level account of concept blending, presenting the theoretical background together with the main model characteristics, as well as two case studies.

Creativity and Concept Blending

The term “combinatorial creativity” (Boden 2003) refers to creativity which arises from a combinatorial process joining familiar ideas (in the form of, for instance, concepts, theories, or artworks) in an unfamiliar way, thereby producing novel ideas. But although the overall idea of combining preexisting ideas into new ones seems fairly intuitive and straightforward, computationally modeling this form of creativity turns out to be surprisingly complicated: When looking at it from a more formal perspective at the current stage neither can a precise algorithmic characterization be given, nor are at least the details of a possible computational-level theory describing the process(es) at work well understood.

Still, in recent years a proposal by (Fauconnier and Turner 1998) called concept blending (or conceptual integration) has influenced and reinvigorated studies trying to unravel the general cognitive principles operating during creative thought. In their theory, concept blending constitutes a cognitive process which allows for the combination of certain elements (and their relations) from originally distinct conceptual spaces into a new unified space combining these previously separate elements and allowing the performance of reasoning and inference over the combination.

Unfortunately, a proper computational modeling of concept blending as cognitive capacity again is lacking. Neither do (Fauconnier and Turner 1998) provide a fully worked out and formalized theory themselves, nor does their informal account capture key properties and functionalities as, for example, the retrieval of input spaces, the selection and transfer of elements from the input into the blend space, or the

further combination of possibly mutually contradictory elements in the blend.

These shortcomings notwithstanding, several researchers in AI and computational cognitive modeling have used the provided conceptual descriptions as a starting point for proposing possible refinements and implementations: (Goguen and Harrell 2010) propose a concept blending-based approach to the analysis of the style of multimedia content in terms of blending principles and also provide an experimental implementation, (Pereira 2007) tries to develop a computationally plausible model of several hypothesized sub-parts of concept blending, (Thagard and Stewart 2011) exemplify how creative thinking could arise from using convolution to combine neural patterns into ones which are potentially novel and useful, and (Veale and O’Donoghue 2000) present their computational model of conceptual integration and propose several extensions to the (at that time prevailing) view on concept blending.

Since 2013, another attempt at developing a computationally feasible, cognitively-inspired formal model of concept creation, grounded on a sound mathematical theory of concepts and implemented in a generic, creative computational system is undertaken by a European research consortium in the so called Concept Invention Theory (COINVENT) project (Schorlemmer et al. 2014)¹. One of the main goals of the COINVENT research program is the development of a computational-level and algorithmic-level account of concept blending based on insights from psychology, AI, and cognitive modeling, the heart of which are made up by results from cognitive systems studies on computational analogy-making and knowledge transfer and combination (i.e., the computation of so called *amalgams*) from case-based reasoning. In the following we present an analogy-inspired perspective on the COINVENT core model for concept blending and show how the respective mechanisms and systems interact.

Two Mechanisms at the Heart of COINVENT: Generalization-Based Analogy and Amalgams

As analogy seems to play a crucial role in human cognition (Gentner and Smith 2013), researchers on the computa-

¹Also see <http://www.coinvent-project.eu> for details on the consortium and the project.

tional side of cognitive science and in AI also very quickly got interested in the topic and have been creating computational models of analogy-making since the advent of computer systems, among others giving rise to (Winston 1980)'s work on analogy and learning, (Hofstadter and Mitchell 1994)'s Copycat system, or (Falkenhainer, Forbus, and Gentner 1989)'s well-known Structure-Mapping Engine (SME).

Generally speaking there are (at least) two families of computational analogy models: one family is based on a (generalization-free) direct mapping approach, the other one relies on a two-step procedure with a generalization stage followed by a subsequent mapping stage. While the former type of analogy engine is, among others, exemplified in the SME and its immediate pairwise mapping of domain elements between elements of source and target of the potential analogy, followed by the accumulation of individual mappings into more complex structures, the latter category is represented by the Heuristic-Driven Theory Projection (HDTP) framework (Schmidt et al. 2014). As COINVENT, for principled conceptual reasons (see the section on the idea(s) behind concept blending in COINVENT below), relies on the generalization-based view on analogy-making, we shortly introduce this model category in the following subsection.

In a conceptually related, but mostly independently conducted line of work researchers in case-based reasoning (CBR) have been trying to develop problem solving methodologies based on the principle that similar problems tend to have similar solutions. CBR tries to solve problems by retrieving one or several cases relevant for the issue at hand from a case-base with already solved previous problems (*cases*), and then reusing the past case(s) to also solve the new task (Aamodt and Plaza 1994). While the retrieval stage has received significant attention over the last two decades, the transfer and combination of knowledge from the retrieved case to the current problem has been studied in an domain-specific way, with (Ontanón and Plaza 2012) being a recent attempt at also gaining insights on this phase of the CBR cycle by suggesting the framework of amalgams (Ontanón and Plaza 2010) as a formal model for reuse of multiple cases. The second subsection gives an overview of amalgams as used in COINVENT.

Generalization-Based Models of Analogy

Generalization-based models of analogy-making share a close conceptual connection to models of inductive generalization (Smaling 2003). Similar to these, the basic principle is the recognition of a common core between source and target of the potential analogy, which is then used for guiding the formation process of the analogy and the subsequent content transfer and reasoning steps. Fig. 1 gives a schematic overview: The common conceptual elements between source S and target T correspond to a shared generalization G (subsuming both, S and T), which also induces mappings between the respective domain elements, establishing an analogical relation. These mappings, governed by the generalization, then also subsequently define how (previously unmatched) knowledge from the source domain can be transferred to and integrated into the target do-




Figure 1: A schematic overview of a generalization-based approach to analogy.

main, namely by converting elements from S into their corresponding counterparts within T .

The precise nature of the subsumption relation between generalization and source or target domain, respectively, is defined by the specific analogy model, possibly ranging from semantic subsumption in a suitable ontology, through taxonomic subsumption based on names and labels, logical subsumption in a model-theoretic sense, to purely syntactic subsumption in a formal language.

One example for a generalization-based computational analogy-model (and the system used in COINVENT) is the already aforementioned HDTP (Schmidt et al. 2014). The framework has been conceived as a mathematically sound theoretical model and implemented engine for computational analogy-making, on a syntax basis computing analogical relations and inferences for domains which are presented in (when allowing for re-representation possibly different) many-sorted first-order logic (FOL) languages. Source and target domains are handed over to the system in terms of finite axiomatizations and HDTP tries to compute a generalization between both domains. This is done by aligning pairs of formulae from the two domains by means of restricted higher-order anti-unification (Schwering et al. 2009): Given two terms, one from each domain, HDTP computes an anti-instance in which distinct subterms have been replaced by variables so that the anti-instance can be seen as a meaningful generalization of the input terms. As already indicated by the name, the class of admissible substitution operations is limited. On each expression, only renamings, fixations, argument insertions, and permutations may be performed. By this process, HDTP tries to find the least general generalization of the input terms, which (due to the higher-order nature of the anti-unification) is not unique. In order to solve this problem, current implementations of HDTP rank possible generalizations according to a complexity measure on the chain of substitutions — the respective values of which are taken as heuristic costs — and returns the least expensive solution as the preferred one.

HDTP extends the notion of generalization from terms to formulae by basically treating formulae in clause form and terms alike. Finally, as analogies rarely rely exclusively on one isolated pair of formulae from source and target domain, but usually encompass sets of formulae (possibly completely covering one or even both input domains), a process iteratively selecting pairs of formulae for generalization has been included. The selection of formulae is again based on a heuristic component. Mappings in which substitutions can be reused get assigned a lower cost than isolated substitu-

tions, leading to a preference for coherent over incoherent mappings.

Due to the use of many-sorted FOL as an expressive representation language, and the purely syntax-based generalization approach underlying HDTP, over the last years the framework has shown remarkable generalizability and generality. Having originally been conceived and applied for modelling the Rutherford analogy and poetic metaphors, as well as for providing an alternate account of (Falkenhainer, Forbus, and Gentner 1989)’s heat-flow analogy in (Schwering et al. 2009), without major changes to the model HDTP has by now been applied to different tasks from different domains, such as modeling a potential inductive analogy-based process for establishing the fundamental concepts of arithmetics (Guhe et al. 2010), or studies applying the framework to modeling analogy use in education and teaching situations (Besold 2014).

Combining Conceptual Theories Using Amalgams

The notion of amalgams was developed in the context of CBR (Ontanón and Plaza 2010), where new problems are solved based on previously solved problems (or cases, residing on a case base). Solving a new problem often requires more than one case from the case base, so their content has to be combined in some way to solve the new problem. The notion of an amalgam of two cases (two descriptions of problems and their solutions) is a proposal to formalize the ways in which cases can be combined to produce a new, coherent case.

Formally, the notion of amalgams can be defined in any representation language \mathcal{L} for which a subsumption relation \sqsubseteq between the formulae (or descriptions) of \mathcal{L} can be defined. We say that a description I_1 subsumes another description I_2 ($I_1 \sqsubseteq I_2$) when I_1 is more general (or equal) than I_2 . Additionally, we assume that \mathcal{L} contains the infimum element \perp (or ‘any’), and the supremum element \top (or ‘none’) with respect to the subsumption order.

Next, for any two descriptions I_1 and I_2 in \mathcal{L} we can define their *unification*, $(I_1 \sqcup I_2)$, which is the *most general specialization* of two given descriptions, and their *anti-unification*, $(I_1 \sqcap I_2)$, defined as the *least general generalization* of two descriptions, representing the most specific description that subsumes both. Intuitively, a unifier is a description that has all the information in both the original descriptions; if joining this information leads to inconsistency, this is equivalent to saying that $I_1 \sqcup I_2 = \top$ (i.e., they have no common specialization except ‘none’). The anti-unification $I_1 \sqcap I_2$ contains all that is common to both I_1 and I_2 ; when they have nothing in common, then $I_1 \sqcap I_2 = \perp$. Depending on \mathcal{L} anti-unification and unification might be unique or not.

The notion of an *amalgam* can be conceived of as a generalization of the notion of unification: as ‘partial unification’ (Ontanón and Plaza 2010). Unification means that what is true for I_1 or I_2 is also true for $I_1 \sqcup I_2$; e.g., if I_1 describes ‘a red vehicle’ and I_2 describes ‘a German minivan’ then their unification yields a common specialization like ‘a red German minivan.’ Two descriptions may contain information that produces an inconsistency when unified; for instance




Figure 2: A diagram of an amalgam A from inputs I_1 and I_2 where $A = \bar{I}_1 \sqcup \bar{I}_2$.




Figure 3: A diagram that transfers content from source S to a target T via an asymmetric amalgam A .

‘a red French sedan’ and ‘a blue German minivan’ have no common specialization except \top . An *amalgam* of two descriptions is a new description that contains *parts from these two descriptions*. For instance, an amalgam of ‘a red French sedan’ and ‘a blue German minivan’ is ‘a red German sedan’; clearly there are always multiple possibilities for amalgams, like ‘a blue French minivan’.

For the purposes of this paper we can define an amalgam of two input descriptions as follows:

Definition 1 (Amalgam) *A description $A \in \mathcal{L}$ is an amalgam of two inputs I_1 and I_2 (with anti-unification $G = I_1 \sqcap I_2$) if there exist two generalizations \bar{I}_1 and \bar{I}_2 such that (1) $G \sqsubseteq \bar{I}_1 \sqsubseteq I_1$, (2) $G \sqsubseteq \bar{I}_2 \sqsubseteq I_2$, and (3) $A = \bar{I}_1 \sqcup \bar{I}_2$*

When \bar{I}_1 and \bar{I}_2 have no common specialization then trivially $A = \top$, since their only unifier is “none”. For our purpose we will be only interested in non-trivial amalgams.

This definition is illustrated in Fig. 2, where the anti-unification of the inputs is indicated as G , and the amalgam A is the unification of two concrete generalizations \bar{I}_1 and \bar{I}_2 of the inputs. Equality here should be understood as \sqsubseteq -equivalence: $X \equiv Y$ iff $X \sqsubseteq Y$ and $Y \sqsubseteq X$. Conventionally, we call the *space of amalgams* of I_1 and I_2 the set of all amalgams A that satisfy Definition 1.

Usually we are interested only in maximal amalgams of two input descriptions, i.e., those amalgams that contain maximal parts of their inputs that can be unified into a new coherent description. Formally, an amalgam A of inputs I_1 and I_2 is maximal if there is no other non-trivial amalgam A' of inputs I_1 and I_2 such that $A \sqsubset A'$. The reason why we are interested in maximal amalgams is very simple: a non-maximal amalgam $\bar{A} \sqsubset A$ preserves less compatible information from the inputs than the maximal amalgam A . Conversely, any non-maximal amalgam \bar{A} can be obtained by generalizing a maximal amalgam A , since $\bar{A} \sqsubset A$.

There is a special case of particular interest that is called

an *asymmetric amalgam*, in which the two inputs play different roles. The inputs are called source and target, and while the source is allowed to be generalized, the target is not.

Definition 2 (Asymmetric Amalgam) An asymmetric amalgam $A \in \mathcal{L}$ of two inputs S (source) and T (target) satisfies that $A = S' \sqcup T$ for some generalization of the source $S' \sqsubseteq S$.

As shown in Fig. 3, the content of target T is transferred completely into the asymmetric amalgam, while the source S is generalized. The result is a form of partial unification that preserves all information in T while relaxing S by generalization and then unifying one of those generalizations S' with T itself. As before, we will usually be interested in maximal amalgams: in this case, a maximal amalgam corresponds to transferring maximal content from S to T while keeping the resulting amalgam A consistent. For these reasons asymmetric amalgams can be seen as models of a form of analogical inference, transferring information from the source to the target by creating a new amalgam that enriches the latter with the content of S' (Ontanón and Plaza 2012).

Analogy-Based Concept Blending in COINVENT

Taking the concept of generalization-based analogies (and HDTP as suitable framework for the computation of the latter) together with the notion of asymmetric amalgams, we now can introduce the core idea(s) behind concept blending as performed in COINVENT in the next subsection, subsequently also showing the feasibility of the approach in two examples. The general suitability of the approach is demonstrated revisiting the “sign forest” metaphor from (Kutz et al. 2012), an implementation using HDTP is exemplified (re-)constructing the concept of a foldable toothbrush.

The Core Model: An Analogy-Inspired View

One of the early formal accounts on concept blending, which is especially influential to the approach applied in COINVENT, is the classical work by Goguen using notions from algebraic specification and category theory (Goguen 2006). This version of concept blending can be described by the diagram in Fig. 4, where each node stands for a representation an agent has of some concept or conceptual domain. We will call these representations “conceptual spaces” and in some cases abuse terminology by using the word “concept” to really refer to its representation by the agent. The arrows stand for morphisms, that is, functions that preserve at least part of the internal structure of the related conceptual spaces. The idea is that, given two conceptual spaces I_1 and I_2 as input, we look for a generalization G and then construct a blend space B in such a way as to preserve as many as possible structural alignments between I_1 and I_2 established by the generalization. This may involve taking the functions to B to be *partial*, in that not all the structure from I_1 and I_2 might be mapped to B . In any case, as the blend respects (to the largest possible extent) the relationship between I_1 and I_2 , the diagram will commute.

Concept invention by concept blending can then be phrased as the following task: given two representations of




Figure 4: A conceptual overview of (Goguen 2006)’s account of conceptual blending.

two domain theories I_1 and I_2 , we need first, to compute a generalized theory G of I_1 and I_2 (which codes the commonalities between I_1 and I_2) and second, to compute the blend theory B in a structure preserving way such that new properties hold in B . Ideally, these new properties in B are considered to be (moderately) interesting properties. In what follows, for reasons of simplicity and without loss of generality we assume that the additional properties are just provided by one of the two domains, i.e., we align the situation with a standard setting in computational analogy-making by renaming I_1 and I_2 . The domain providing the additional properties for the concept blend will be called source S , the domain providing the conceptual basis and receiving the additional features will be called target T .

In COINVENT’s account, the reasoning process is then triggered by the computation of the generalization G (generic space), where for concept invention we will only need the mapping mechanism and replace the transfer phase by a new blending algorithm. The mapping is achieved via the usual generalization process between S and T , in which a generalized theory is created that reflects common aspects of both spaces. The generalized theory can be projected back into the original spaces by specializations ϕ_S and ϕ_T , respectively. As S and T might contain elements which are not reflected in the shared generalization, it holds that $\phi_S(G) \subseteq S$ and $\phi_T(G) \subseteq T$. While in analogy making the analogical relations are used in the transfer phase to translate additional uncovered knowledge from the source to the target space, blending combines additional facts (i.e., elements from $S \setminus S_C$ or $T \setminus T_C$) from one or both spaces. Therefore the process of blending can build on the generalization and specializations provided by the analogy engine, but has to include a new mechanism for transfer and concept combination. Here, amalgams naturally come into play: The set of specializations can be inverted and applied to generalize the original source theory S into a more general version S' (forming a superset of the shared generalization G , also including previously uncovered knowledge from the source) which then can be combined into an asymmetric amalgam with the target theory T , forming the (possibly underspecified) proto-blend T' of both. In a final step, T' is then completed into the blended theory and output of the process T_B by applying corresponding specialization steps stored from the generalization process between S and T (see also Fig. 5).

If we now take the domains to be represented in the form of finite axiomatizations as processed by HDTP, in an im-

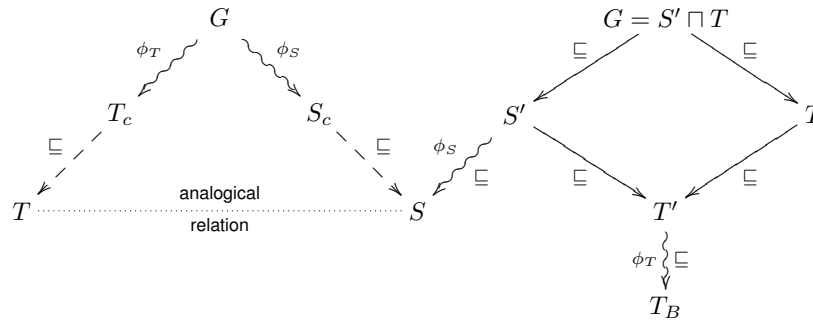


Figure 5: A general overview of COINVENT’s account of concept blending using generalization-based analogy and asymmetric amalgams: The shared generalization G from S and T is computed with $\phi_S(G) = S_c$. The relation ϕ_S is subsequently re-used in the generalization of S into S' , which is then combined in an asymmetric amalgam with T into the proto-blend $T' = S' \sqcup T$ and finally, by application of ϕ_T , completed into the blended output theory T_B . (Here \sqsubseteq indicates subsumption between theories in the direction of the respective arrows.)

plementation of the general model we can use the analogy-engine for computing the generalizations and deriving the corresponding substitutions. In the generalization step between S and T , as usual pairs of formulas from the source and target spaces are anti-unified for deriving the generalized theory G , and the specializations ϕ_S and ϕ_T become substitutions which are computed during anti-unification.

Example 1: The Sign Forest

We now want to revisit the example of the blend *sign forest* discussed in (Kutz et al. 2012), providing an interpretation of the concept from a metaphor-centered perspective and showing how the general COINVENT model can serve for reconstructing the blending process. In what follows we consider *sign forest* equivalent to the expression “a forest of signs”, that shows more clearly its metaphorical nature.

The original *sign forest* blend was defined in the context of blending ontologies, which means that the involved inputs for blending were ontological descriptions of trees, forests, and (traffic) signs. This approach views a concept such as *tree* defined as an ontological specification of the concept of tree: a specification that is ideally so general as to cover *all kinds of trees*; the same can be said about forest, and (traffic) signs. As such, certain properties and relations are selected to form these specifications that are useful for an ontology framework. However, our approach follows the notion that concepts in human cognition can often be viewed, in cognitive science, as bundles of their most typical properties (albeit typicality may certainly be context-dependent). This view is also taken in examples by (Fauconnier and Turner 1998) that are used to show how conceptual blending works: a boathouse has typical properties of boat and house —but not other properties that may appear in an ontological specification of boat and house.

Thus, in this approach, the concept of tree is typically formed by a plant having roots, a trunk and a crown (even if there may be plants categorized as trees that do not have a trunk, this is ignored as it does not belong to the bundle of properties that are typical); this view is depicted as I_2 in the bottom right of Fig. 6, where other properties are included,

like plants being not mobile and the roots fixing the (typical) tree to the ground. Finally, a forest is commonsensically defined as a group of trees. The second concept, (traffic) sign, may come in many forms (as we know from own experience), but the first that comes to mind is the most typical one: the signpost. The signpost is typically fixed on the ground near a road, and has a post supporting a surface panel depicting some traffic related information (labeled I_1 in the lower left corner of Fig. 6). The cognitive advantage of a signpost is that it has a recognizable physical structure, while “traffic sign” is so generic as to be a merely functional-based concept: any kind of surface panel depicting some traffic-related information is a traffic sign.

The generic space G of conceptual blending corresponds to the anti-unification shown as $G = I_1 \sqcap I_2$ in Fig. 6; G depicts common structure between a signpost and a tree: a stem-like object, fixed to the ground, and supporting another object on top. As discussed later, this common structure is the basis for a metaphor like “a forest of signs” to *make sense* — in contradistinction to a metaphor that does not *make sense* such as “a forest of chairs”, even when a typical chair is made of wood.

Now, the construction of the blended metaphor for *sign forest* can be interpreted easily in the combined generalization-based analogy and amalgam framework: the input spaces can be generalized in different ways (although always satisfying what they already have in common, namely G). Different generalizations would yield different amalgams, but the one we are considering here can be seen as generalizing I_2 into \bar{I}_2 , as shown in Fig. 6. Now this generalization \bar{I}_2 can directly be unified with I_1 , since \bar{I}_1 is identical to I_1 ; this unification yields the amalgam $A = \bar{I}_1 \sqcup \bar{I}_2$ that, as shown in Fig. 6, represents a “forest of signposts”. Moreover, since $I_1 \equiv \bar{I}_1$, this model is an asymmetric amalgam, as evidenced by the fact we generalize the source (Forest) until it unifies with the target (Signpost), while the latter remains fixed (i.e., is not generalized).

In order to support our perspective that a metaphor (viewed as an analogy and amalgam combination in natural language) is based on some (strong enough) common

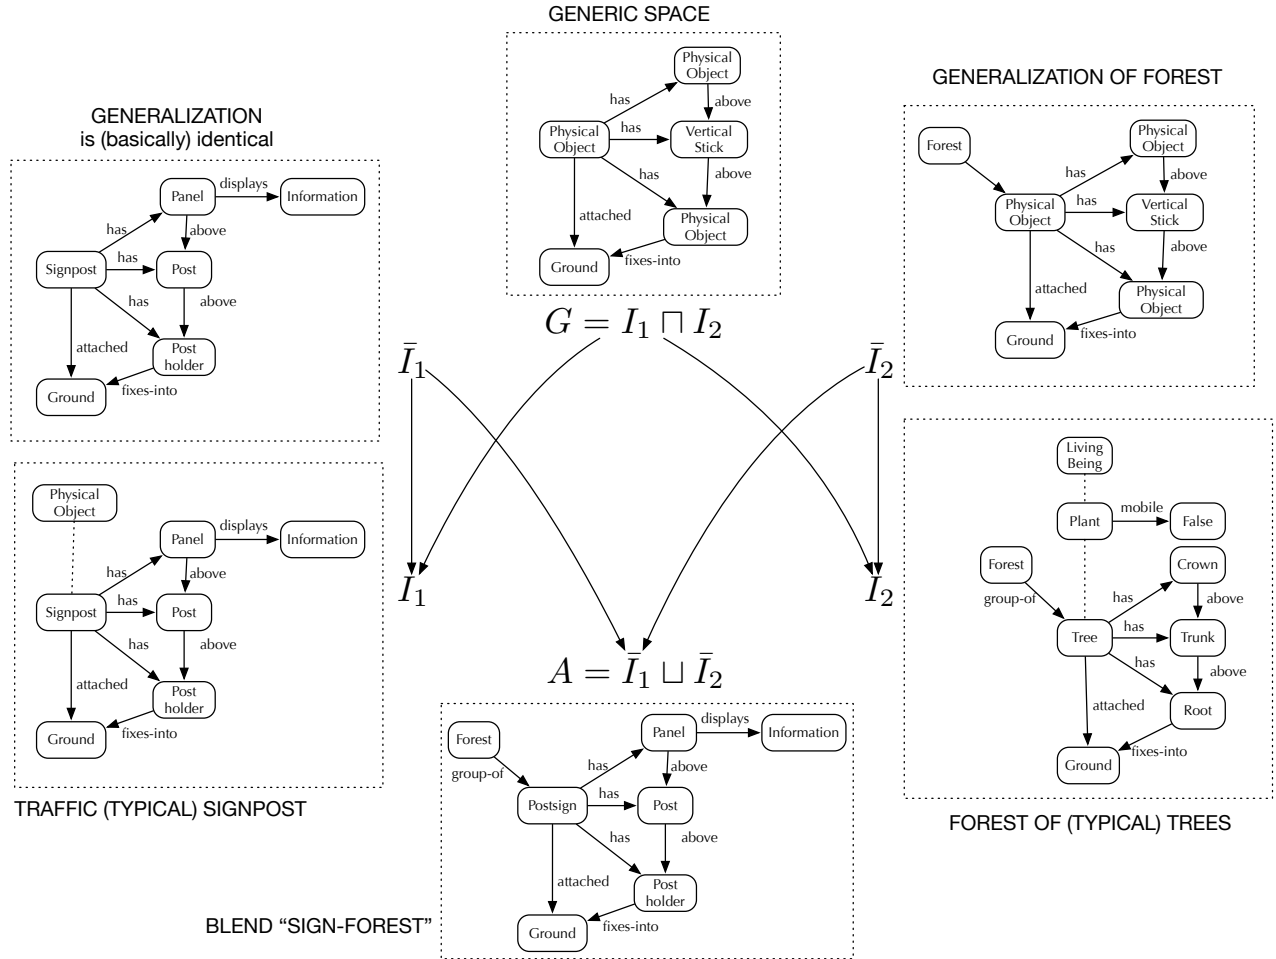


Figure 6: Blending schema for “Sign Forest” when inputs are typical concepts for “Sign” (traffic signpost) and “Forest” (forest of typical trees); the arrows indicate subsumption (\sqsubseteq) as in Figure 2.

structure of the typical concepts participating in the blending process, we checked if other metaphors can be constructed, or better yet, have already been constructed, that are based on the same kind of generic space G . We used Google’s n-grams database to search for existing phrases in which “forest of X ” is used metaphorically². Most n-grams starting with “forest of” were about places or kinds of trees, as is to be expected; still, we found the following metaphors used on the web: (1) forest of spears, (2) forest of masts, and (3) forest of marble columns. These three cases have a generic space that is very similar to G : they all represent a multitude of vertical stem-like objects. Some differences are: while masts and columns are fixed, spears are not fixed to the ground, but may be used in a context where they are vertical and immobile stems, supporting a pointed tip; masts

²Google’s 3-grams starting with “fo” are available at: <http://storage.googleapis.com/books/ngrams/books/googlebooks-eng-all-3gram-20120701-fo.gz>

and columns support different kinds of objects, but all three examples have generic spaces resembling G in Fig. 6.

What about counterexamples? We did not find “forest of chairs” of course, and there were other metaphors on forest, but they were based on different generic spaces and different input spaces; we found these metaphors: “forest of X ”, where X could be *opinions*, *possibilities*, *desires*, *words*, *human experience*. Clearly, these metaphors were not based on the trees being elements of a “forest”, but on the human experience of (walking in) the forest as a place of multiplicity of paths, options, destinations. We think they are not counterexamples, but rather examples of blends from different input spaces.

Example 2: The Folding Toothbrush

Having given an example for the general model in the previous subsection, we now want to also exemplify a concrete implementation of the approach using HDTP as anal-



Figure 7: Brillo, an example of a foldable toothbrush as produced by Metaphys.

ogy framework. As application example, we will use the blending-driven (re-)invention of foldable toothbrushes as, for instance, the one depicted in Fig. 7.

Currently, when using HDTP, the required subsumption relation between theories is given by logical semantic consequence \models , i.e., $A \sqsubseteq A'$ if $A' \models A$ for any two theories A and A' . In order to make sure that this relationship is preserved by HDTP's syntax-based operations, the range of admissible substitutions for restricted higher-order anti-unifications has to be further constrained to only allow for fixations and renamings.

Foldable toothbrushes are a conceptual combination between a typical toothbrush and a folding mechanism like that of pocketknives. In order to reconstruct the underlying blending process, we start with the stereotypical characterizations of a toothbrush and a pocketknife in a many-sorted first-order logic representation from Table 1.

Sorts:	
<i>entity, part, functionality</i>	
Entities:	
<i>toothbrush, pocketknife : entity</i>	<i>handle, brush.head, blade, hinge : part</i>
<i>brush, cut, fold : functionality</i>	
Predicates:	
<i>has_part : entity \times part, has_functionality : entity \times functionality</i>	
Laws of the pocketknife characterization:	
(α_1) <i>has_part(pocketknife, handle)</i>	(α_2) <i>has_part(pocketknife, blade)</i>
(α_3) <i>has_functionality(pocketknife, cut)</i>	(α_4) <i>has_part(pocketknife, hinge)</i>
(α_5) <i>has_functionality(pocketknife, fold)</i>	
Laws of the horse characterization:	
(β_1) <i>has_part(toothbrush, handle)</i>	(β_2) <i>has_part(toothbrush, brush.head)</i>
(β_3) <i>has_functionality(toothbrush, brush)</i>	

Table 1: Example formalizations of stereotypical characterizations for a pocketknife S and a toothbrush T .

Given these characterizations, HDTP can be used for

finding a common generalization of both, for instance (due to the syntactic similarities and the system's heuristics) aligning and generalizing α_1 with β_1 , α_2 with β_2 , and α_3 with β_3 . Subsequently, reusing the same anti-unifications (corresponding to ϕ_S), the source theory S is generalized into S' as given in Table 2: γ_1 corresponds to α_1/β_1 , γ_2 to α_2/β_2 , γ_3 to α_3/β_3 , and γ_4 and γ_5 are obtained by generalizing α_4 and α_5 , respectively.

Entities:	
<i>E : entity, P : part, F : functionality</i>	
Laws:	
(γ_1) <i>has_part(E, handle)</i>	(γ_2) <i>has_part(E, P)</i>
(γ_3) <i>has_functionality(E, F)</i>	
(γ_4^*) <i>has_part(E, hinge)</i>	(γ_5^*) <i>has_functionality(E, fold)</i>

Table 2: Abbreviated representation of the generalized source theory S' based on the stereotypical characterizations for a toothbrush and a pocketknife (axioms not obtained from the covered subset S_c are highlighted by *).

Computing the asymmetric amalgam of S' with the (fixed) target theory T , we obtain the proto-blend T' from Table 3. As T' still features axioms containing non-instantiated variables, ϕ_T is applied to the theory resulting in the (with respect to ϕ_T) fully instantiated blend theory T_B from Table 4, describing the concept of a hinge-equipped toothbrush that can be folded.

Entities:	
<i>E : entity</i>	
Laws:	
(δ_1) <i>has_part(toothbrush, handle)</i>	(δ_2) <i>has_part(toothbrush, brush.head)</i>
(δ_3) <i>has_functionality(toothbrush, brush)</i>	
(δ_4) <i>has_part(E, hinge)</i>	(δ_5) <i>has_functionality(E, fold)</i>

Table 3: Abbreviated representation of the proto-blend T' obtained from computing the asymmetric amalgam between S' and T .

Laws:	
(δ_1) <i>has_part(toothbrush, handle)</i>	(δ_2) <i>has_part(toothbrush, brush.head)</i>
(δ_3) <i>has_functionality(toothbrush, brush)</i>	
(δ_4) <i>has_part(toothbrush, hinge)</i>	(δ_5) <i>has_functionality(toothbrush, fold)</i>

Table 4: Abbreviated representation of $T_B = \phi_T(T')$.

Conclusions

We presented a perspective on the blending of concept theories building on generalization-based analogy and the amalgam framework: Building upon analogy models of generalization and domain matching, asymmetric amalgams allow to provide a sound model for the controlled computation of the concept blend(s) of two input theories.

Clearly, this is not the only attempt at developing a computational model of (some facet of) concept blending: (Martinez et al. 2014) present an algorithmic approach for blending mathematical theories, (Kutz et al. 2015) give an account

of ontological blending, (Li et al. 2012) describe the goal- and context-sensitive blending-based production of creative artifacts, and (Martinez et al. 2012) consider concept blending in a human-level AI context. Still, in combining the generality of generalization-based analogies and the amalgam framework, COINVENT's approach stands out as high-level, cognitively-inspired perspective on concept blending.

Acknowledgements

The authors acknowledge the financial support of the Future and Emerging Technologies within the 7th Framework Programme for Research of the European Commission, under FET-Open grant 611553 (COINVENT).

References

- Aamodt, A., and Plaza, E. 1994. Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. *AI Communications* 7(1):39–59.
- Besold, T. R. 2014. Sensorimotor Analogies in Learning Abstract Skills and Knowledge: Modeling Analogy-Supported Education in Mathematics and Physics. In *Proc. of the AAAI Fall 2014 Symposium on Modeling Changing Perspectives: Reconceptualizing Sensorimotor Experiences*, volume FS-14-05 of *AAAI Press Technical Reports*.
- Boden, M. A. 2003. *The Creative Mind: Myths and Mechanisms*. Routledge.
- Falkenhainer, B.; Forbus, K.; and Gentner, D. 1989. The structure-mapping engine: Algorithm and examples. *Artificial Intelligence* 41(1):1–63.
- Fauconnier, G., and Turner, M. 1998. Conceptual Integration Networks. *Cognitive Science* 22(2):133–187.
- Gentner, D., and Smith, L. A. 2013. Analogical learning and reasoning. In Reisberg, D., ed., *The Oxford Handbook of Cognitive Psychology*. Oxford University Press. 668–681.
- Goguen, J. A., and Harrell, D. F. 2010. Style: A Computational and Conceptual Blending-Based Approach. In Argamon, S.; Burns, K.; and Dubnov, S., eds., *The Structure of Style*. Springer. 291–316.
- Goguen, J. 2006. Mathematical models of cognitive space and time. In Andler, D.; Ogawa, Y.; Okada, M.; and Watanabe, S., eds., *Reasoning and Cognition; Proc. of the Interdisciplinary Conference Series on Reasoning Studies*, 125–128.
- Guhe, M.; Pease, A.; Smaill, A.; Schmidt, M.; Gust, H.; Kühnberger, K.-U.; and Krumnack, U. 2010. Mathematical reasoning with higher-order anti-unification. In *Proc. of the 32nd Annual Conference of the Cognitive Science Society*, 1992–1997. Cognitive Science Society.
- Hofstadter, D., and Mitchell, M. 1994. The Copycat project: a model of mental fluidity and analogy-making. In *Advances in Connectionist and Neural Computation Theory*, volume 2: Analogical Connections, 31–112. Ablex.
- Kutz, O.; Mossakowski, T.; Hois, J.; Bhatt, M.; and Bateman, J. 2012. Ontological Blending in DOL. In *Proc. of the 1st International Workshop on “Computational Creativity, Concept Invention, and General Intelligence”*, Publications of the Institute of Cognitive Science, Univ. of Osnabrück.
- Kutz, O.; Bateman, J.; Neuhaus, F.; Mossakowski, T.; and Bhatt, M. 2015. E Pluribus Unum. In Besold, T. R.; Schorlemmer, M.; and Smaill, A., eds., *Computational Creativity Research: Towards Creative Machines*, volume 7 of *Atlantis Thinking Machines*. Atlantis Press. 167–196.
- Li, B.; Zook, A.; Davis, N.; and Riedl, M. 2012. Goal-Driven Conceptual Blending: A Computational Approach for Creativity. In *Proc. of the Third International Conference on Computational Creativity*, 9–16.
- Martinez, M.; Besold, T. R.; Abdel-Fattah, A.; Gust, H.; Schmidt, M.; Krumnack, U.; and Kühnberger, K.-U. 2012. Theory Blending as a Framework for Creativity in Systems for General Intelligence. In Wang, P., and Goertzel, B., eds., *Theoretical Foundations of Artificial General Intelligence*. Atlantis Press. 219–239.
- Martinez, M.; Krumnack, U.; Smaill, A.; Besold, T. R.; Abdel-Fattah, A. M.; Schmidt, M.; Gust, H.; Kühnberger, K.-U.; Guhe, M.; and Pease, A. 2014. Algorithmic Aspects of Theory Blending. In Aranda-Corral, G.; Calmet, J.; and Martín-Mateos, F., eds., *Artificial Intelligence and Symbolic Computation*, volume 8884 of *LNCS*. Springer. 180–192.
- Ontanón, S., and Plaza, E. 2010. Amalgams: A Formal Approach for Combining Multiple Case Solutions. In Bichindaritz, I., and Montani, S., eds., *Case-Based Reasoning. Research and Development*, volume 6176 of *LNCS*. Springer. 257–271.
- Ontanón, S., and Plaza, E. 2012. On Knowledge Transfer in Case-Based Inference. In Agudo, B. D., and Watson, I., eds., *Case-Based Reasoning Research and Development*, volume 7466 of *LNCS*. Springer. 312–326.
- Pereira, F. C. 2007. *Creativity and AI: A Conceptual Blending Approach*. Mouton de Gruyter.
- Schmidt, M.; Krumnack, U.; Gust, H.; and Kühnberger, K.-U. 2014. Heuristic-Driven Theory Projection: An Overview. In Prade, H., and Richard, G., eds., *Computational Approaches to Analogical Reasoning: Current Trends*. Springer. 163–194.
- 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. 2009. Syntactic Principles of Heuristic-Driven Theory Projection. *Journal of Cognitive Systems Research* 10(3):251–269.
- Smaling, A. 2003. Inductive, analogical, and communicative generalization. *International Journal of Qualitative Methods* 2(1).
- Thagard, P., and Stewart, T. C. 2011. The AHA! Experience: Creativity Through Emergent Binding in Neural Networks. *Cognitive Science* 35(1):1–33.
- Veale, T., and O’Donoghue, D. 2000. Computation and Blending. *Cognitive Linguistics* 11(3/4):253–281.
- Winston, P. H. 1980. Learning and Reasoning by Analogy. *Commun. ACM* 23(12):689–703.