

Coherent Concept Invention

Marco Schorlemmer, Roberto Confalonieri, Enric Plaza

Artificial Intelligence Research Institute, IIIA-CSIC
Bellaterra (Barcelona), Catalonia, Spain
`{marco, confalonieri, enric}@iiia.csic.es`

Abstract. We address the problem on how newly invented concepts are evaluated with respect to a background ontology of conceptual knowledge so as to decide which of them are to be accepted into a system of familiar concepts, and how this, in turn, may affect the previously accepted conceptualisation. As technique to tackle this problem we explore the applicability of Paul Thagard’s computational theory of coherence. In particular, we propose a formalisation of Thagard’s notion of *conceptual coherence* for concepts represented in the \mathcal{AL} description logic and explore by means of an illustrative example the role coherence may play in the process of conceptual blending.

Keywords: conceptual blending, coherence, description logics

1 Introduction

Combinational creativity —when novel ideas (concepts, theories, solutions, works of art) are produced through unfamiliar combinations of familiar ideas— is, of the three forms of creativity put forward by Boden, the most difficult to capture computationally [2]. Putting concepts together to generate new concepts is, in principle, not a difficult task; but doing this in a computationally tractable way, and being able to recognise the value of newly invented concepts for better understanding a certain domain, is not as straightforward.

An important development that has significantly influenced the current understanding of the general cognitive principles operating during concept invention is Fauconnier and Turner’s theory of *conceptual blending* [6, 7]. Fauconnier and Turner proposed conceptual blending as the fundamental cognitive operation underlying much of everyday thought and language, and modelled it as a process by which humans subconsciously combine particular elements and their relations of originally separate conceptual spaces into a unified space, in which new elements and relations emerge, and new inferences can be drawn.

The theory has been primarily applied as an analytic tool for describing already existing blends of ideas and concepts in a varied number of fields, such as linguistics, music theory, poetics, mathematics, theory of art, political science, discourse analysis, philosophy, anthropology, and the study of gesture and of material culture [20]. But it has been also widely recognised to be a theory that can serve as a basis for computational models of creativity [5, 9, 10, 15, 21].

To guide the concept invention process, in addition to the blending mechanism *per se*, at least two additional dimensions need to be considered, namely the origin and destination of concept invention, i.e., from where (and how) input concepts are selected and to whom the concept invention is headed. Confalonieri et al. have proposed a process model for concept invention in which these dimensions are taken into account [3]. Inputs are selected based on a similarity measure that is computed relative to a Rich Background, and blends are evaluated using an argumentation framework based on value preferences of the audience for which concepts are invented.

In this paper, we aim at showing how Thagard’s computational theory of coherence [19] could also serve as an additional mechanism for triggering concept invention and evaluating newly blended concepts. In [18], Thagard suggested to use coherence as a model for the closely related cognitive process of *conceptual combination*, where the focus is primarily on language compositionality such as noun-noun or adjective-noun combinations [17]. Kunda and Thagard, for instance, show how conceptual coherence can be used for describing how we reason with social stereotypes [12].

Building upon Thagard’s intuitions and principles for modelling coherence, we propose a formalisation of Thagard’s notion of conceptual coherence for concepts represented in a description logic —we take the basic description logic \mathcal{AL} as a start— and further explore its applicability to conceptual blending. But instead of interpreting coherence or incoherence based on statistical correlations or causal relations (i.e., on frequencies of positive or negative association), we determine coherence and incoherence as dependent on how concept descriptions are stated. Failure to find conceptual blends that cohere with some given background knowledge leads to a search for alternative conceptual blends that eventually increase the overall coherence of the blend with the background knowledge.

The paper is organised as follows: In Section 2 we give a brief overview of Thagard’s computational theory of coherence, in Section 3 we introduce some core definitions regarding coherence and coherence graphs, and in Section 4 we provide a formalisation of conceptual coherence for the description logic \mathcal{AL} . Conceptual blending in \mathcal{AL} is described in Section 5, and coherence is applied to blending in Section 6. We conclude in Section 7.

2 Thagard’s Computational Theory of Coherence

Thagard addresses the problem of determining which pieces of information, such as hypotheses, beliefs, propositions or concepts, to accept and which to reject based on how they cohere and incohere among them, given that, when two elements cohere, they tend to be accepted together or rejected together; and when two elements incohere, one tends to be accepted while the other tends to be rejected [19].

This can be reformulated as a constraint satisfaction problem as follows. Pairs of elements that cohere between them form positive constraints, and pairs of elements that incohere between them form negative constraints. If we partition

the set of pieces of information we are dealing with into a set of accepted elements and a set of rejected elements, then a positive constraint is satisfied if both elements of the constraint are either among the accepted elements or among the rejected ones; and a negative constraint is satisfied if one element of the constraint is among the accepted ones and the other is among the rejected ones. The coherence problem is to find the partition that maximises the number of satisfied constraints.

Note that in general we may not be able to partition a set of elements as to satisfy *all* constraints, thus ending up accepting elements that incohere between them or rejecting an element that coheres with an accepted one. The objective is to minimise these undesired cases. The coherence problem is known to be NP-complete, though there exist algorithms that find good enough solutions of the coherence problem while remaining fairly efficient.

Depending on the kind of pieces of information we start from, and on the way the coherence and incoherence between these pieces of information is determined, we will be dealing with different kinds of coherence problems. So, in *explanatory coherence* we seek to determine the acceptance or rejection of hypotheses based on how they cohere and incohere with given evidence or with competing hypotheses; in *deductive coherence* we seek to determine the acceptance or rejection of beliefs based on how they cohere and incohere due to deductive entailment or contradiction; in *analogical coherence* we seek to determine the acceptance or rejection of mapping hypotheses based on how they cohere or incohere in terms of structure; and in *conceptual coherence* we seek to determine the acceptance or rejection of concepts based on how they cohere or incohere as the result of the positive or negative associations that can be established between them. Thagard discusses these and other kinds of coherence.

Although Thagard provides a clear technical description of the coherence problem as a constraint satisfaction problem, and he enumerates concrete principles that characterise different kinds of coherences, such as those discussed later in Section 4 for conceptual coherence, he does not clarify the actual nature of the coherence and incoherence relations that arise between pieces of information, nor does he suggest a precise formalisation of the principles he discusses. Joseph et al. have proposed a concrete formalisation and realisation of deductive coherence [11], which they applied to tackle the problem of norm adoption in normative multi-agent systems. In this paper, we shall focus on the problem of conceptual coherence and its applicability to conceptual blending.

3 Preliminaries: Coherence Graphs

In this section we give precise definitions of the concepts intuitively introduced in the previous section.

Definition 1. A coherence graph is an edge-weighted, undirected graph $G = \langle V, E, w \rangle$, where:

1. V is a finite set of nodes representing pieces of information.

2. $E \subseteq V^{(2)}$ (where $V^{(2)} = \{\{u, v\} \mid u, v \in V\}$) is a finite set of edges representing the coherence or incoherence between pieces of information.
3. $w : E \rightarrow [-1, 1] \setminus \{0\}$ is an edge-weighted function that assigns a value to the coherence between pieces of information.

Edges of coherence graphs are also called constraints.

When we partition the set V of vertices of a coherence graph (i.e., the set of pieces of information) into a set A of accepted elements and a set $R = V \setminus A$ of rejected elements, then we can say when a constraint—an edge between vertices—is satisfied or not by the partition.

Definition 2. Given a coherence graph $G = \langle V, E, w \rangle$, and a partition (A, R) of V , the set of satisfied constraints $C_{(A,R)} \subseteq E$ is given by:

$$C_{(A,R)} = \left\{ \{u, v\} \in E \mid \begin{array}{l} u \in A \text{ iff } v \in A, \text{ whenever } w(\{u, v\}) > 0 \\ u \in A \text{ iff } v \in R, \text{ whenever } w(\{u, v\}) < 0 \end{array} \right\}$$

All other constraints (i.e., those in $E \setminus C_{(A,R)}$) are said to be unsatisfied.

The coherence problem is to find the partition of vertices that satisfies as much constraints as possible, i.e., to find the partition that maximises the coherence value as defined as follows, which makes coherence to be independent of the size of the coherence graph.

Definition 3. Given a coherence graph $G = \langle V, E, w \rangle$, the coherence of a partition (A, R) of V is given by

$$\kappa(G, (A, R)) = \frac{\sum_{\{u,v\} \in C_{(A,R)}} |w(\{u, v\})|}{|E|}$$

Notice that there may not exist a unique partition with a maximum coherence value. Actually, at least two partitions have the same coherence value, since $\kappa(G, (A, R)) = \kappa(G, (R, A))$ for any partition (A, R) of V .

4 Conceptual Coherence in Description Logics

Thagard characterises conceptual coherence with these principles [19]:

Symmetry: Conceptual coherence is a symmetric relation between pairs of concepts.

Association: A concept coheres with another concept if they are positively associated, i.e., if there are objects to which they both apply.

Given Concepts: The applicability of a concept to an object may be given perceptually or by some other reliable source.

Negative Association: A concept incoheres with another concept if they are negatively associated, i.e., if an object falling under one concept tends not to fall under the other concept.

Acceptance: The applicability of a concept to an object depends on the applicability of other concepts.

To provide a precise account of these principles we shall formalise *Association* and *Negative Association* between concepts expressed in a description logic, since these are the principles defining coherence and incoherence. We shall assume coherence between two concept descriptions when we have explicitly stated that one subsumes the other (“there are objects to which both apply”); and we shall assume incoherence when we have explicitly stated that they are disjoint (“an object falling under one concept tends not to fall under the other concept”).

Definition 4. Given a Tbox \mathcal{T} in description logic \mathcal{AL} and a pair of concept descriptions $C, D \notin \{\top, \perp\}$, we will say that:

- C coheres with D , if $C \sqsubseteq D \in \mathcal{T}$, and that
- C incoheres with D , if $C \sqsubseteq \neg D \in \mathcal{T}$ or $C \sqcap D \sqsubseteq \perp \in \mathcal{T}$.

In addition, coherence and incoherence between concept descriptions depend on the concept constructors used, and we will say that, for all atomic concepts A , atomic roles R , and concept descriptions $C, D \notin \{\top, \perp\}$:

- $\neg A$ incoheres with A ;
- $C \sqcap D$ coheres both with C and with D ;
- $\forall R.C$ coheres (or incoheres) with $\forall R.D$, if C coheres (or incoheres) with D .¹

Symmetry follows from the definition above, and *Acceptance* is captured by the aim of maximising coherence in a coherence graph. For this we need to define how a TBox determines a coherence graph, and, in order to keep the graph finite, we express coherence and incoherence only between non-trivial concept descriptions (i.e., excluding \top and \perp) that are explicitly stated in the TBox.

Definition 5. Let \mathcal{T} be a TBox in \mathcal{AL} . The set of non-trivial subconcepts of \mathcal{T} is given as

$$\text{sub}(\mathcal{T}) = \bigcup_{C \sqsubseteq D \in \mathcal{T}} \text{sub}(C) \cup \text{sub}(D)$$

where *sub* is defined over the structure of concept descriptions as follows:

$$\begin{aligned} \text{sub}(A) &= \{A\} \\ \text{sub}(\perp) &= \emptyset \\ \text{sub}(\top) &= \emptyset \\ \text{sub}(\neg A) &= \{\neg A, A\} \\ \text{sub}(C \sqcap D) &= \{C \sqcap D\} \cup \text{sub}(C) \cup \text{sub}(D) \\ \text{sub}(\forall R.C) &= \{\forall R.C\} \cup \text{sub}(C) \\ \text{sub}(\exists R.\top) &= \{\exists R.\top\} \end{aligned}$$

¹ Note that since \mathcal{AL} allows only for limited existential quantification we cannot provide a general rule for coherence between concept descriptions of the form $\exists R.\top$.

Definition 6. The coherence graph of a TBox \mathcal{T} is the edge-weighted, undirected graph $G = \langle V, E, w \rangle$ whose vertices are non-trivial subconcepts of \mathcal{T} (i.e., $V = \text{sub}(\mathcal{T})$), whose edges link subconcepts that either cohere or incohere according to Definition 4, and whose edge-weight function w is given as follows:

$$w(\{C, D\}) = \begin{cases} 1 & \text{if } C \text{ and } D \text{ cohere} \\ -1 & \text{if } C \text{ and } D \text{ incohere} \end{cases}$$

5 Conceptual Blending in \mathcal{AL}

We follow the modelling principles and techniques of [4], where the process of conceptual blending is characterised by the notion of amalgams [1, 14]. According to this approach, the process of conceptual blending can be described as follows:

1. We take a taxonomy of concepts described in a background ontology expressed as a Tbox \mathcal{T} .
2. A mental space of an atomic concept A is modelled, for the purpose of conceptual blending, by means of a subsumption $A \sqsubseteq C$ specifying the necessary conditions we are focusing on.
3. The new concept to be invented is represented by the concept description that conjoins the atomic concepts to be blended.
4. With amalgams we generalise the input spaces based on the taxonomy in our TBox until a satisfactory blend is generated.

Formally, the notion of amalgams can be defined in any representation language \mathcal{L} for which a subsumption relation between formulas (or descriptions) of \mathcal{L} can be defined, and therefore also in the set of all \mathcal{AL} concept descriptions with the subsumption relation $\sqsubseteq_{\mathcal{T}}$.

To formally specify an amalgam we first need to introduce some notions. Let N_C be a set of concept names, N_R be a set of role names, and $\mathcal{L}(\mathcal{T})$ be the finite set of all \mathcal{AL} concept descriptions that can be formed with the concept and role names occurring in an \mathcal{AL} TBox \mathcal{T} . Then:

Definition 7. Given two descriptions $C_1, C_2 \in \mathcal{L}(\mathcal{T})$:

- A most general specialisation (MGS) is a description C_{mgs} such that $C_{mgs} \sqsubseteq_{\mathcal{T}} C_1$ and $C_{mgs} \sqsubseteq_{\mathcal{T}} C_2$ and for any other description D such that $D \sqsubseteq_{\mathcal{T}} C_1$ and $D \sqsubseteq_{\mathcal{T}} C_2$, then $D \sqsubseteq_{\mathcal{T}} C_{mgs}$.
- A least general generalisation (LGG) is a description C_{lgg} such that $C_1 \sqsubseteq_{\mathcal{T}} C_{lgg}$ and $C_2 \sqsubseteq_{\mathcal{T}} C_{lgg}$ and for any other description D such that $C_1 \sqsubseteq_{\mathcal{T}} D$ and $C_2 \sqsubseteq_{\mathcal{T}} D$, then $C_{lgg} \sqsubseteq_{\mathcal{T}} D$.

Intuitively, an MGS is a description that has some of the information from both original descriptions C_1 and C_2 , while an LGG contains what is common to them.

An *amalgam* or *blend* of two descriptions is a new description that contains parts from these original descriptions and it can be formally defined as follows.

Definition 8 (Amalgam). Let \mathcal{T} be an \mathcal{AL} TBox. A description $C_{am} \in \mathcal{L}(\mathcal{T})$ is an amalgam of two descriptions C_1 and C_2 (with LGG C_{lgg}) if there exist two descriptions $\overline{C_1}$ and $\overline{C_2}$ such that: $C_1 \sqsubseteq_{\mathcal{T}} \overline{C_1} \sqsubseteq_{\mathcal{T}} C_{lgg}$, $C_2 \sqsubseteq_{\mathcal{T}} \overline{C_2} \sqsubseteq_{\mathcal{T}} C_{lgg}$, and C_{am} is an MGS of $\overline{C_1}$ and $\overline{C_2}$.

The number of blends that satisfies the above definition can be very large and selection criteria for filtering and ordering them are therefore needed. Fauconnier and Turner discussed optimality principles [7], however, these principles are difficult to capture in a computational way, and other selection strategies need to be explored. Since we use a logical theory such as \mathcal{AL} , one way to evaluate a blend is consistency checking. Another alternative, that we will investigate in this paper, is to evaluate blends in terms of conceptual coherence.

The LGG and the generalised descriptions, needed to compute the amalgam as defined above, are obtained by means of a generalisation refinement operator that allows us to find generalisations of \mathcal{AL} concept descriptions.

5.1 Generalising \mathcal{AL} descriptions

Roughly speaking, a generalisation operator takes a concept C as input and returns a set of descriptions that are more general than C by taking a Tbox \mathcal{T} into account.

In order to define a generalisation refinement operator for \mathcal{AL} , we define the upward cover set of atomic concepts. In the following definition, $\text{sub}(\mathcal{T})$ (Definition 5) guarantees the following upward cover set to be finite.

Definition 9. Let \mathcal{T} be an \mathcal{AL} TBox with concept names from N_C . The upward cover set of an atomic concept $A \in N_C \cup \{\top, \perp\}$ with respect to \mathcal{T} is given as:

$$\begin{aligned} \text{UpCov}(A) := \{C \in \text{sub}(\mathcal{T}) \cup \{\top, \perp\} \mid A \sqsubseteq_{\mathcal{T}} C \\ \text{and there is no } C' \in \text{sub}(\mathcal{T}) \cup \{\top, \perp\} \\ \text{such that } A \sqsubset_{\mathcal{T}} C' \sqsubset_{\mathcal{T}} C\} \end{aligned} \quad (1)$$

We can now define our generalisation refinement operator for \mathcal{AL} as follows.

Definition 10. Let \mathcal{T} be an \mathcal{AL} TBox. We define the generalisation refinement operator γ inductively over the structure of concept descriptions as follows:

$$\begin{aligned} \gamma(A) &= \text{UpCov}(A) \\ \gamma(\top) &= \text{UpCov}(\top) = \emptyset \\ \gamma(\perp) &= \text{UpCov}(\perp) \\ \gamma(C \sqcap D) &= \{C' \sqcap D \mid C' \in \gamma(C)\} \cup \{C \sqcap D' \mid D' \in \gamma(D)\} \cup \{C, D\} \\ \gamma(\forall r.C) &= \begin{cases} \{\forall r.C' \mid C' \in \gamma(C)\} & \text{whenever } \gamma(C) \neq \emptyset \\ \{\top\} & \text{otherwise.} \end{cases} \\ \gamma(\exists r.C) &= \emptyset \end{aligned}$$

House \sqsubseteq Object	Resident \sqsubseteq Person
Boat \sqsubseteq Object	Passenger \sqsubseteq Person
Land \sqsubseteq Medium	Person \sqcap Medium $\sqsubseteq \perp$
Water \sqsubseteq Medium	Object \sqcap Medium $\sqsubseteq \perp$
Water \sqcap Land $\sqsubseteq \perp$	Object \sqcap Person $\sqsubseteq \perp$

Fig. 1. The background ontology of the House and Boat.

We should notice at this point that γ can return concept descriptions that are equivalent to the concept being generalised. One possible way to avoid this situation is to discard these generalisations [4]. Given a generalisation refinement operator γ , \mathcal{AL} concepts are related by refinement paths as described next.

Definition 11. A finite sequence C_1, \dots, C_n of \mathcal{AL} concepts is a concept refinement path $C_1 \xrightarrow{\gamma} C_n$ from C_1 to C_n of the generalisation refinement operator γ iff $C_{i+1} \in \gamma(C_i)$ for all $i : 1 \leq i < n$. $\gamma^*(C)$ denotes the set of all concepts that can be reached from C by means of γ in a finite number of steps.

The repetitive application of the generalisation refinement operator allows us to find a description that represents the properties that two or more \mathcal{AL} concepts have in common. This description is a common generalisation of \mathcal{AL} concepts, the so-called *generic space* that is used in conceptual blending.

Definition 12. An \mathcal{AL} concept description G is a generic space of the \mathcal{AL} concept descriptions C_1, \dots, C_n if and only if $G \in \gamma^*(C_i)$ for all $i = 1, \dots, n$.

5.2 An Example: The House-Boat Blend

The process of conceptual blending in terms of amalgams can be illustrated by means of a typical blend example: the *house-boat* [7, 8]. The precise formalisation is not critical at this point, different ones exist [9, 15], but all provide similar distinctions.

The \mathcal{AL} theories for House and Boat introduce the axioms modelling the mental spaces for *house* and *boat*.

$$\begin{aligned} \text{House} &\sqsubseteq \forall \text{usedBy.Resident} \sqcap \forall \text{on.Land} \\ \text{Boat} &\sqsubseteq \forall \text{usedBy.Passenger} \sqcap \forall \text{on.Water} \end{aligned}$$

The House and Boat theories cannot be directly blended since they generate an inconsistency. This is due to the background ontology stating that the medium on which an object is situated cannot be *land* and *water* at the same time (Figure 1). Therefore, some parts of the House and Boat descriptions need to be generalised in a controlled manner before these concepts can be blended. The generic space between a house and a boat—an object that is on a *medium* and *used-by* a *person*—is a lower bound in the space of generalisations that need to be explored in order to generalise these concepts and to blend them into a *house-boat*. The generic space is obtained according to Definition 12 by applying the refinement operator γ .

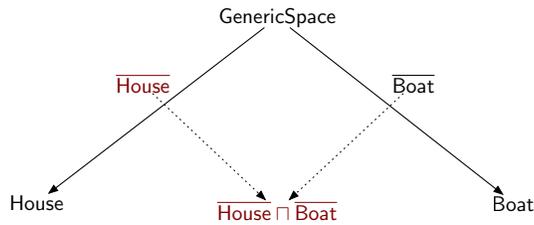


Fig. 2. A diagram of an amalgam $\overline{\text{HouseBoat}}$ from descriptions House and Boat and their respective generalisations $\overline{\text{House}}$ and $\overline{\text{Boat}}$. Arrows indicate the subsumption of the target by the source of the arrow.

Example 1. Let us consider the House and Boat concepts. Their generic space is: $\forall \text{usedBy. Person} \sqcap \forall \text{on. Medium}$ and is obtained as follows. In the House concept, the subconcepts $\forall \text{usedBy. Resident}$ and $\forall \text{on. Land}$ are generalised to $\forall \text{usedBy. Person}$ and $\forall \text{on. Medium}$ respectively. In the Boat concept, the subconcepts $\forall \text{usedBy. Passenger}$ and $\forall \text{on. Water}$ are generalised in a similar way.

From a conceptual blending point of view, the *house-boat* blend can be created when the medium on which a house is situated (land) becomes the medium on which boat is situated (water), and the resident of the house becomes the passenger of the boat. This blend can be obtained when the input concepts house and boat are generalised as follows:

$$\begin{aligned} \overline{\text{House}} &\sqsubseteq \forall \text{usedBy. Resident} \sqcap \forall \text{on. Medium} \\ \overline{\text{Boat}} &\sqsubseteq \forall \text{usedBy. Person} \sqcap \forall \text{on. Water} \end{aligned}$$

The *house-boat* blend is obtained by conjoining the generalised mental spaces $\overline{\text{House}}$ and $\overline{\text{Boat}}$ (Figure 2). It is easy to see that $\overline{\text{House}} \sqcap \overline{\text{Boat}}$ is an amalgam according to Definition 8.

6 Evaluating the Coherence of Conceptual Blends

This section describes how coherence is used to evaluate blends. That is, how coherence graphs are built, and how the different coherence values are to be interpreted. The overall idea is to compute the coherence graph and maximising partitions for each blend, and use the maximal coherence degree of the coherence graphs to rank the blends.

Let \mathcal{T} be the TBox of the background ontology, let $A \sqsubseteq C$ and $B \sqsubseteq D$ be the axioms representing our mental spaces, and let $A \sqcap B$ be the new concept we would like to invent. The process of evaluating blends according to conceptual coherence can be described as follows:

1. Given the mental spaces, we generate a candidate blend according to Definition 8.

2. We form the coherence graph for $\mathcal{T} \cup \{A \sqsubseteq C, B \sqsubseteq D\}$, including node $A \sqcap B$, according to Definition 6.
3. We compute the coherence maximising partitions according to Definition 3 and we associate it to the blend.
4. We repeat this procedure for all the blends that can be generated from the mental spaces.

Once the maximising partitions are computed, the coherence of the blend could be measured in terms of the coherence value of the coherence-maximising partitions. The degree of the coherence graph directly measures how much a blend coheres with the background ontology.

Definition 13. Let $G = \langle V, E, w \rangle$ the coherence graph of a blend B and let \mathcal{P} the set of partitions of G . The maximal coherence value of B of G is $\deg(B) = \max_{P \in \mathcal{P}} \{\kappa(G, P)\}$.

This maximal coherence value can be used to rank blends as follows.

Definition 14. Let \mathcal{T} be a *TBox* of a background ontology, let $A \sqsubseteq C$ and $B \sqsubseteq D$ be the axioms representing mental spaces, let \mathcal{B} be the set of blends that can be generated from them. For each $b_1, b_2 \in \mathcal{B}$, we say that b_1 is preferred to b_2 ($b_1 \succeq b_2$) if and only if $\deg(b_1) \geq \deg(b_2)$.

To exemplify how the coherence degree can be used to evaluate blends, we consider the *house-boat* example. According to the amalgams process of conceptual blending described in the previous section, several blends can be generated by blending the mental space of **House** and **Boat**. In particular, the concept $\text{House} \sqcap \text{Boat}$ is a valid blend.

The coherence graph blending the **House** and **Boat** directly is shown in Figure 3. As expected the concepts **House** and **Boat** positively coheres with the axioms representing the mental spaces and with the concept $\text{House} \sqcap \text{Boat}$, which is representing the blend. The incoherence relation between $\forall \text{on.Land}$ and $\forall \text{on.Water}$ is due to the fact that the concepts **Water** and **Land** incohere, since the background ontology contains the disjointness axiom $\text{Water} \sqcap \text{Land} \sqsubseteq \perp$. The coherence graph of **House** and **Boat** has a maximal coherence value of 0.84.

For the sake of our example, we generate new blends by generalising the axioms modelling our mental spaces. For instance, by applying the generalisations seen in the previous section that lead to the creation of the *house-boat* blend, we obtain the coherence graph in Figure 4.² The coherence graph of blending $\overline{\text{House}}$ and $\overline{\text{Boat}}$ has a maximal coherence value of 0.9. This graph yields a higher coherence degree since generalising $\forall \text{on.Land}$ to $\forall \text{on.Medium}$ prevents the appearance of the incoherence relation between $\forall \text{on.Land}$ and $\forall \text{on.Water}$.

It is easy to see that the blend $\overline{\text{House}} \sqcap \overline{\text{Boat}}$ is preferred to $\text{House} \sqcap \text{Boat}$ since it has a maximal coherence degree that is higher.

² Concepts belonging to the background ontology are omitted.

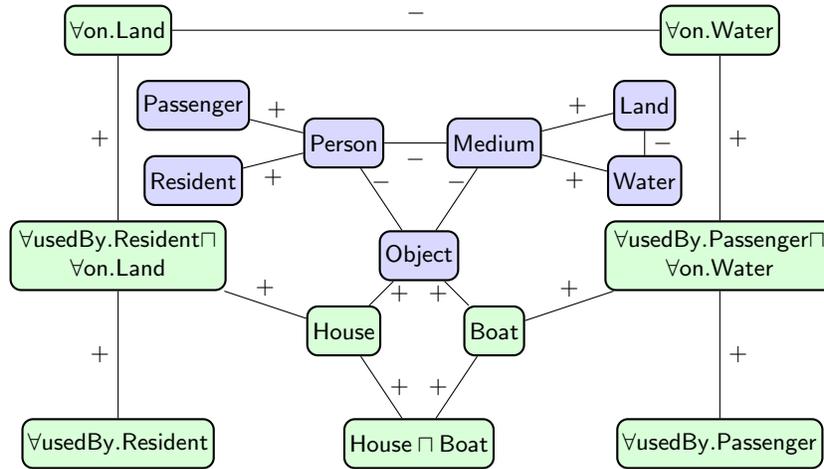


Fig. 3. The coherence graph of the $\text{House} \sqcup \text{Boat}$ blend, showing the main concepts and their coherence relations. Blue and green coloured boxes represent concepts belonging to the background ontology and to the input mental spaces respectively.

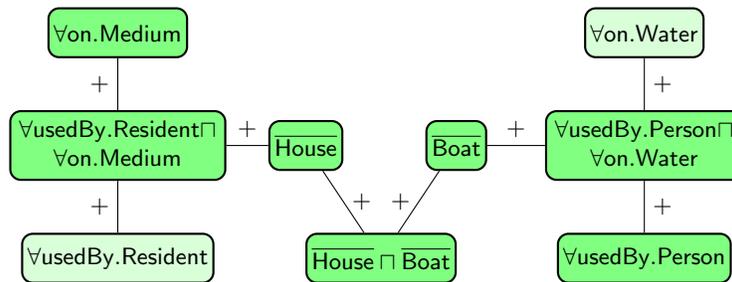


Fig. 4. The coherence graph of the $\overline{\text{House}} \sqcup \overline{\text{Boat}}$ blend, showing the main concepts and coherence relations. Generalised concepts are displayed in a darker tonality.

7 Conclusion

This paper should be seen as a first attempt to (a) provide a formal account of conceptual coherence for a particular concept representation language, and (b) to explore its applicability for guiding the process of conceptual blending.

With respect to (a), we proposed a formalisation of conceptual coherence between concept descriptions expressed in the basic \mathcal{AL} description logic. This is only a starting point, and obviously this formalisation exercise should be carried out for more expressive concept representation languages. Usually, coherence and incoherence are not treated only in binary terms, but it is also natural to take certain degrees of coherence or incoherence into account. This, for instance, has also been the approach of Joseph et al. when formalising deductive coherence [11]. Although there is not an obvious way to do so with the formalisation of

conceptual coherence of \mathcal{AL} proposed in this paper, we do not discard that this could be done for more expressive concept representation languages. One could imagine that description logics with number restrictions or nominals, such as *SRCIQ* for instance, would allow for expressing degrees of concept overlap that could be interpreted as degrees of coherence or incoherence.

With respect to (b), we have so far only focused on how the coherence values of a graph of concept descriptions were evolving dependent on how these descriptions were changing in our amalgam-based conceptual blending process. However, we have not discussed yet another important aspect of coherence theory, namely how to interpret the two parts of a coherence-maximising partition: the set of accepted and of rejected concepts. The information that a particular concept description falls in the set of accepted concepts or in the set of rejected concepts could also be taken into account to decide the acceptance or rejection of newly invented concepts; or even of already existing concepts in the background knowledge, in the light of newly invented concepts. With the formalisation in \mathcal{AL} given in this paper we could not see yet a clear way to provide such an interpretation of acceptance and rejection, but we think this aspect might become clearer as a wider range of concept representation languages is explored.

In this paper we attempted to see how coherence could be used as another tool for guiding the process of conceptual blending and for evaluating conceptual blends in the task of concept invention; an additional technique to those already proposed, such as optimality principles [16], logical consistency [13], and values of audiences [3]. We believe it is worth to further study the proper combination of these techniques and to carry out a comprehensive evaluation.

An implementation of conceptual coherence presented in this paper using the OWL API and Answer Set Programming is available at: <https://rconfalonieri@bitbucket.org/rconfalonieri/coinvent-coherence.git>.

References

- [1] T. R. Besold and E. Plaza. Generalize and Blend: Concept Blending Based on Generalization, Analogy, and Amalgams. In *Proceedings of the 6th International Conference on Computational Creativity, ICC15*, 2015.
- [2] M. A. Boden. *The Creative Mind: Myths and Mechanisms*. George Weidenfeld and Nicolson Ltd., 1990.
- [3] R. Confalonieri, E. Plaza, and M. Schorlemmer. A Process Model for Concept Invention. In *Proc. of the 7th International Conference on Computational Creativity, ICC16*, 2016.
- [4] R. Confalonieri, M. Schorlemmer, O. Kutz, R. Peñaloza, E. Plaza, and M. Eppe. Conceptual blending in EL++. In *Proc. of 29th Int. Workshop on Description Logics*, volume 1577 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2016.
- [5] M. Eppe, R. Confalonieri, E. Maclean, M. A. Kaliakatsos-Papakostas, E. Cambouropoulos, W. M. Schorlemmer, M. Codescu, and K. Kühnberger. Computational Invention of Cadences and Chord Progressions by Conceptual Chord-Blending. In *Proc. of the 24th Int. Joint Conf. on Artificial Intelligence, IJCAI 2015*, pages 2445–2451. AAAI Press, 2015.

- [6] G. Fauconnier and M. Turner. Conceptual integration networks. *Cognitive Science*, 22(2):133–187, 1998.
- [7] G. Fauconnier and M. Turner. *The Way We Think: Conceptual Blending and the Mind’s Hidden Complexities*. Basic Books, New York, 2003.
- [8] J. Goguen. An introduction to algebraic semiotics, with application to user interface design. In *Computation for Metaphors, Analogy, and Agents*, volume 1562 of *Lecture Notes in Computer Science*, pages 242–291. Springer, 1999.
- [9] J. A. Goguen and D. F. Harrell. Style: A computational and conceptual blending-based approach. In S. Argamon, K. Burns, and S. Dubnov, editors, *The Structure of Style*, chapter 12, pages 291–316. Springer, 2010.
- [10] M. Guhe, A. Pease, A. Smaill, M. Martínez, M. Schmidt, H. Gust, K.-U. Kühnberger, and U. Krumnack. A computational account of conceptual blending in basic mathematics. *Cognitive Systems Research*, 12(3–4):249–265, 2011.
- [11] S. Joseph, C. Sierra, M. Schorlemmer, and P. Dellunde. Deductive coherence and norm adoption. *Logic Journal of the IGPL*, 18(1):118–156, 2010.
- [12] Z. Kunda and P. Thagard. Forming unpressions from stereotypes, traits, and behaviours: A parallel-constraint-satisfaction theory. *Psychological Review*, 103(2):284–308, 1996.
- [13] F. Neuhaus, O. Kutz, M. Codescu, and T. Mossakowski. Fabricating monsters is hard. towards the automation of conceptual blending. In *Proceedings of the Workshop “Computational Creativity, Concept Invention, and General Intelligence” 2014*, volume 01-2014 of *Publications of the Institute of Cognitive Science*, 2014.
- [14] S. Ontañón and E. Plaza. Amalgams: A Formal Approach for Combining Multiple Case Solutions. In *Proceedings of the International Conference on Case Base Reasoning*, volume 6176 of *Lecture Notes in Computer Science*, pages 257–271. Springer, 2010.
- [15] F. C. Pereira. *Creativity and Artificial Intelligence: A Conceptual Blending Approach*. Mouton de Gruyter, 2007a.
- [16] F. C. Pereira and A. Cardoso. Optimality principles for conceptual blending: A first computational approach. *AISB Journal*, 1(4), 2003.
- [17] B. Ran and P. R. Duimering. Conceptual combination: Models, theories and controversies. *International Journal of Cognitive Linguistics*, 1(1):65–90, 2010.
- [18] P. Thagard. Coherent and creative conceptual combinations. In *Creative thought: An investigation of conceptual structures and processes*, pages 129–141. American Psychological Association, 1997.
- [19] P. Thagard. *Coherence in thought and action*. The MIT Press, 2000.
- [20] M. Turner. Blending and conceptual integration. <http://markturner.org/blending.html>. Last checked on June 20, 2016.
- [21] T. Veale and D. O’Donoghue. Computation and blending. *Cognitive Linguistics*, 11(3/4):253–281, 2000.