Coherence, Similarity, and Concept Generalisation

Roberto Confalonieri[1], Oliver Kutz[1], Pietro Galliani[1], Rafael Peñaloza[1], Daniele Porello[1], Marco Schorlemmer[2], and Nicolas Troquard[1]

[1] Free University of Bozen-Bolzano, Italy,
{firstname.lastname}@unibz.it
[2] Artificial Intelligence Research Institute, IIIA-CSIC, Spain
marco@iiia.csic.es

Abstract. We address the problem of analysing the joint coherence of a number of concepts with respect to a background ontology. To address this problem, we explore the applicability of Paul Thagard's computational theory of coherence, in combination with semantic similarity between concepts based on a generalisation operator. In particular, given the input concepts, our approach computes maximally coherent subsets of these concepts following Thagard's partitioning approach, whilst returning a number of possible generalisations of these concepts as justification of why these concepts cohere.

1 Introduction

In this paper, we aim at showing how Thagard's computational theory of coherence [15] could serve as an analytical tool to analyse the coherence of concepts that are inconsistent w.r.t. a background ontology. In [14], 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 [11]. Kunda and Thagard, for instance, show how conceptual coherence can be used for describing how we reason with social stereotypes [5].

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 the description logic $\mathcal{ALC}$ [1], and further explore its applicability to justify why a selection of concepts, even when jointly inconsistent, can be seen to cohere. But instead of interpreting coherence or incoherence on the basis of statistical correlations or causal relations (i.e., on frequencies of positive or negative association), we determine coherence and incoherence as dependent on the semantic similarity between concepts.

Given a set of input concepts, our approach computes maximally coherent subsets of these concepts following Thagard's computational model for coherence as a constraint satisfaction problem. We show how these maximising partitions not only suggest which of the input concepts can jointly cohere, but also how inconsistent concepts can be repaired using generalisations, letting them become consistent w.r.t. the background ontology.

To generalise these concepts, we propose a generalisation refinement operator which is inductively defined over the structure of $\mathcal{ALC}$ concept descriptions. Generalising DL concepts have been addressed in the DL literature in the context of the non-standard reasoning tasks of finding the least common subsumer of different concepts [2]. Finding a least common subsumer is a challenging research question, but, in practice, a common generalisation, w.r.t. the finite set of subformulas generated from the axioms contained in a (finite) TBox will suffice.

2 Thagard's Theory of Coherence

Thagard addresses the problem of determining which pieces of information 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 [15].

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 of 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 enumerates concrete principles that characterise different kinds of coherences, 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. In [4], a concrete formalisation and realisation of deductive coherence was proposed in order to tackle the problem of norm adoption in normative multi-agent system. In this paper, we shall focus on the problem of conceptual coherence and its applicability to analyse the coherence of $\mathcal{ALC}$ concepts.

3 Preliminaries

3.1 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 of the form $G = \langle V, E, w \rangle$, 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. Edges of coherence graphs are also called* constraints.
3. *$w : E \to [-1, 1] \setminus \{0\}$ is an edge-weighted function that assigns a value to the coherence between pieces of information.*

When we partition the set V of vertices of a coherence graph (i.e., the set of pieces of information) into the sets A and $R = V \setminus A$ of accepted and rejected elements respectively, 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:*

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

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 many 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:*

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

Notice that there may not exist a unique partition with a maximum coherence value. In fact, 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.

3.2 Generalising $\mathcal{ALC}$ Descriptions

To generalise $\mathcal{ALC}$ concept descriptions, we propose a generalisation refinement operator that extends our previous work in [3].

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 $\mathcal{ALC}$, we need some auxiliary definitions. In the following, we assume the TBox and set of concepts N_C to be finite.[3]

Definition 4. *Let $\mathcal{T}$ be a $\mathcal{ALC}$ TBox with concept names N_C. The set of* non-trivial subconcepts *of $\mathcal{T}$ is given as*

$$\mathsf{sub}(\mathcal{T}) = \{\top, \bot\} \cup \bigcup_{C \sqsubseteq D \in \mathcal{T}} \mathsf{sub}(C) \cup \mathsf{sub}(D)$$

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

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

Based on $\mathsf{sub}(\mathcal{T})$, we define the upward and downward cover sets of atomic concepts. In the following, we will assume that complex concepts C are rewritten into negation normal form, and that thus negation only appears in front of atomic concepts. For the following definition, $\mathsf{sub}(\mathcal{T})$ (Definition 4) guarantees the upward and downward cover sets to be finite. Intuitively, the upward set of A collects the most specific subconcepts found in the Tbox $\mathcal{T}$ that are more general (subsume) A; conversely, the downward set of A collects the most general subconcepts from $\mathcal{T}$ that are subsumed by A. The downcover is only needed for the base case of generalising a negated atom.

---

[3] To avoid any confusion, we point out that we use $=$ and $\neq$ between $\mathcal{ALC}$ concepts to denote syntactic identity and difference, respectively.

Definition 5. *Let $\mathcal{T}$ be an $\mathcal{ALC}$ TBox with concept names from N_C. The* upward cover set *of an atomic concept $A \in N_C \cup \{\top, \bot\}$ with respect to $\mathcal{T}$ is given as:*

$$\begin{aligned} \mathsf{UpCov}(A) := \{C \in \mathsf{sub}(\mathcal{T}) \mid {} & A \sqsubseteq_{\mathcal{T}} C \\ & \textit{and there is no } C' \in \mathsf{sub}(\mathcal{T}) \\ & \textit{such that } A \sqsubset_{\mathcal{T}} C' \sqsubset_{\mathcal{T}} C\} \end{aligned} \tag{1}$$

The downward cover set *of an atomic concept $A \in N_C \cup \{\top, \bot\}$ with respect to $\mathcal{T}$ is given as:*

$$\begin{aligned} \mathsf{DownCov}(A) := \{B \in N_C \cup \{\top, \bot\} \mid {} & B \sqsubseteq_{\mathcal{T}} A \\ & \textit{and there is no } B' \in N_C \cup \{\top, \bot\} \\ & \textit{such that } B \sqsubset_{\mathcal{T}} B' \sqsubset_{\mathcal{T}} A\} \end{aligned} \tag{2}$$

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

Definition 6. *Let $\mathcal{T}$ be an $\mathcal{ALC}$ TBox. We define the* generalisation refinement operator γ *inductively over the structure of concept descriptions as:*

$$\begin{aligned}
\gamma(A) &= \mathsf{UpCov}(A) \\
\gamma(\neg A) &= \begin{cases} \{\neg B \mid B \in \mathsf{DownCov}(A)\} & \textit{if } \mathsf{DownCov}(A) \neq \{\bot\} \\ \{\top\} & \textit{otherwise.} \end{cases} \\
\gamma(\top) &= \{\top\} \\
\gamma(\bot) &= \mathsf{UpCov}(\bot) \\
\gamma(C \sqcap D) &= \begin{cases} \{C' \sqcap D | C' \in \gamma(C)\} \cup \{C \sqcap D' | D' \in \gamma(D)\} \cup \{C, D\} & \textit{if } C \textit{ or } D \neq \top \\ \{\top\} & \textit{otherwise.} \end{cases} \\
\gamma(C \sqcup D) &= \begin{cases} \{C' \sqcup D \mid C' \in \gamma(C)\} \cup \{C \sqcup D' \mid D' \in \gamma(D)\} & \textit{if } C \textit{ and } D \neq \top \\ \{\top\} & \textit{otherwise.} \end{cases} \\
\gamma(\forall R.C) &= \begin{cases} \{\forall R.C' \mid C' \in \gamma(C)\} & \textit{if } C \neq \top \\ \{\top\} & \textit{otherwise.} \end{cases} \\
\gamma(\exists R.C) &= \begin{cases} \{\exists R.C' \mid C' \in \gamma(C)\} & \textit{if } C \neq \top \\ \{\top\} & \textit{otherwise.} \end{cases}
\end{aligned}$$

Lemma 1. *γ is a finite operator; i.e., for any given complex $\mathcal{ALC}$ concept C, the set $\gamma(C)$ is finite.*

Given a generalisation refinement operator γ, $\mathcal{ALC}$ concepts are related by refinement paths as described next.

Definition 7. *A finite sequence $C_1, \ldots, C_n$ of $\mathcal{ALC}$ concepts is a* generalisation 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 $i : 1 \leq i < n$. Then:*

– $\gamma^*(C)$ *denotes the set of all concepts that can be reached from* C *by means of* γ *in zero or a finite number of steps.*
– $\lambda(C \xrightarrow{\gamma} D)$ *denotes the minimal number of generalisations to be applied in order to generalise* C *to* D *when* $D \in \gamma^\star(C)$*.*

The repeated application of the generalisation refinement operator allows us to find descriptions that represent the properties that two or more $\mathcal{ALC}$ concepts have in common. This description is a common generalisation of $\mathcal{ALC}$ concepts.

For the sake of this paper, we are interested in common generalisations that have minimal distance from the concepts, or in case their distance is equal, the ones that are far from $\top$.

Definition 8. *An* $\mathcal{ALC}$ *concept description* G *is a common generalisation of* C_1 *and* C_2 *if* $G \in \gamma^*(C_1) \cap \gamma^*(C_2)$ *and, furthermore,* G *is such that for any other* $G' \in \gamma^*(C_1) \cap \gamma^*(C_2)$ *with* $(G' \neq G)$ *we have:*

– $\lambda(C_1 \xrightarrow{\gamma} G) + \lambda(C_2 \xrightarrow{\gamma} G) < \lambda(C_1 \xrightarrow{\gamma} G') + \lambda(C_2 \xrightarrow{\gamma} G')$*, or*
– $\lambda(C_1 \xrightarrow{\gamma} G) + \lambda(C_2 \xrightarrow{\gamma} G) = \lambda(C_1 \xrightarrow{\gamma} G') + \lambda(C_2 \xrightarrow{\gamma} G')$ *and* $\lambda(G \xrightarrow{\gamma} \top) \geq \lambda(G' \xrightarrow{\gamma} \top)$

At this point we should notice that common generalisations, as per the above definition, are not unique. However, for any two common generalisations G and G' of C_1 and C_2, $\lambda(C_1 \xrightarrow{\gamma} G) + \lambda(C_2 \xrightarrow{\gamma} G) = \lambda(C_1 \xrightarrow{\gamma} G') + \lambda(C_2 \xrightarrow{\gamma} G')$ and $\lambda(G \xrightarrow{\gamma} \top) = \lambda(G' \xrightarrow{\gamma} \top)$. Any one of them will result in the same value for our generalisation-based similarity measure between concepts, and therefore in the same coherence or incoherence judgements. In the following, we denote $C_1 \blacktriangle C_2$ a common generalisation of C_1 and C_2; $C_1 \blacktriangle C_2$ is a concept that always exists.

3.3 Concept Similarity

The common generalisation of two concepts C and D can be used to measure the similarity between concepts in a quantitative way. To estimate the quantity of information of any description C we take into account the length of the minimal generalisation path that leads from C to the most general term $\top$.

In order to define a similarity measure, we need to compare what is common to C and D with what is not common. The length $\lambda(C \blacktriangle D \xrightarrow{\gamma} \top)$ estimates the informational content that is common to C and D, and the lengths $\lambda(C \xrightarrow{\gamma} C \blacktriangle D)$ and $\lambda(D \xrightarrow{\gamma} C \blacktriangle D)$ measures how much C and D are different. Then, the common generalisation-based similarity measure can be defined as follows [9].

Definition 9. *The similarity between two concepts* C*,* D*, denoted by* $S_\lambda(C, D)$*, is defined as:*

$$S_\lambda(C, D) = \begin{cases} \dfrac{\lambda(C \blacktriangle D \xrightarrow{\gamma} \top)}{\lambda(C \blacktriangle D \xrightarrow{\gamma} \top) + \lambda(C \xrightarrow{\gamma} C \blacktriangle D) + \lambda(D \xrightarrow{\gamma} C \blacktriangle D)} & \text{if } C \text{ or } D \neq \top \\ 1 & \text{otherwise.} \end{cases}$$

The measure S_λ estimates the ratio between the amount of information that is shared and the total information content. The range of the similarity function is the interval $[0, 1]$, where 0 represents the minimal similarity between concepts (when their common generalisation is equal to $\top$), and 1 represents maximal similarity (when the concepts are equivalent).

4 Similarity-based Conceptual Coherence

Thagard characterises conceptual coherence with these principles [15]:

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 formal 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 they are sufficiently similar so that "there are objects to which both apply;" and we shall assume incoherence when they are not sufficiently similar so that "an object falling under one concept tends not to fall under the other concept."

In this formalisation of conceptual coherence, we determine this 'sufficiently similar' condition by taking into account the minimal length of the generalisation path from the common generalisation of two concepts to $\top$. The intuition behind the following definition is that similar concepts whose common generalisation is far from $\top$ should cohere, and incohere otherwise.

Definition 10 (Coherence Relations). *Given a set $\{C_1, \ldots, C_n\}$ of $\mathcal{ALC}$ concepts, we will say for each pair of concepts $\langle C_i, C_j \rangle$ $(1 \leq i, j \leq n, i \neq j)$:*

- *C_i coheres with C_j, if $S_\lambda(C_i, C_j) > 1 - \delta$*
- *C_i incoheres with C_j, if $S_\lambda(C_i, C_j) \leq 1 - \delta$*

where

$$\delta = \frac{\lambda(C_i \blacktriangle C_j \xrightarrow{\gamma} \top)}{\max\{\lambda(C_i \xrightarrow{\gamma} \top), \lambda(C_j \xrightarrow{\gamma} \top)\}}$$

In this definition, $\lambda(C_i \blacktriangle C_j \xrightarrow{\gamma} \top)$ is normalised to the interval $[0, 1]$ in order to make it comparable with the similarity measure. This is done by considering the range of values that $\lambda(C_i \blacktriangle C_j \xrightarrow{\gamma} \top)$ can assume. Since the maximal value corresponds to the case in which the common generalisation of C_i and C_j is $\top$, and the minimal value is 0, this interval is $[0, \max\{\lambda(C_i \xrightarrow{\gamma} \top), \lambda(C_j \xrightarrow{\gamma} \top)\}]$.

Returning to the Thargardian principles, *Symmetry* follows from the definition above, and *Acceptance* is captured by the aim of maximising coherence in a coherence graph.

Definition 11 (Thagardian Coherence Graph). *The* coherence graph *for the set of $\mathcal{ALC}$ concepts $\{C_1, \ldots, C_n\}$ is the edge-weighted and undirected graph $G = \langle V, E, w\rangle$ whose vertices are $C_1, \ldots, C_n$, whose edges link concepts that either cohere or incohere according to Definition 10, and whose edge-weight function w is given as follows:*

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

This definition creates a concept graph in the sense of Thagard where only binary values 'coheres' or 'incoheres' are recorded, represented by '+1' and '-1', respectively. However, it should be noted that Def. 10 can give rise also to graded versions of coherence graphs, which we will explore in future work.

5 Analysing the Coherence of Concepts

This section describes how we use coherence to analyse the joint coherence of a number of concepts with respect to a background ontology.

The overall idea is to compute the coherence graph and the maximising partitions for the input concepts, and use them to decide which concepts to keep and which ones to discard. The pairwise comparison and the maximising coherence degree partitions will give us the biggest subsets of coherent input concepts. Then, we compute the nearest common generalisations of the accepted concepts, to convey a justification of why certain concepts were partitioned together.

Given an $\mathcal{ALC}$ TBox representing a background ontology, and a set of $\mathcal{ALC}$ concepts $\{C_1, \ldots, C_n\}$ as input, the process of evaluating the coherence of concepts can be described as follows:

1. We form the coherence graph for the input concepts $C_1, \ldots, C_n$ according to Definition 11.
2. We compute the coherence maximising partitions according to Definition 3.
3. We use the partitions to decide which concepts to accept.
4. For each maximising partition of accepted concepts, we compute the nearest common generalisations, and present them as justifications of why these concepts were accepted.

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

$\mathsf{White} \sqsubseteq \mathsf{GrayScale}$	,	$\mathsf{Cats} \sqsubseteq \mathsf{Pets}$
$\mathsf{Black} \sqsubseteq \mathsf{GrayScale}$	,	$\mathsf{Pets} \sqsubseteq \mathsf{Animals}$
$\mathsf{GrayScale} \sqsubseteq \mathsf{Colours}$	,	$\mathsf{Black} \sqcap \mathsf{White} \sqsubseteq \bot$
$\mathsf{Integers} \sqsubseteq \mathsf{Numbers}$	,	$\mathsf{Numbers} \sqsubseteq \mathsf{Abstract\_Objects}$
$\mathsf{Animals} \sqsubseteq \mathsf{Physical\_Objects}$	,	$\mathsf{Physical\_Objects} \sqcap \mathsf{Abstract\_Objects} \sqsubseteq \bot$
$\mathsf{Qualities} \sqcap \mathsf{Abstract\_Objects} \sqsubseteq \bot$	,	$\mathsf{Primeness} \sqsubseteq \mathsf{Qualities}$
$\mathsf{Colours} \sqsubseteq \mathsf{Qualities}$	,	$\mathsf{Range(hasColour)} = \mathsf{Colours}$
$\mathsf{Domain(hasColour)} = \mathsf{Physical\_Object}$	,	$\mathsf{Range(hasQuality)} = \mathsf{Qualities}$
$\mathsf{Domain(hasQuality)} = \top$		

Fig. 1. The background ontology of Black_Cats, White_Cats, and Prime_Numbers.

It is worth noticing that according to our definition of coherence relation, inconsistent concepts can cohere provided that they are sufficiently similar and their common generalisation is far from the $\top$ concept.

Example. Let us consider the $\mathcal{ALC}$ theory in the TBox in Figure 1 and the following three input concepts:

$$\begin{array}{ll} \mathsf{Black\_Cats} & \equiv \mathsf{Cats} \sqcap \forall \mathsf{hasColour.Black} \sqcap \exists \mathsf{hasColour.Black} \\ \mathsf{White\_Cats} & \equiv \mathsf{Cats} \sqcap \forall \mathsf{hasColour.White} \sqcap \exists \mathsf{hasColour.White} \\ \mathsf{Prime\_Numbers} & \equiv \mathsf{Integers} \sqcap \exists \mathsf{hasQuality.Primeness} \end{array}$$

Black_Cats and White_Cats define black cats and white cats as cats that are coloured black and white respectively, whereas the Prime_Numbers defines the concept of prime numbers. We want to know whether these concepts cohere together or not.

Intuitively, Black_Cats and White_Cats, although inconsistent according to the background ontology, should cohere, since they "talk" about the same objects, i.e., cats. The Prime_Numbers concept, instead, should incohere with Black_Cats and White_Cats, since the objects it applies to are essentially different.

The coherence graph for these three concepts is computed as follows and it is shown in Figure 2:

- Black_Cats and White_Cats:
  - $\mathsf{Black\_Cats} \blacktriangle \mathsf{White\_Cats} = \mathsf{Cats} \sqcap \forall \mathsf{hasColour.GrayScale} \sqcap \exists \mathsf{hasColour.GrayScale}$
  - $\lambda(\mathsf{Black\_Cats} \blacktriangle \mathsf{White\_Cats} \xrightarrow{\gamma} \top) = 13$
  - $\lambda(\mathsf{Black\_Cats} \xrightarrow{\gamma} \mathsf{Black\_Cats} \blacktriangle \mathsf{White\_Cats}) = 2$
  - $\lambda(\mathsf{White\_Cats} \xrightarrow{\gamma} \mathsf{Black\_Cats} \blacktriangle \mathsf{White\_Cats}) = 2$
  - $S_\lambda(\mathsf{Black\_Cats}, \mathsf{White\_Cats}) = \frac{13}{17} = 0.76$
  - δ is: $\frac{13}{15} = 0.87$
  - Since $S_\lambda(\mathsf{Black\_Cats}, \mathsf{White\_Cats}) > 1 - \delta$, Black_Cats and White_Cats cohere.
- Black_Cats and Prime_Numbers
  - $\mathsf{Black\_Cats} \blacktriangle \mathsf{Prime\_Numbers} = \top$
  - $\lambda(\mathsf{Black\_Cats} \blacktriangle \mathsf{Prime\_Numbers} \xrightarrow{\gamma} \top) = 0$
  - $\lambda(\mathsf{Black\_Cats} \xrightarrow{\gamma} \mathsf{Black\_Cats} \blacktriangle \mathsf{Prime\_Numbers}) = 14$

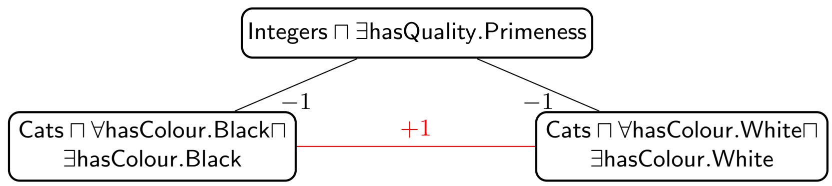

Fig. 2. The coherence graph of the Black_Cats, White_Cats, Prime_Numbers. A red coloured edge represents that the connected concepts are inconsistent.

  - $\lambda(\mathsf{Prime\_Numbers} \xrightarrow{\gamma} \mathsf{Black\_Cats} \blacktriangle \mathsf{Prime\_Numbers}) = 6$
  - $S_\lambda(\mathsf{Black\_Cats}, \mathsf{Prime\_Numbers})) = 0$
  - δ is: 0
  - Since $S_\lambda(\mathsf{Black\_Cats}, \mathsf{Prime\_Numbers})) \leq 1 - \delta$, we have that Black_Cats and Prime_Numbers incohere.
- White_Cats and Prime_Numbers incohere (similar to the previous case).

The maximising partitions of coherence graph are $A = \{\mathsf{Black\_Cats}, \mathsf{White\_Cats}\}$ and $R = \{\mathsf{Prime\_Numbers}\}$, and all constraints are satisfied (so $\kappa(G, (A, R)) = 1$). These concepts can still cohere, since they can be generalised in different ways.

For instance, by generalising Black and White to GrayScale, we can obtain the concept Cats ⊓ ∀hasColour.GrayScale ⊓ ∃hasColour.GrayScale that represents the category of gray-coloured cats. Or, by generalising Black and White to Colours, we can obtain the concept Cats ⊓ ∀hasColour.Colour ⊓ ∃hasColour.Colour that represents the category of coloured cats.

These generalisations, which can be obtained by applying our refinement operator γ, are to be considered explanations of why these concepts can cohere together. Therefore, our approach can improve the given 'coherence claim' by presenting the best generalisations that let the concepts be consistent.

As far as complexity is concerned, since subsumption reasoning in $\mathcal{ALC}$ is exponential, our proposed methodology stays in the same complexity class of coherence theory as a constraint satisfaction problem, namely, NP.

6 Discussion

This paper is a preliminary work, and the obvious next step would be to test it extensively on a number of real use cases, to verify the degree up to which our approach agrees with human intuitions.

There exist many possible variants of our definitions of $S_\lambda(C, D)$ and δ, and there is a very rich literature about similarity measures between concepts in ontologies [6], and especially w.r.t. Gene Ontology annotations [8].

For instance, one possible criticism of our approach (as well as of any "edge-based" similarity measure) is that the path length λ between two concepts depends on the *granularity* of our upward cover set: in particular, it is possible that it contains a high number of elements along the shortest path from C to $C \blacktriangle D$ (and, consequently, for the similarity between C and D to be low, and for them to possibly be incoherent) merely because our TBox (and, therefore, our cover set) contains many more assertions about generalisations of C than about generalisations of $C \blacktriangle D$.

It is not clear the degree up to which this is a problem in practice: for instance, it might be possible to reply that this is working as intended, since if our TBox contains many claims about concepts between C and $C \blacktriangle D$ then the differences between these concepts are indeed particularly relevant in our context, and hence it is appropriate for C and D to be comparatively less coherent.

In any case, if necessary, there exist ways around this: for instance, given an ABox of *facts* about various entities, we might employ *semantic similarity* measures such as Resnik's [12] or Lin's [7], which measure the similarity between C and D by comparing the *information content* (in brief, the information content of a concept is the logarithm of the occurrence probability of it or of any generalisation) of C, D and of $C \blacktriangle D$. Such similarity measures could be adopted in our approach easily enough; but there is little point in doing so unless we first establish a baseline and a way to compare their predictions.

We leave such issues, as well as the more general question of the evaluation of joint coherence analysis approaches, to further work. Here we contented ourselves with establishing a first such approach, which may then be tweaked according to its performance and to its user's needs.

7 Conclusion and Future Perspectives

In this work, we introduced a novel approach to the problem of analysing the joint coherence of a number of concepts with respect to a background ontology. This paper should be seen as an attempt to (a) provide a formal account of conceptual coherence for a particular concept representation language, and (b) to explore its applicability for analysing the joint coherence of a number of concepts with respect to a background ontology.

With respect to (a), a previous attempt to formalise conceptual coherence in the $\mathcal{AL}$ description logic is [13], where the authors attempted to see how coherence could be used as a tool for guiding the process of conceptual blending and for evaluating conceptual blends in the task of concept invention. Here, we proposed a formalisation of conceptual coherence between concept descriptions expressed in the $\mathcal{ALC}$ description logic. This is only a starting point, and obviously this formalisation exercise should be carried out for more expressive concept representation languages. Moreover, 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 [4]. As already remarked in the paper, our definitions can be extended to graded coherence and incoherence relations, and we aim at exploring this in the future.

With respect to (b), we have only focused on how maximally coherent sets suggest why inconsistent concepts can cohere together, and the way in which these concepts can be generalised to steer inconsistency debugging. In the future, we will investigate the relation of our approach to ontology debugging [10].

References

1. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F. (eds.): The Description Logic Handbook: Theory, Implementation, and Applications. Cambridge University Press, New York, NY, USA (2003)
2. Baader, F., Küsters, R.: Non-standard Inferences in Description Logics: The Story So Far. In: Mathematical Problems from Applied Logic I, International Mathematical Series, vol. 4, pp. 1–75. Springer New York (2006)
3. Confalonieri, R., Eppe, M., Schorlemmer, M., Kutz, O., Peñaloza, R., Plaza, E.: Upward Refinement Operators for Conceptual Blending in $\mathcal{EL}^{++}$. Annals of Mathematics and Artificial Intelligence (2016), doi:10.1007/s10472-016-9524-8
4. Joseph, S., Sierra, C., Schorlemmer, M., Dellunde, P.: Deductive coherence and norm adoption. Logic Journal of the IGPL 18(1), 118–156 (2010)
5. Kunda, Z., Thagard, P.: Forming unpressions from stereotypes, traits, and behaviours: A parallel-constraint-satisfaction theory. Psychological Review 103(2), 284–308 (1996)
6. Lee, W.N., Shah, N., Sundlass, K., Musen, M.A.: Comparison of ontology-based semantic-similarity measures. In: AMIA (2008)
7. Lin, D., et al.: An information-theoretic definition of similarity. In: ICML. vol. 98, pp. 296–304. Citeseer (1998)
8. Lord, P.W., Stevens, R.D., Brass, A., Goble, C.A.: Semantic similarity measures as tools for exploring the gene ontology. In: Pacific symposium on biocomputing. vol. 8, pp. 601–612 (2003)
9. Ontañón, S., Plaza, E.: Similarity measures over refinement graphs. Machine Learning Journal 87(1), 57–92 (2012)
10. Peñaloza, R.: Axiom pinpointing in description logics and beyond. Ph.D. thesis, Dresden University of Technology (2009)
11. Ran, B., Duimering, P.R.: Conceptual combination: Models, theories and controversies. International Journal of Cognitive Linguistics 1(1), 65–90 (2010)
12. Resnik, P.: Using information content to evaluate semantic similarity in a taxonomy. arXiv preprint cmp-lg/9511007 (1995)
13. Schorlemmer, M., Confalonieri, R., Plaza, E.: Coherent concept invention. In: Proceedings of the Workshop on Computational Creativity, Concept Invention, and General Intelligence (C3GI 2016) (2016)
14. Thagard, P.: Coherent and creative conceptual combinations. In: Creative thought: An investigation of conceptual structures and processes, pp. 129–141. American Psychological Association (1997)
15. Thagard, P.: Coherence in thought and action. Life and Mind: Philosophical Issues in Biology and Psychology, The MIT Press (2000)