

A Case-Based Approach to Mutual Adaptation of Taxonomic Ontologies

Sergio Manzano¹, Santiago Ontañón², and Enric Plaza¹

¹ IIIA, Artificial Intelligence Research Institute
CSIC, Spanish Council for Scientific Research
Campus UAB, 08193 Bellaterra, Catalonia (Spain)

{sergio,enric}@iia.csic.es

² Computer Science Department,

Drexel University
Philadelphia, 19104 PA, USA
santi@cs.drexel.edu

Abstract. We present a general framework for addressing the problem of semantic intelligibility among artificial agents based on concepts integral to the case-based reasoning research program. For this purpose, we define case-based semiotics (CBS) (based on the well known notion of the semiotic triangle) as the model that defines semantic intelligibility. We show how traditional CBR notions like transformational adaptation can be used in the problem of two agents achieving mutual intelligibility over a collection of concepts (defined in CBS).

1 Introduction

We propose an approach based on case-based semiotics (CBS) to determine problems in consistency or ambiguity based on the well known notion of the semiotic triangle. This approach aims at supporting the participating agents in evolving their individual ontologies on-demand, in a way that is enough to coordinate their activity in a particular task or subdomain. This participatory ontology is understood as an adaptation of the individual ontologies that converges into a shared mutually consistent and unambiguous ontology guided by our case-based approach to semiotics. Our approach is based on two basic assumptions:

(i) *Case-based Assumption:* participating agents share their environment and are capable of understanding their case description language(s). They either (1) share the case description language or (2) they share some basic ontology and language that allows them to explain their case description language(s)¹.

(ii) *Taxonomy Assumption:* Concepts in an ontology are organized in a hierarchy. More complex structures of ontologies are left for future work. In particular, DL-based ontologies require further development of inductive generalization methods before supporting this kind of approach.

This work is a generalization of the *concept convergence* approach [6], in which two agents deliberate about the meaning of a concept using their case-bases in

¹ How to achieve (2) is beyond the scope of this paper, but see [9].

order to achieve a shared, agreed-upon meaning of a concept. This generalization is due to the fact that concepts do not exist in isolation, but are related to other neighboring concepts in what we currently call an ontology. Concept convergence introduced the use of the semiotic triangle (see Fig. 1) to define concept meaning in a case-based agent. Our view with respect to case-based semiotics (CBS) is that *specific cases* are needed to perform certain forms of reasoning. This paper, in particular, focuses on the problem of mutual intelligibility for artificial agents endowed with a domain ontology. We think that the process by which two agents can adapt their ontologies to achieve mutual intelligibility requires *reasoning about cases*. Specifically, we think that a purely logic-based approach is not sufficient, and that a view of concept meaning based on classical logical semantics is not sufficient. This issue is a long-standing philosophical debate between the logic-based semantic view of meaning and the semiotic view of meaning (see e.g. [2]). We propose that concept meaning is better modeled by the semiotics approach that has a two-layer description of concepts: the intentional description or definition of a concept (in some formalism) and an extensional description of a concept (as is classically used in CBR).

Although this paper does not deal with *case-based problem solving* (in which new problems are solved using precedents or solved cases), our approach explicitly deals with case-based reasoning in the general sense: performing intelligent tasks by *reasoning about cases*. Moreover, we will show how mutual intelligibility about concepts can be seen as a process of mutual adaptation of taxonomies, using a process that is equivalent to transformational adaptation.

2 Background and Related Work

Most approaches use the *ontology alignment* metaphor to deal with the relationship between two different ontologies; it's a metaphor in that it originates by analogy with molecular sequence alignment [3]. Intuitively, ontology alignment (or matching) is a process that aims at finding "classes of data" that are *semantically equivalent*. Ontology alignment has been studied on database schemas, XML schemas, taxonomies, formal languages, entity-relationship models, and dictionaries. Formally, while matching is the process of finding relationships or correspondences between entities of different ontologies, alignment is a set of correspondences between two (or more) ontologies (by analogy with molecular sequence alignment) [3]. Thus, the alignment is the output of the matching process, which is very similar (in a conceptual sense) to partial matching in CBR retrieval and to structure-mapping in analogical mechanisms. Notice also that ontology alignment is different from ontology merging: ontology merging takes as input two (or more) source ontologies and returns a merged ontology based on the given source ontologies.

There are two families of approaches to ontology alignment, commonly called syntactic and semantic approaches. Syntactic approaches establish matchings among predicates, terms or other structural properties of a formalism, essentially focusing on a notion of similarity. Semantic approaches establish logical

equivalence correspondences among ontology terms, essentially focusing on a notion of semantic equivalence—in the logical sense of “semantic”. We propose a third approach, a semiotic viewpoint that takes into account both the extensional and intensional definitions of a concept. Related to our approach are methods that work on “populated ontologies,” i.e. ontologies that also contain instances of their concepts. Some approaches use instances to compute similarities among them in order to help them determine which concepts match. Although this is related to CBR, this is not the path taken here.

Another related approach is [8, 7], where a combination of Formal Concept Analysis (FCA) with Information Flow models for modeling and sharing common semantics is proposed. Their use of FCA is interestingly related with the approach taken here by the case-based semiotics for representing concepts. FCA has a two-layer representation of concepts, as we have in CBS with the intensional level and the extensional, that in FCA are called the intent and extent respectively of the concept. FCA, however, works only on attribute-value representations of instances and the intensional representations are subsets of attribute-value pairs, while our approach is more general, only requiring a representation formalism that has the subsumption operation. Similarly, FCA-merge [10] uses FCA over a common set of shared instances to merge two ontologies expressed as FCA lattices. Finally, “mutual online ontology alignment” [11] uses clustering and interchange of cases, but only uses the extensional description of concepts, while [1] proposes a similarity that takes into account the three dimensions of the semiotic triangle (see Fig. 1).

3 Adapting Taxonomies

A well known tenet in CBR is that after partial matching (i.e. retrieval), we need to adapt what’s matched (because it is only *partially* matched) in order to reuse it for some purpose. Thus, while ontology alignment/matching is related to CBR retrieval and to structure-mapping in analogical mechanisms, the partiality of this process requires a second process: adaptation for reuse. This is the focus of this paper: retrieve for reuse, and in particular mutual adaptation of ontologies for reaching a shared, participatory ontology (or more precisely a fragment of an ontology).

In particular, we envision a context-dependent mutual adaptation of two ontologies held by two agents. These two agents aim at performing a particular task or goal, which defines a context in which (part of) their taxonomies need to be mutually intelligible. This does not mean an agent has to modify forever its ontology, only create a modified version for working within a particular context. Our approach can be summarized in the following schema:

$$O_1 \leftrightarrow T_1 \xrightarrow{\text{adapt}} T'_1 \Leftrightarrow T'_2 \xleftarrow{\text{adapt}} T_2 \leftrightarrow O_2$$

where two agents, with ontologies O_1 and O_2 , in order to perform some task, select (\leftrightarrow) a segment of their ontologies (T_1 and T_2) as relevant to the task, and




Fig. 1. The classic semiotic triangle on the left and a CBS concept C as a triplet $\langle \lambda, I, E \rangle$ on the right

then need to create two adapted versions of these segments ($T_1 \xrightarrow{adapt} T'_1$), such that are mutually intelligible ($T'_1 \rightleftharpoons T'_2$). In this view, the agents do not renounce or change their core ontologies, but they are capable of adapting (segments of) them to a particular context. In this paper we will focus on the adaptation process, assuming the agents are capable of previously agreeing on the context (i.e. the goal to achieve and the part of the ontology that is relevant).

We will propose a transformational adaptation approach to achieve mutually intelligible ontologies, but this approach has some limitations. Specifically, we will encompass only hierarchical ontologies (henceforth taxonomies). Moreover, while denotational semantics are commonly used in logic-based ontologies, we will propose a case-based approach to defining meaning and mutual intelligibility of concepts. This approach, based on the semiotics approach to meaning, takes into account not only the “abstract” definition of a concept but also the “experiences” with concrete episodes where this concept is used. The next subsection presents this case-based semiotics (CBS) approach to meaning and mutual intelligibility of concepts.

3.1 CBS Taxonomies

Our representation of hierarchical ontologies (taxonomies for short) is based on the semiotic triangle for concepts. We will define a CBS concept for a language \mathcal{L} that possesses a subsumption relation (\sqsubseteq) among \mathcal{L} 's formulas. The language describing cases, without loss of generality, will be sublanguage $\mathcal{L}_c \subseteq \mathcal{L}$.

Definition 1. (CBS Concept) A CBS concept C is a triplet $\langle \lambda, I, E \rangle$, given a signature $\langle \mathcal{L}, \sqsubseteq \rangle$ and a set of labels Λ , where:

1. $\lambda \in \Lambda$; where λ is a label (the name for the concept) from the set of labels Λ ,
2. $I \in \mathcal{L}$ (I is a formula in a the language \mathcal{L} ; where I is called the intensional definition of λ , also noted as $I(C)$),
3. $E = \{e_1, \dots, e_n\}$ is a non-empty set of cases such that $\forall e_i \in E : e_i \in \mathcal{L}_c$; where the set E is called the extensional definition of λ , also noted as $E(C)$, and
4. $\forall e_i \in E : I \sqsubseteq e_i$




Fig. 2. For a taxonomy with , A and two children B and C , the intensional relations are shown at the left while the extensional relations are shown at the right

That is to say, a concept in CBS has a name, an intensional definition (that is a formula in some language), and a set of cases belonging to that concept (the extensional definition of that concept). For simplicity, we will sometimes denote a concept triplet by a symbol $C = \langle \lambda, I, E \rangle$ and we will use $I(C)$ and $E(C)$ to denote its intensional (I) and extensional definitions (E).

However, for an ontology we will need a discriminant definition; for this purpose we will use the notion of contrast set. We will say that a concept C is defined over a set of cases E whenever $E(C) \subseteq E$.

Definition 2. (*Contrast Set*) Given a set of cases $E = \{e_1, \dots, e_n\}$ we say a set of concepts (C_1, \dots, C_m) defined over E is a contrast set whenever $\forall i = 1, \dots, m$:

$$\left(\forall e_j \in E(C_i) : I(C_i) \sqsubseteq e_j \right) \wedge \left(\forall e_k \in \bigcup_{j=1, \dots, m, j \neq i} E(C_j) : I(C_i) \not\sqsubseteq e_k \right)$$

That is to say, a case in E belongs (is subsumed by) at most one concept in the contrast set (C_1, \dots, C_m) . However, not all cases need be members of a concept, which requires the contrast set to be a partition.

Definition 3. (*Conceptual Partition*) A contrast set (C_1, \dots, C_m) defined over a set of cases E is a conceptual partition $\Pi((C_1, \dots, C_m), E)$ iff $\forall e_i \in E, \exists C_j : I(C_j) \sqsubseteq e_i$.

That is to say, a conceptual partition of a set of cases is an exhaustive classification of the set of cases where all cases belong to only one of the concepts.

We turn now to define a CBS hierarchical ontology (or CBS taxonomy for short). The taxonomy of concepts can be seen as a tree where nodes are CBS concepts. More formally, a taxonomy is an *arborescence*, i.e. a directed graph in which, for a vertex x called the *root* and any other vertex y , there is exactly one directed path from x to y . We will denote an arborescence by $\langle \mathbf{C}, \mathbf{A} \rangle$ where \mathbf{C} is a set of nodes (or vertices) and \mathbf{A} is set of arcs; for a $C \in \mathbf{C}$, we will denote the children of C as $A(C)$.

Definition 4. (*CBS Taxonomy*) Given a collection of concepts $\mathbf{C} = \{C_1, \dots, C_m\}$ defined over a set of cases $E = \{e_1, \dots, e_n\}$, and an arborescence $\langle \mathbf{C}, \mathbf{A} \rangle$ with root C_1 , the triple $\langle \mathbf{C}, \mathbf{A}, E \rangle$ is a CBS Taxonomy whenever:




Fig. 3. Example of a concept B' in taxonomy T' converging with respect to the concept B in taxonomy T

1. $E(C_1) = E$ and $\forall e_j \in E : I(C_1) \sqsubseteq e_j$ (root is sound and complete w.r.t. E)
2. $\forall C_i \in \mathbf{C}' : \Pi(A(C_i), E(C_i))$ is a conceptual partition,
3. $\forall C_i \in \mathbf{C}' \wedge \forall C_j \in A(C_i) : I(C_i) \sqsubset I(C_j)$ (intensional subsumption)

where $\mathbf{C}' \subset \mathbf{C}$ is the set of non-terminal concepts in the taxonomy.

Fig. 2 shows some of these properties in a small example of taxonomy with root A and two children B and C . The subsumption relation is established among intensional descriptions of concepts, while extensional descriptions are related by set inclusion. Moreover, the two concepts at the same level, B and C , form a partition upon the cases in the extension of its father A .

Two concepts labels are *mutually intelligible* (a.k.a. aligned) when their CBS concepts converge. Conceptual convergence is defined as follows.

Definition 5. (CBS Concept Convergence) Two CBS concepts C_i and C_j belonging to conceptual partitions $\Pi(\mathbf{C}_i, E_i)$ and $\Pi(\mathbf{C}_j, E_j)$ in taxonomies T_i and T_j respectively, and with C_i^\blacktriangle and C_j^\blacktriangle the parents of C_i and C_j converge with respect to taxonomy T_i whenever:

1. $\forall e \in E(C_i) : I(C_j) \sqsubseteq e$
2. $\forall e \in E(C_i), K \in \mathbf{C}_j - \{C_j\} : I(K) \not\sqsubseteq e$
3. $\forall e \in E(C_i) : I(C_j^\blacktriangle) \sqsubseteq e$

When the dual properties of 1 to 3 are satisfied, C_i and C_j converge with respect to T_j . When C_i and C_j converge w.r.t. both T_i and T_j , we say C_i and C_j are conceptually convergent, noted as $(C_i \cong C_j)$. Moreover, we say their labels are mutually intelligible $(\lambda_i \leftrightarrow \lambda_j)$ for T_i and T_j .

Property 1 states that C_j is consistent with C_i 's extensional description, Property 2 that partition $\Pi(\mathbf{C}_j, E_j)$ is consistent with C_i 's extensional description, and Property 3 that C_j 's parent is consistent with C_i 's extensional description.

Thus, convergence of two concepts occurs when both concepts converge with respect to the *other* taxonomy. Figure 3 shows an example of a concept B' in taxonomy T' converging with respect to the taxonomy T . Intuitively, the example in Fig. 3 means that $B' \sqsubseteq B$, since B' covers the cases in $E(B)$ and none of the cases in the remaining concepts of the partition under A . Thus, if two agents




Fig. 4. Two adaptation operations over Hierarchical Ontologies: *Identification* (a case is identified as belonging to concept C) and *Categorization* (a case is identified as belonging to a new, previously non-existent, concept D)

Ag and Ag' are using taxonomies T and T' respectively, we say concept B is *intelligible for agent Ag'* — in the sense that there will be no misunderstanding or disagreement for agent Ag' with respect to concept B of agent Ag . Definition 5 states that two concepts converge w.r.t. both T and T' we have both $B' \sqsubseteq B$ and $B \sqsubseteq B'$, and thus they are equivalent ($B \cong B'$) w.r.t. to CBS. Consequently, when two agents communicate with each other using their labels ($\lambda \leftrightarrow \lambda'$) for the “same concept,” their usage will be mutually intelligible.

Notice however, that the equivalence ($C_i \cong C_j$) in Definition 5 does not mean they are logically equivalent; what is assured is that they are equivalent w.r.t. the known set of known cases relevant to C_i and C_j , namely $E(C_i^\Delta) \cup E(C_j^\Delta)$ (the set of observed cases in the contrasts sets to which C_i and C_j belong). Indeed, previously unseen cases can be identified or not as belonging to (C_i or C_j), leading to a disagreement that would require *adapting* again their taxonomies.

Thus, ontology matching and convergence is an evolving process according to case-based semiotics. Any agreement on the meaning of a sign or label is first participatory (applying to the involved agents) and contextual (depending on the finite knowledge of the world of the agents expressed as the set of cases grounding the concept’s meaning). Finally, notice that our form of concept *alignment* is that of concepts being mutually intelligible w.r.t. to CBS. Thus, the alignment of two ontologies is to be defined as convergence of their concepts.

Definition 6. (*CBS Taxonomy Convergence*) Two taxonomies $\langle \mathbf{C}, \mathbf{A}, E \rangle$ and $\langle \mathbf{C}', \mathbf{A}', E' \rangle$ with roots A and A' are CBS-convergent whenever $\forall C \in \mathbf{C}, \exists C' \in \mathbf{C}'$ such that $C \cong C'$ and $\forall C' \in \mathbf{C}', \exists C \in \mathbf{C}$ such that $C \cong C'$.

3.2 Adaptation Operators

We will define several operations of transformational adaptation over the space of hierarchical ontologies. These operations are Identification, Categorization, Split, and Merge, and are shown in Figures 4, 5 and 6. These operators are similar to (and inspired from) the ones on the CobWeb unsupervised learning system [4], the main difference being derived from our distinction between cases (at the extensional level) and concepts (at the intensional level), which is nonexistent in CobWeb. Figure 4 shows on the left a new case X that is already classified




Fig. 5. The adaptation operation *Split*: concept *C* is “split” into its subconcepts that are promoted to the higher level, while the case *X* is later identified to one of the promoted concepts




Fig. 6. The adaptation operation *Merge*: concepts *C* and *D* are “merged” into a higher level super-concept *E* and they are demoted to the lower level, while the case *X* is yet to be processed below the new concept *E*

as being a member of concept *A*; in other words $I(A) \sqsubseteq X$ (the intensional definition of *A* subsumes *X*). Applying the operator Identification we obtain the tree shown in the middle of Fig. 4. This operator characterizes the situation where a new case is identified as member of a concept (e.g. the concept *C*) with no further change required except maybe to generalize to insure $I(C) \sqsubseteq X$. However, consider the case where generalizing $I(C)$ to include *X* would mean that $I(C)$ also subsumes cases under *B*; this means that *X* cannot be identified as member of *C*. If this is also the case for *B* then *X* cannot be identified as member of *B* or *C* and (as shown at the right of Fig. 4) we need a new concept, let’s call it *D*, that encompasses *X*. This situation is characterized by the operator Categorize, that creates a “new category” for a case *X*. Thus, the result of operator Categorize is moving from a partition (*B*, *C*) of the extension of *A* to a partition (*B*, *C*, *D*).

4 Mutual Adaptation of Taxonomies

Two agents communicate and deliberate about the meaning of their taxonomies, or more specifically, about a fragment of their taxonomies starting from a common root. If the root under discussion is the taxonomies root then the agents will deliberate about the meaning of all concepts in their taxonomies. For this purpose, agents need to recognize situations where there is no agreement and then apply some adaptation operators. This approach is similar to goal-driven learning (GDL) [5], Goal-driven learning decomposes the learning problem in

three steps: blame assignment, learning goal generation, and repair (or learning) strategy. GDL considers a single agent reasoning introspectively about detecting its own failures (blame assignment), deciding what needs to be learnt to correct it (learning goal generation), and determining a way to achieve this goal (repair strategy).

4.1 Non-structural Adaptations

In non-structural adaptations, disagreements involve mismatches between intensional and extensional definitions that do not require transforming the *is-a* relationship between concepts (as can be seen in [6]).

Generalization. This situation is characterized as follows: agent Ag_1 has a case X subsumed by concept B , while agent Ag_2 has a concept B' that subsumes most cases in B but not X . Moreover, the other concepts in the partition K' where B' is located in T' do not subsume X either. Thus, the partition K' does not account for X and since Ag_1 knows it should be covered by B' , Ag_2 should change the definition of B' . Therefore Ag_1 sends the argument “your concept B' should also cover X ” to Ag_2 . Then, Ag_2 generalizes $I(B')$ to cover X while not covering any case subsumed by the other concepts in partition K' .

Specialization. This situation is characterized as follows: agent Ag_1 has a case X subsumed by concept B , while agent Ag_2 has a concept C' different from B' that does subsume X . Since Ag_1 current hypothesis is that B and B' should converge while B and C' should not, Ag_2 should change the definition of C' . Therefore Ag_1 sends the argument “your concept C' should not cover X , which should be covered by B' ” to Ag_2 . Consequently, Ag_2 has to specialize concept’s intension $I(C')$ so that C' it does no longer cover X . Additionally, B' may or may not cover X . If not, Ag_2 generalizes $I(B')$ to cover X while not covering any case subsumed by the other concepts in partition K' .

4.2 Structural Adaptations

Structural adaptations are triggered by mismatches in the way the cases are sorted by partitions, and require the transformation of partitions; that is to say transforming the tree of *is-a* relationships among concepts, including the creation of new concepts. Let us start with the second situation in Fig. 4: categorization. Let us assume, e.g., an agent Ag_2 sends case X to agent Ag_1 and eventually Ag_1 ’s ontology is adapted by including a new concept D to encompass case X .

The scenario requires some initial conditions, as follows. Assume the agents have already achieved concept convergence over the root A ; therefore both agree that X (sent from Ag_2 to Ag_1) can be identified as belonging to A . However, they do not agree under which concept, in the partition set under A , should case X be allocated. In order to move from the left part of Fig. 4 to the right part, the agents have to agree on the following: a) X is not under B (or B' for the second agent); and X is not under C (or C') either. Thus, since X should be under A but is not part of B or C , a new concept is needed for Ag_1 ; since Ag_2 already

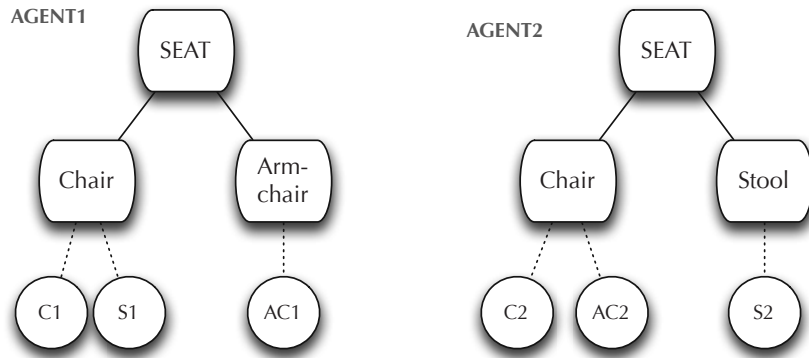


Fig. 7. The initial state of two agents taxonomies in the *Seat* domain




Fig. 8. The two taxonomies in the *Seat* domain

has identified X under a concept D' (with label λ'_D), agent Ag_1 will create this new concept D (with label λ_D), and X will be situated under D . Since they are mutually intelligible ($\lambda_D \leftrightarrow \lambda'_D$) the adaptation process ends there.

To explain the Split adaptation operator we will introduce an example shown in Fig. 7. The *Seat* domain is very simple but will illustrate the kind of mismatches that can be found and resolved by mutual adaptation. Agent Ag_1 in Fig. 7 knows two kind of seats: chairs and armchairs, while agent Ag_2 knows two kind of seats: chairs and stools. Ag_1 divides seats depending on whether they have arms or not, while Ag_2 divides seats depending on whether they have backs or not, as shown in Fig. 8². Notice that both agents have cases that are stools ($S1$ and $S2$), chairs ($C1$ and $C2$) and armchairs ($AC1$ and $AC2$); they just choose to conceptualize them differently.

We can easily imagine a convergence of both ontologies into one shared by both agents and that has the three involved concepts: stools, chairs and armchairs. Indeed, we will follow the deliberation and adaptations that achieve that, but notice that the particular solution achieved is not unique. As we will show, the agents reach an ontology with *Seat* as root and three children (*Stool*, *Chair* and *Armchair*); nevertheless, other ontology structures are possible and correct results, for instance an ontology where *Seat* is the root with children *Stool* and *Chair*, in which *Chair* has two children concepts: *Armchair* and *SimpleChair* (i.e. a seat with back and no arms).

² The second ontology is the one found in the Wikipedia (armchair is a subtype of chair) while some of the authors claim they feel more intuitive the first one is more intuitive, and to classify an armchair not as a chair and see it as a kind couch.

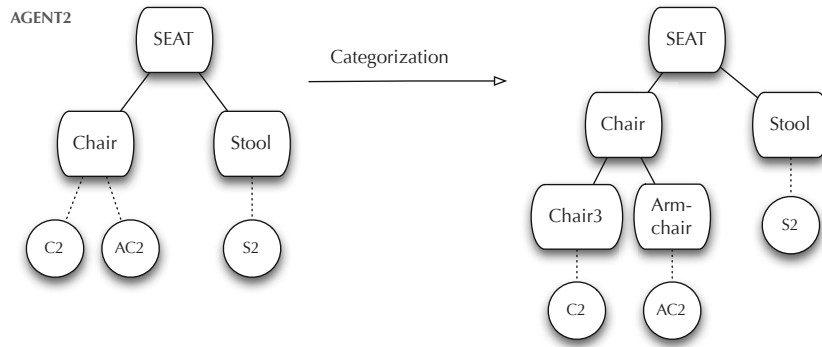


Fig. 9. Adaptation of the taxonomy of agent Ag_2 by adding concept $Armchair$

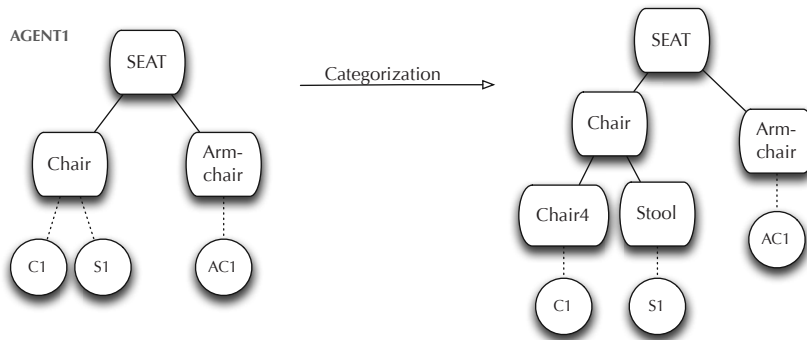


Fig. 10. Adaptation of the taxonomy of agent Ag_1 by adding concept $Stool$

Agent Ag_1 in Fig. 7 has received the intensional definition of concepts $Chair_2$ and $Stool_2$ from Ag_2 , and agent Ag_2 the intensional description of Ag_1 's concepts. Considering first Ag_1 , the agent has found the following disagreements:

- 1) $Stool_2$ covers case $S1$ that is covered by the intensional definition of concept $Chair_1$; thus a chair like $S1$ is a stool of Ag_2 , a concept not existing in Ab_1
- 2) $Chair_2$ covers case $AC1$ that, according to Ag_1 , is not a chair but is under concept $Armchair_1$, a concept that is not present in Ag_2 's taxonomy.

In order to proceed, Ag_1 asks Ag_2 to create and include the concept $Armchair$ in its taxonomy. Ag_2 accepts, which implies the following:

- 1) a new concept using the intensional definition of $Armchair_1$ has to be created, and call it $Armchair_2$ (thus $I(Armchair_2) := I(Armchair_1)$)
- 2) Ag_2 determines that $Armchair_2$ covers case $AC2$ but not case $C2$, thus the adaptation operation Categorization can be applied to concept $Chair_2$ creating $Armchair_2$ as a subconcept of $Chair_2$,
- 3) however, now the children of $Chair$ (case $C2$ and $Armchair_2$) do not form a partition (since case $C2$ is not a concept). Thus a new concept $Chair_3$ is created to cover case $C2$; as shown in Fig. 9 the children of $Chair$ now form a partition.

Notice that in the example we show only one case per concept just for brevity's sake. In general, adding $Armchair_2$ under $Chair_2$ would mean that all cases subsumed by (the intensional definition of) $Chair_2$ that are also subsumed by

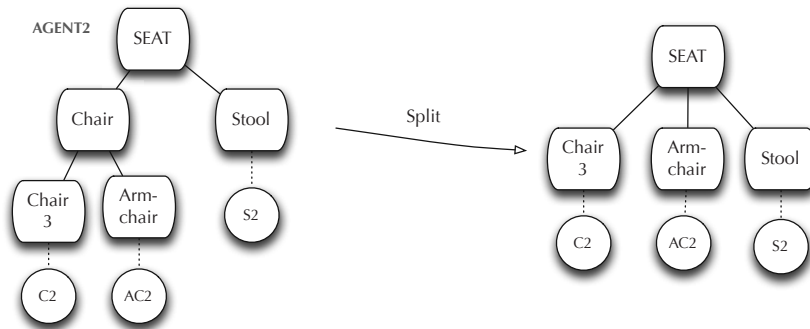


Fig. 11. Adaptation of the taxonomy of agent *Ag2* by splitting the old concept *Chair* and promoting concept *Chair3* and *Armchair* as subconcepts of *Seat*

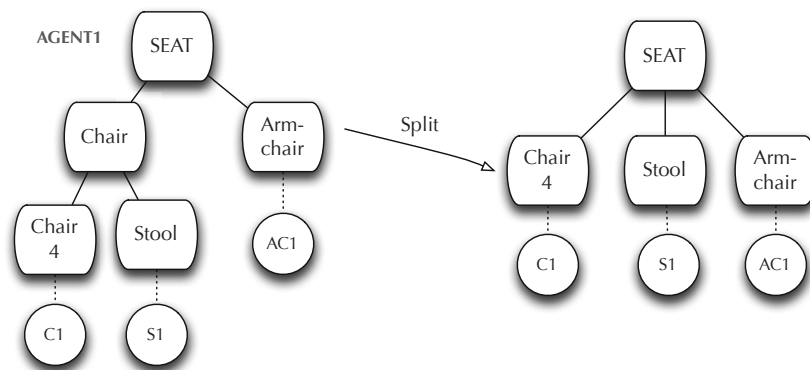


Fig. 12. Adaptation of the taxonomy of agent *Ag1* by splitting the old concept *Chair* and promoting concept *Chair3* and *Armchair* as subconcepts of *Seat*

(the intensional definition of) *Armchair₂* become the extensional definition of *Armchair₂*, i.e. $E(Armchair_2) = \{c \in E(Chair_2) | I(Armchair_2) \sqsubseteq c\}$, while the rest become the extensional definition for a new concept to complete the partition: $E(Chair_3) = E(Chair_1) - E(Armchair_2)$. Finally, the intensional definition $I(Chair_3)$ of the new concept is inferred by induction over the cases of the extensional definition $E(Chair_3)$.

A similar process is carried out when agent *Ag₂* asks *Ag₁* to include the *Stool* concept, as shown in Fig. 10. Now both agents have incorporated a new concept coming from the other agent refining their respective ontologies. However, as can be observed comparing Fig. 9 and Fig. 10 their ontologies do not match: although their lower level concepts converge (since *Chair*, *Stool* and *Armchair* partition the extensional definition of the overall concept *Seat* the same way), the intermediate concepts (the “old” concepts of *Chair* in both agents) do not converge. This disagreement can be resolved applying adaptation operator split (Fig. 5) to the “old” concepts of *Chair* in both agents. Figures 11 and 12 show that the same result is obtained by both agents using the split operation.

Finally, the Merge adaptation operation works in a similar way to Split. Recalling Fig. 6, we see Merge would be applied when one agent has an intermediate concept that the other has not. We will not develop the example in full, but it is easy to see how merge can be used in the example of Figure 13. Given the state of agents *Ag₁* and *Ag₂* in Fig. 13, when *Ag₂* applies Merge to concepts *Chair*

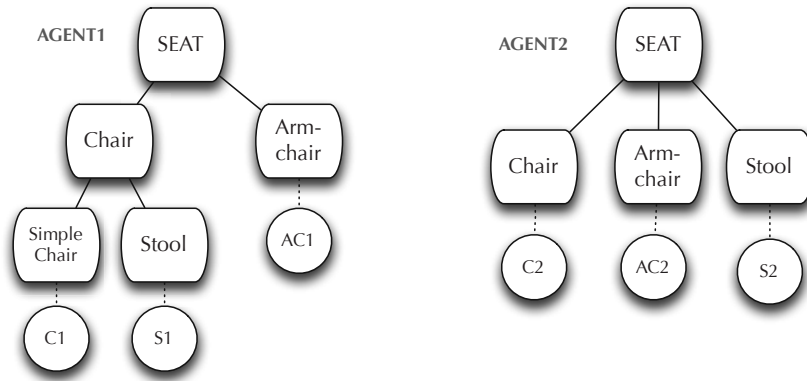


Fig. 13. An state where the Merge operator would make two taxonomies convergent

and *Armchair* creating a new superconcept *NewChair* the two taxonomies converge. Specifically, they have the following alignments: (*Chair* \leftrightarrow *NewChair*), (*SimpleChair* \leftrightarrow *Chair*), (*Stool* \leftrightarrow *Stool*), (*Armchair* \leftrightarrow *Armchair*). Clearly, this is not the only configuration that leads to a convergence. A second, equivalent solution is that agent Ag_1 applies Split to *Chair* (promoting *SimpleChair* and *Stool* to the level of *Armchair*) thus reaching a taxonomy convergent with that of Ag_2 . Both solutions are equally adequate from the point of view of CBS.

5 Mutual Adaptation as Search

The CBS approach allows us to characterize (1) disagreements in the intended meaning of concepts in two taxonomies and (2) the transformations upon ontologies performed by adaptation operations. Thus, mutual adaptation of ontologies is viewed as a search process over the space of possible taxonomies under case-based semiotics. We say that two concepts from T and T' are in coincidence when, although they do not converge, they both subsume a subset of the cases subsumed by the other.

Definition 7. (*CBS Coincident Concepts*) Two CBS concepts C_i and C_j in taxonomies T and T' respectively, and with parents C_i^\blacktriangle and C_j^\blacktriangle such that $C_i^\blacktriangle \cong C_j^\blacktriangle$ are in coincidence ($C_i \equiv C_j$) whenever $\exists K_i \subseteq E(C_i), K_j \subseteq E(C_j)$ (with $K_i \neq \emptyset$ and $K_j \neq \emptyset$) such that $I(C_i) \sqsubseteq K_j$ and $I(C_j) \sqsubseteq K_i$.

Two concepts that are in coincidence are basically candidates for converging if the current disagreement or mismatches are solved by adaptation operators. The search process maintains a list of coincident concept pairs and constitute the candidates to which adaptation operators can be applied.

Figure 14 and Figure 15 show some examples of the CBS typology of disagreements for taxonomies to which some adaptation operators may be applied. For instance, Fig. 14 shows on the left the situation where a case X' of taxonomy T' is covered by concept A in T ($I(A) \sqsubseteq X'$) but none of the concepts in the conceptual partition (B, C) cover X' (i.e. $I(B) \not\sqsubseteq X'$ and $I(C) \not\sqsubseteq X'$). This may be due to two different situations depending on X' in taxonomy T' , shown




Fig. 14. A type of concept disagreement in which a case X' of T' is not covered by a conceptual partition in T




Fig. 15. A type of concept disagreement in which a case X' of T' is covered by a different concept in T

to the right of Fig. 14: either (1) X' is covered by a concept, say B' and thus $B \not\cong B'$, or (2) X' is covered by concept that does not exist in T , say D' in T' . To solve this disagreement and achieve convergence, in situation (1) the intensional definition is changed using the Generalization adaptation operation on B , while in situation (2) the Categorization adaptation operation is used to include a new concept D in the conceptual partition. Another instance of disagreement is shown in Fig. 15, where a case X' that belongs to concept B' in T' is however covered in taxonomy T by a concept that is not the coincident concept B .

6 Discussion

We have presented a general framework for addressing the problem of semantic intelligibility among artificial agents based on concepts integral to the case-based reasoning research program. Mutual intelligibility of concepts should be grounded, in our approach, to collections of cases (i.e. descriptions of objects or situations). Using a semiotic viewpoint instead of a classical logic semantics allows us to work with cases in a principled way, that we have formalized as CBS (case-based semiotics), in which a concept has a label and two (mutually dependent) levels of description: the intensional level and the extensional level.

Mutual intelligibility of concepts is moreover modeled as a process of mutual adaptation, in which artificial agents modify their knowledge structures to reach a convergent model (in the CBS framework) of the concepts they need to share. This mutual adaptation process is viewed as a search process performed by adaptation operators, as is classically obtained by transformational adaptation in CBR. However, the situation is here more complex than in classical transformational adaptation, since there are two agents involved. A particular

interaction protocol to implementing search in the space of possible taxonomies remains future work, although the adaptation operators that define the search space have already been defined here.

Finally, the CBS framework allows the acquisition of new concepts in a natural way. A new concept implies either the reorganization of the partition of the cases known to an agent or the acquisition of a new, unknown case. In the CBS approach, learning from cases and adapting the knowledge structure commonly called “ontology” are seamlessly integrated in the same process. As part of the future work we intend to show that two agents using our adaptation operators can always converge on a shared taxonomy, even when they have concepts and cases unknown to one another.

Acknowledgments. This research was partially supported by projects Next-CBR (TIN2009-13692-C03-01) and Agreement Technologies (CONSOLIDER CSD2007-0022).

References

- [1] Aimé, X., Furst, F., Kuntz, P., Trichet, F.: SemioSem: A Semiotic-Based Similarity Measure. In: Meersman, R., Herrero, P., Dillon, T. (eds.) OTM 2009 Workshops. LNCS, vol. 5872, pp. 584–593. Springer, Heidelberg (2009)
- [2] Eco, U.: *I limiti dell’interpretazione*. Bompiani (1990)
- [3] Euzenat, J., Shvaiko, P.: *Ontology Matching*. Springer (2007)
- [4] Fisher, D.H.: Knowledge acquisition via incremental conceptual clustering. *Machine Learning* 2(2), 139–172 (1987)
- [5] Leake, D.B., Ram, A. (eds.): *Goal-Driven Learning*. MIT Press (1995)
- [6] Ontañón, S., Plaza, E.: Concept convergence in empirical domains. In: Pfahringer, B., Holmes, G., Hoffmann, A. (eds.) DS 2010. LNCS (LNAI), vol. 6332, pp. 281–295. Springer, Heidelberg (2010)
- [7] Schorlemmer, M., Kalfoglou, Y.: Progressive ontology alignment for meaning coordination: An information-theoretic foundation. In: Proc. AAMAS 2005, pp. 737–744. ACM Press (2005)
- [8] Schorlemmer, M., Kalfoglou, Y., Atencia, M.: A formal foundation for ontology-alignment interaction models. *International Journal on Semantic Web and Information Systems* 3(2), 50–68 (2007)
- [9] Steels, L.: Why We Need Evolutionary Semantics. In: Bach, J., Edelkamp, S. (eds.) KI 2011. LNCS, vol. 7006, pp. 14–25. Springer, Heidelberg (2011)
- [10] Stumme, G., Maedche, A.: FCA-MERGE: Bottom-up merging of ontologies. In: Proc. IJCAI 2001, pp. 225–230. Morgan Kaufmann Publishers Inc. (2001)
- [11] Wang, J., Gasser, L.: Mutual online ontology alignment. In: Proc. AAMAS 2002 Workshop on Ontologies in Agent Systems, OAS 2002. CEUR Workshop Proceedings, vol. 66 (1990)