Epistemic Alignment Repair

Line van den Berg[1]

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, F-38000 Grenoble France
`line.van-den-berg@inria.fr`

Abstract. Ontology alignments enable interoperability between heterogeneous information resources. The Alignment Repair Game (ARG) specifically provides a way for agents to simultaneously communicate and improve the alignment when a communication failure occurs. This is achieved through applying adaptation operators that provide a revision strategy for agents to resolve failures with minimum information loss. In this paper, we explore how closely these operators resemble logical dynamics. We develop a variant of Dynamic Epistemic Logic called DEOL to capture the dynamics of ARG by modeling ontologies as knowledge and alignments as belief with respect to the plausibility relation. The dynamics of ARG are then achieved through announcements and conservative upgrades. With the representation of ARG in DEOL, we formally establish the limitations and the redundancy of the adaptation operators. More precisely, that for a complete logical reasoner, `replace`, `addjoin` and `refine` are redundant for one or both agents in the game and that `add` would be replaced by `addjoin` in all cases.

Keywords: Ontology alignment · Alignment repair · Dynamic Epistemic Logic

1   Introduction

Ontology alignments are a widely-used tool to facilitate interoperability while preserving heterogeneity of information [7]. Ontology matching is the process of determining the correspondences of the alignment. Many different matching algorithms have been developed [13]. These do however not suffice in any realistic context of multi-agent systems where agents have access to different ontologies and are required to communicate with incomplete alignments or ontologies may be adopted to new information in the environment. This requires a dynamic view on matching algorithms instead of the originally static approach where the alignment is determined initially and provided as input to the communicating agents.

The dynamic perspective has led to various proposals for ontology matching on multi-agent systems. Amongst others: gossiping [1], logical repair [12] and conservativity principles [9]. These proposals have been integrated within multi-agent systems via specific protocols [1,10]. In this paper, we focus on a specific approach to the dynamic perspective to evolve alignments *through their use*: the Alignment Repair Game (ARG) [5,6]. In this approach, which is based

on ideas of cultural evolution [14], agents communicate and overcome failures in communication (due to an incorrect correspondence) by repairing the alignment. This repair is done via adaptation operators that specify how agents adopt their knowledge in case of an incorrect correspondence with minimum information loss. Because this repair is executed during communication, ARG avoids the problematic assumption that alignments need to be determined before any communication can take place.

In this paper we explore the connection between the adaptation operators and logical dynamics. We introduce a variant of Dynamic Epistemic Logic called Dynamic Epistemic Ontology Logic (DEOL) to capture the dynamics of ARG by modeling ontologies as knowledge and alignments as belief with respect to the plausibility relation. This logic is an extension of work on Modal Description Logics, see [3] for more information. The communication in ARG is achieved through announcements and conservative upgrades. With this representation of ARG in DEOL, we formally establish the limitations and the redundancy of the adaptation operators in failing to capture all the available information in ARG and supplying information that is already known or believed by agents. More precisely, that for a complete logical reasoner, the adaptation operators `replace`, `addjoin` and `refine` are redundant for one or both agents in the game and that adaptation operator `add` would be replaced by `addjoin` to preserve consistency. Even though these results do heavily rely on the reasoning capacities of agents, they do not depend on any specific feature of the representation in DEOL. In fact, any logic whose ontology semantics is that of OWL will guarantee the results.

That said, the precise proposal in this paper has its own limitations. We use strong requirements on DEOL models that facilitate some form of maximal consistency over the interpretation function for classes. Furthermore, our focus is that of a logical understanding of ARG and hence a detailed investigation of the specific properties and validities of the proposed logic still remains to complete.

2   Alignment Repair Game

An ontology typically provides a vocabulary of a domain of interest and a specification of the meaning of terms in that vocabulary via semantic relations such as specialization ($\sqsubseteq$), equivalence ($\equiv$), exclusion ($\oplus$) and membership ($\in$) [7]. Formally, an ontology is a knowledge base (a TBox) in description logic. For precise definitions, consult [2].

The intuition we use in this paper is that an ontology is a quintuple $\mathcal{O} = \langle \mathcal{D}, \mathcal{C}, \sqsubseteq, \oplus, \in \rangle$ where $\mathcal{D}$ is a set of object names, $\mathcal{C}$ is a non-empty set of class names, $\sqsubseteq, \equiv, \oplus \subseteq \mathcal{C} \times \mathcal{C}$ are the semantic relations and $\in \subseteq \mathcal{D} \times \mathcal{C}$ is the membership relation. To give meaning to these relations an interpretation $I = \langle \Delta, \cdot^I \rangle$ is given for an ontology $\mathcal{O}$ that provides a domain Δ and a function $\cdot^I$ assigning to objects $o \in \mathcal{D}$ an interpretation in the domain Δ and to each class $C \in \mathcal{C}$ a set of objects of the domain [7]. We then say that *"C is subsumed by D"* ($C \sqsubseteq D$) iff $C^I \subseteq D^I$, *"C and D are equivalent"* ($C \equiv D$) iff $C \sqsubseteq D$ and $D \sqsubseteq C$ and *"C and*

D are disjoint" $(C \oplus D)$ iff $C^I \cap D^I = \varnothing$. Lastly we say that *"o is a member of C"* $(o \in C)$ iff $o^I \in C^I$. For two classes C, D that are not disjoint we also write $C \between D$ and in each ontology $\top$ is the class such that $\top^I = \Delta$. From classes C, D, we also form the classes $C \sqcup D$, $C \sqcap D$ and $\neg C$ that represent their union, intersection and complement. To denote an ontology (without an interpretation) we also say the *signature* of an ontology. The signature of $\mathcal{O} = \langle \mathcal{D}, \mathcal{C}, \sqsubseteq, \oplus, \in \rangle$ is the set of class names $C \in \mathcal{C}$ and object names $o \in \mathcal{D}$, without any specification of the relations that hold between them.

An alignment A_{ab} between two ontologies $\mathcal{O}_a$, $\mathcal{O}_b$ is a set of correspondences between classes of the two ontologies [7]. Such a correspondence is a triple $\langle C_a, C_b, R \rangle$ where C_a, C_b are classes of $\mathcal{O}_a, \mathcal{O}_b$, respectively, and $R \in \{\sqsubseteq, \sqsupseteq, \equiv, \oplus\}$ is a semantic relation that is asserted to hold between C_a and C_b. We also write $C_a R C_b$ for $\langle C_a, C_b, R \rangle$.

The Alignment Repair Game (ARG) is a game designed for artificial agents to evolve alignments between their ontologies through their use [5,6]. In ARG, agents are given individual (private) ontologies, a set of randomly generated (possibly incorrect) alignments and a common domain. Each round of ARG is played privately between two agents and is described in Definition 1. The aim of ARG is to improve communication between agents by repairing the alignment when failures occur. Through application of the adaptation operators the alignment between the ontologies evolves and converges towards a state closer to the *reference alignment*, the maximal truthful alignment between the ontologies that serves to determine the performance of the operators. Note that the reference alignment is not available to the agents themselves.

For simplicity, in ARG it is assumed that the objects of the domain are described by a finite set of Boolean features (more specifically, by the presence or absence of each feature) and that ontologies are strict hierarchies in which each level adds one constraint (positive or negative) related to exactly one feature. This means that classes not in subsumption relation are disjoint. Ontologies are assumed to be incomplete by having one level less than the environment has features and classes in the ontology correspond to the conjunction of the features of its ancestors. For instance, the bottom-leftmost class in Example 1 is defined by $Square_a \sqcap Small_a$.

Definition 1 (Alignment Repair Game (ARG)). *The* Alignment Repair Game *(ARG) is played by a set of agents $\mathcal{A}$ with a common domain $\mathcal{D}$ of objects. Each agent $a \in \mathcal{A}$ is associated with an ontology $\mathcal{O}_a$, an interpretation I_a for this ontology and a randomly generated set of shared alignments A_{ab} is given between any two ontologies $\mathcal{O}_a$, $\mathcal{O}_b$ that at least includes $\top_a \equiv \top_b$. We write $D_i^o \in \mathcal{O}_i$ for the most specific class ($\sqsubseteq$-wise) of object $o \in \mathcal{D}$ available in $\mathcal{O}_i$ and we assume that D_a^o, D_b^o are not part of the initial alignment.*

At each round of the game:

1. *Two agents $a, b \in \mathcal{A}$ and an object $o \in \Delta$ are picked at random.*
2. *Agent a asks agent b to which class the object o belongs and agent b answers $C_b \in \mathcal{O}_b$ s.t. $D_b^o \sqsubseteq C_b$ and $C_a R C_b \in A_{ab}$ with $R \in \{\sqsupseteq, \equiv\}$.*

3. *Agent a compares C_a ($C_a R C_b \in A_{ab}$) with D_a^o.*
4. *If $D_a^o \sqsubseteq C_a$ belongs to $\mathcal{O}_a$ then the round is a success.*
5. *Else (then by construction $D_a^o \oplus C_a$ belongs to $\mathcal{O}_a$) the round is a failure and an adaptation operator is applied to the alignment A_{ab}.*

ARG consists of a fixed number of rounds as described above for a chosen operator and then the alignments are compared to the reference alignments. By playing ARG with different operators, their effects can be compared.

Note that because of the assumptions that we imposed on our ontologies, the most specific class D_a^o of object o in ontology $\mathcal{O}_a$ is always unique. For an intuition, consider the following example.

Example 1. Let agent a and agent b play ARG where their ontologies $\mathcal{O}_a$ and $\mathcal{O}_b$ are described in Figure 1 with the domain $\Delta = \{\square, \triangle, \blacksquare, \blacktriangle, \square, \triangle, \blacksquare, \blacktriangle\}$. The initial alignment A_{ab} is represented by the dotted red correspondences between classes of their ontologies. Now, if the object $\blacktriangle$ is drawn in a round, this round will be considered a success, whereas the object $\triangle$ would yield a failure and an adaptation operator will be applied.


Fig. 1: *The ontologies (black) of agent a (left) and agent b (right) and the alignment (red, dashed) between them of Example 1.*

The repair in case of a failure is performed via adaptation operators [5,6]. These operators provide a strategy for the agents to revise the failed correspondence.

Definition 2 (Adaptation Operator). *We have the following* Adaptation Operators *for ARG (for a schematic view of the effect of the operators, refer to Figure 2):*

- `delete`*: delete the initial correspondence;*
- `replace`*: replace the initial correspondence by $C_a \sqsubseteq C_b$;*
- `add`*: in addition to* `replace`*, add a $\sqsubseteq$-correspondence between C_b and the immediate superclass C_a' of C_a: $C_b \sqsubseteq C_a'$;*
- `addjoin`*: in addition to* `replace`*, add a $\sqsubseteq$-correspondence between C_b and the lowest superclass C_a' of C_a that is compatible with D_a^o (i.e. $D_a^o \sqsubseteq C_a'$): $C_b \sqsubseteq C_a'$;*

- refine: *in addition to* replace, *add $\sqsubseteq$-correspondences between the sub-classes C_b' of C_b that do not subsume the actual class of the object (i.e. $D_b^o \not\sqsubseteq C_b'$): $C_b' \sqsubseteq C_a$;*
- refadd: *the combination of* addjoin *and* refine.

Note that the order of the actions that are performed by the adaptation operators does not matter.

Fig. 2: *Schematic diagram of the deleted (red) and added correspondences (blue, green) by the different the adaptation operators in ARG [5,6].*

3   Dynamic Epistemic Ontology Logic

We model the game in a version of Dynamic Epistemic Logic where the propositions are elements of a Description Logic language: object classifications ($C(x)$) and class relations ($C \equiv D$, $C \sqsubseteq D$ and $C \oplus D$). We call this extension (multi-agent) Dynamic Epistemic Ontology Logic (DEOL).

Definition 3 (Syntax of DEOL). *The* syntax *of (multi-agent) DEOL is defined in the following way.*

$$\phi ::= C(x) \mid \top(x) \mid CRD \mid \phi \wedge \psi \mid \neg\phi \mid K_i\phi \mid B_i\phi \mid [\dagger\phi]\psi$$

$$R \in \{\sqsubseteq, \equiv, \oplus\}, \ \dagger \in \{!, \uparrow\}$$

Where $C, \top \in \mathcal{C}$, K_i and B_i are the knowledge and belief operators for each agent i, $\dagger\phi$ with $\dagger \in \{!, \uparrow\}$ are dynamic upgrades and x is a variable.

The connectives $\rightarrow$ and $\vee$ and the duals $\hat{K}_i, \hat{B}_i, \langle\dagger\phi\rangle$ are defined in the usual way. Models of DEOL are plausibility models.

Definition 4 (DEOL Model). *A* model *of (multi-agent) DEOL is a 6-tuple $\mathfrak{M} = \langle W, (\geqslant_i)_{i\in\mathcal{A}}, w^*, \Delta, I \rangle$ where*

- *W is the set of states, or worlds;*

- $(\geqslant_i)_{i\in\mathcal{A}}$ *are the plausibility relations on* W*, one for each agent, that are converse well-founded, locally connected preorders;*
- w^* *is the actual world;*
- Δ *is the domain of interpretation (a set of objects);*
- *and* I *is an* interpretation function *that assigns to each state* w *a function* $\cdot^{I_w}$ *interpreting each class name* $C \in \mathcal{C}$ *as a set of objects of the domain of interpretation, i.e.* $C^{I_w} \subseteq \Delta$*, such that* $\top^{I_w} = \Delta$*,* $(C \sqcap D)^{I_w} = C^{I_w} \cap D^{I_w}$ *and* $(\neg C)^{I_w} = \Delta \setminus C^{I_w}$ *for each* $w \in W$*.*

We also write $C \sqcup D$ *for the class defined by* $\neg(\neg C \sqcap \neg D)$*,* $\sqcap\{C_i\}$ *for the class defined by* $C_1 \sqcap C_2 \sqcap \ldots$ *and* $\sqcup\{C_i\}$ *for the class defined by* $C_1 \sqcup C_2 \sqcup \ldots$*, and their interpretations at* w *are given by* $C^{I_w} \cup D^{I_w}$*,* $\bigcap C_i^{I_w}$ *and* $\bigcup C_i^{I_w}$*, respectively. In each DEOL model* $\bot := \neg\top$ *is the empty class.*

The plausibility relation $w \geqslant_i v$ reads as "w is at least as plausible as v for agent i". From this, we can define the epistemic and doxastic relations on W as follows:

$$w \sim_i v := w \ (\leqslant_i \cup \geqslant_i) \ v \tag{1}$$

$$w \to_i v := v \in Max_{\leqslant_i} w(i) \tag{2}$$

Where $w(i) := \{v \in W \mid w \sim_i v\}$ is the *information cell* of agent i at state w. It follows from the properties of $\leqslant_i$ and $\geqslant_i$ that the relations $\sim_i$ are reflexive, transitive and symmetric, and the relations $\to_i$ are transitive, serial and Euclidean. Therefore they satisfy the usual properties of knowledge and belief, respectively.

Definition 5 (Semantics of DEOL). *The* semantics *of DEOL is defined in the following way:*

- $\mathcal{M}, w \vDash C(x)$ *iff* $x \in C^{I_w}$
- $\mathcal{M}, w \vDash C \sqsubseteq D$ *iff* $C^{I_w} \subseteq D^{I_w}$
- $\mathcal{M}, w \vDash C \equiv D$ *iff* $C^{I_w} = D^{I_w}$
- $\mathcal{M}, w \vDash C \oplus D$ *iff* $C^{I_w} \cap D^{I_w} = \varnothing$
- $\mathcal{M}, w \vDash \phi \wedge \psi$ *iff* $\mathcal{M}, w \vDash \phi$ *and* $\mathcal{M}, w \vDash \psi$
- $\mathcal{M}, w \vDash \neg\phi$ *iff* $\mathcal{M}, w \nvDash \phi$
- $\mathcal{M}, w \vDash K_i\phi$ *iff* $\forall v$ *s.t.* $w \sim_i v$: $\mathcal{M}, v \vDash \phi$
- $\mathcal{M}, w \vDash B_i\phi$ *iff* $\forall v$ *s.t.* $w \to_i v$: $\mathcal{M}, v \vDash \phi$
- $\mathcal{M}, w \vDash [!\phi]\psi$ *iff* $\mathcal{M}^{!\phi}, w \vDash \psi$
- $\mathcal{M}, w \vDash [\uparrow \phi]\psi$ *iff* $\mathcal{M}^{\uparrow\phi}, w \vDash \psi$

Where $!\phi$ *and* $\uparrow \phi$ *act as model transformers* $!\phi := \mathcal{M} \to \mathcal{M}^{!\phi}$ *and* $\uparrow \phi := \mathcal{M} \to \mathcal{M}^{\uparrow\phi}$ *in the following ways, where* $||\phi||_{\mathcal{M}} := \{w \in W \mid \mathcal{M}, w \vDash \phi\}$*:*

- ***Announcement*** $!\phi$*: Delete all '*$\neg\phi$*'-worlds from the model. I.e.* $W^{!\phi} = ||\phi||_{\mathcal{M}}$ *and* $w \geqslant_i^{!\phi} v$ *iff* $w \geqslant_i v$ *and* $w, v \in W^{!\phi}$*.* w^**,* Δ*,* $\mathcal{C}$ *and* I *remain unchanged;*
- ***Conservative upgrade*** $\uparrow \phi$*: Change the plausibility orders so that the best '*ϕ*'-worlds become better than all other worlds, while the old ordering on the rest of the worlds remains. I.e.* $W^{\uparrow\phi} = W$ *and* $w \geqslant_i^{\uparrow\phi} v$ *iff either* $v \in Max_{\leqslant_i}(w(i) \cap ||\phi||_{\mathcal{M}})$ *or* $w \geqslant_i v$*. Again,* w^**,* Δ*,* $\mathcal{C}$ *and* I *remain unchanged.*

The intuition behind these transformations is that the trustworthiness of new information may vary: whereas an announcement is considered as new information from an infallible source, conservative upgrades consider information from sources that are trusted, but that are not infallible. Precisely for this reason conservative upgrades only change the plausibility of worlds without deleting any alternatives.

Note that in both cases w^* remains the actual world of the model. This also means that an announcement $!\phi$ can only be validly performed on a model $\mathcal{M}$ if ϕ is true is the actual world w^*. For more technical details about announcements and conservative upgrades consult [11,4,15,16].

4   ARG Dynamics

We now formalize ARG in DEOL and model the changes in the epistemic-doxastic state of the agents by means of upgrades. The DEOL model $\mathcal{M}$ that describes the initial epistemic-doxastic states of the agents satisfies the following conditions:

- The ontology $\mathcal{O}_a$ is *known* to agent a. That is, the sentences of DEOL that describe $\mathcal{O}_a$ are true in any world accessible by a from the actual world w^* via $\sim_a$;
- The alignment A_{ab} between two ontologies is *believed* by agents a, b. That is, the sentences of DEOL that describe A_{ab} are true in any worlds accessible by a, b from the actual world w^* via $\rightarrow_a$, $\rightarrow_b$;
- The signatures of every agent's ontology is known to every agent (public signature assumption).

For instance, in Example 1 the sentences of DEOL describing ontology $\mathcal{O}_a$ are $(Square_a \sqcap Small_a) \sqsubseteq Small_a$, $Small_a \subseteq \top_a$, $Triangle_a \sqcap Small_a(\blacktriangle)$, $Triangle_a \sqcap Small_a(\triangle)$, $(Triangle_a \sqcap Small_a) \sqsubseteq Small_a$, etc, and the sentences describing the alignment A_{ab} are $Small_a \equiv Black_a$ and $\top_a \equiv \top_b$.

To satisfy the public signature assumption, we require that for each two classes C, D in any other agent's ontology that do not appear in the alignment, agent a considers all combinations of the following alternatives equally plausible (for each $o \in \Delta$):

- $C(o)$ / $\neg C(o)$ - $C \equiv D$ / $C \not\equiv D$ - $C \sqsupseteq D$ / $C \not\sqsupseteq D$
- $D(o)$ / $\neg D(o)$ - $C \sqsubseteq D$ / $C \not\sqsubseteq D$ - $C \oplus D$ / $C \,\slashed{\oplus}\, D$

This is achieved in models of DEOL by ultimately making as many copies of the worlds describing an agent's knowledge of her own ontology and belief of alignments involving her ontology as there are combinations of the alternatives above, ranking them all equally plausible while respecting the order imposed by the alignments. Such a model very rapidly explodes and therefore we often only draw the epistemic-doxastic states of agents. Note that to satisfy the public signature assumption the first requirement already takes care of agents' own

Fig. 3: *Initial knowledge (solid black lines) and belief (dashed red lines) of agent b in Example 1.*

signature and the second requirement of classes appearing in the alignment. In Figure 3, the epistemic-doxastic state of agent b in Example 1 is given.

As a result, both ontologies and alignments are captured in the epistemic-doxastic states of agents in the DEOL model. However, now that agents can reason logically, there is more. Agents can combine their knowledge and beliefs to obtain new beliefs. For instance, $K_a(C_a(o))$ and $B_a(C_a \equiv C_b)$ entails $B_a(C_b(o))$, etc. In other words, everything agent a knows about $C_a \in \mathcal{O}_a$ she also believes about $C_b \in \mathcal{O}_b$ whenever $C_a \equiv C_b$ is part of the alignment A_{ab}, i.e. whenever she believes that C_a and C_b are the same. This is already satisfied in the DEOL model by the closure properties of the interpretation function.

During ARG new information is learned. There are two dynamic acts involved in the learning: the communication of $C_b(o)$ in step 2 of ARG and the adaptation operator applied in step 5 (see Definition 1). How do these acts change the knowledge and beliefs of the agents? And are the adaptation operators as defined by [5,6] necessary to account for these changes?

In order to answer these questions, we formalize the communication acts in ARG with the dynamic upgrades available in our model: announcements $!\phi$ and conservative upgrades $\uparrow \phi$.

Definition 6 (ARG Dynamics in DEOL). *We model each round of ARG as defined in Definition 1 by*

$$!C_b(o); \ if \ D_a^o \oplus C_a \ then \ \texttt{operator} \tag{3}$$

Where operator *denotes the applicable operator and D_a^o was defined as the actual (most specific) class of $o \in \Delta$ in $\mathcal{O}_a$.*

It is clear that the communication of the class of the object should be captured by a announcement. Step 3 of ARG is then the model transformation induced by $!C_b(o)$: $\mathcal{M}$ to $\mathcal{M}^{!C_b(o)}$. Yet, for the adaptation operators, announcements are

not the right tool: these are a form of a *revision policy* and tell the agents how to change the alignment upon a failure in communication. Therefore the effect of the operators should not take place at the level of knowledge but at the level of belief. For this we use conservative upgrades.

Definition 7 (Adaptation Operators as Dynamic Modalities). *Given the failure of the correspondence* $C_b \sqsubseteq C_a (\wedge C_b \sqsupseteq C_a) \in A_{ab}$, *we can model the dynamics of the adaptation operators in the following way (recall that D_a^o was defined as the actual (most specific) class of $o \in \Delta$ in $\mathcal{O}_a$):*

- $\texttt{replace:} \uparrow (C_a \not\sqsupseteq C_b)$
- $\texttt{add:} \uparrow (C_a \not\sqsupseteq C_b \wedge C_b \sqsubseteq C_a^{sup})$
- $\texttt{addjoin:} \uparrow (C_a \not\sqsupseteq C_b \wedge C_b \sqsubseteq C_a^{supD})$
- $\texttt{refine:} \uparrow (C_a \not\sqsupseteq C_b \wedge \sqcup \mathbf{C}_b^{sub} \sqsubseteq C_a)$
- $\texttt{refadd:} \uparrow (C_a \not\sqsupseteq C_b \wedge C_b \sqsubseteq C_a^{supD} \wedge \sqcup \mathbf{C}_b^{sub} \sqsubseteq C_a)$

Where $C_a^{sup} = Min_{\sqsubseteq}\{C \in \mathcal{O}_a \mid C_a \sqsubseteq C\}$, $C_a^{supD} = Min_{\sqsubseteq}\{C \in \mathcal{O}_a \mid C_a \sqsubsetneq C \wedge D_a \sqsubseteq C\}$ *(and by construction of the ontologies, C_a^{sup} and C_a^{supD} are unique) and* $\mathbf{C}_b^{sub} = \{C_b^{sub} \in \mathcal{O}_b \mid C_b^{sub} \sqsubseteq C_b \wedge D_b^o \not\sqsubseteq C_b^{sub}\}$.

5   Redundancy Results

For logical agents, some of the operators of ARG are redundant or even superfluous. In particular, we show redundancy of the operators $\texttt{replace}$ and $\texttt{addjoin}$ with respect to the epistemic-doxastic states of (one of) the agents in $\mathcal{M}^{!C_b(o)}$. We then prove why the operator $\texttt{add}$ should not be considered an option and discuss what the effect is of the representation of ARG in DEOL to arrive at these results. In particular, we explore which properties of DEOL enforce complete reasoning of agents.

Proposition 1 (Redundancy). *The adaptation operator $\texttt{replace}$ is already entailed by the logic. That is, $\mathcal{M}^{!C_b(o)}$ is bisimilar to $\mathcal{M}^{!C_b(o);\texttt{replace}}$. Furthermore, the operators $\texttt{addjoin}$ and $\texttt{refine}$ do not alter the epistemic-doxastic state of agent a and b, respectively.*

Proof. The proof for $\texttt{replace}$ is straightforward as the initial correspondence (belief) $C_a \sqsupseteq C_b (\wedge C_a \sqsubseteq C_b) \in A_{ab}$ is already altered by $!C_b(o)$ in case of a failure to $C_a \sqsupseteq C_b \wedge C_a \not\sqsubseteq C_b \in \mathcal{A}_{ab}^{!C_b(o)}$. For the proof of $\texttt{addjoin}$, we compare the epistemic-doxastic state of agent a before and after the announcement $!C_b(o)$, see Figure 4. The proof of $\texttt{refine}$ is analogous.

Note that through the announcement $!C_b(o)$, agent a does not only revise the alignment, but also acquires new knowledge. This is because as a model transformer $!C_b(o)$ deletes all the worlds in which $C_b(o)$ is false and as a result, agent a comes to know that it is true. Proposition 1 does not address the adaptation operator $\texttt{add}$ because applying $\texttt{add}$ may be inconsistent for agent a.

Proposition 2. *The adaptation operator* `add` *may be inconsistent with the knowledge and belief of agent a and when it is consistent,* `add` *is equivalent to* `addjoin`.

Proof. Assume that $!C_b(o)$ is announced s.t. $D_a^o \oplus C_a$ with $C_b \sqsubseteq C_a \in A_{ab}$ and the adaptation operator `add` is applied. That means that the correspondence $C_b \sqsubseteq C_a$ is deleted and $C_b \sqsubseteq C_a^{sup}$ is added where C_a^{sup} is the immediate superclass of C_a in $\mathcal{O}_a$. However, there is no evidence that C_a^{sup} is compatible with object o: it cannot be ensured that $C_a^{sup} \between D_a^o$. In fact, whenever $C_a^{sup} \between D_a^o$ holds $C_a^{sup} = C_a^{supD}$ (as $D_a^o \sqsubseteq C_a^{sup}$) and therefore this operation is equivalent to `addjoin`. Whenever $C_a^{sup} \oplus D_a^o$ it is contradictory to add $C_b \sqsubseteq C_a^{sup}$ as agent a knows that $C_b(o)$ and $\neg C_a^{sup}(o)$.

Proposition 2 is in line with initial predictions and experimental results[1]: `addjoin` shows faster convergence than `add` because `add` may add false correspondences that later need revision.

In a way, Proposition 1 and Proposition 2 follow from the representation of ARG in DEOL. This is because DEOL enforces a greater reasoning capacity of agents compared to the original setting of ARG. In models of DEOL some closure properties with respect to agents' epistemic-doxastic state are already embedded as a consequence of the transitive plausibility relations, the definition of knowledge as truth in all epistemically accessible worlds and the requirement that the interpretation function satisfies ($C \sqcap D)^{I_w} = C^{I_w} \cap D^{I_w}$ and $(\neg C)^{I_w} = \Delta \backslash C^{I_w}$. Together, the implementation of these definitions result in agents being complete reasoners over the domain of interpretation. Resultingly the adaptation operators that intend to expand the interpretation

Fig. 4: *The knowledge (solid black) and belief (dashed red) of agent a before (above) and after (below) the announcement* $!C_b(o)$ *is made.*

of classes, `replace`, `addjoin` (for agent a) and `refine` (for agent b), naturally become redundant.

The results presented here rely on the reasoning capacity of agents that are a consequence of the semantics of DEOL. However, they do not rely on any specific feature of DEOL. In fact, any logic whose ontology semantics is that of OWL would warrant it because $D_a^o \oplus C_a$, $C_b(o)$, $D_a^o(o)$ imply that $\neg C_a(o)$ and hence $C_a \nsqsubseteq C_b$, etc. Hence, the conclusions about the redundancy of the

---

[1]Available at `https://gforge.inria.fr/plugins/mediawiki/wiki/lazylav/index.php/20180826-NOOR`

operators `replace`, `addjoin` and `refine` and the possible inconsistency of `add` will still hold.

6   Conclusion

We have explored the connection between the adaptation operators in ARG and the logical dynamics in DEOL where ontologies are modeled as knowledge and alignments as belief with respect to the plausibility relation. The dynamics of ARG are achieved through announcements and conservative upgrades in DEOL and we have shown that the changes induced are purely epistemic-doxastic. Whereas the adaptation operators only target the incorrect correspondence in ARG, there is more information available to acquire. This means that the adaptation operators have a limitation regarding the information that can be learned by the agents. Yet, on the other hand, the operators are redundant whenever their task is only to extend the interpretation of concepts. In DEOL, this extension is already entailed and causes, together with the closure properties that follow from the transitivity of the plausibility relations and knowledge as truth in all accessible worlds, agents to be complete reasoners.

More research is needed to address this matter of complete reasoners. Weakening some of the closure properties might yield different results. The main question is then, to what extend do the results presented in this paper still hold and how can these closure properties be weakened to resemble reasoning capacities closer to the original state of ARG? Likewise, now that the basis for a logical analysis of ARG is laid, we can explore whether new operators can be constructed for ARG from the DEOL setting. Exploring more of the properties and validities of the proposed logic would facilitate this purpose.

The results in this paper are not symmetric: `addjoin` is redundant for agent a and not for agent b and vice versa for `refine`. This is because one agent may not know the (super-, sub)class that is used to repair the alignment nor knows the relation between this class and the initial aligned classes. While the latter is a direct consequence of the design of the gameplay of ARG, the former was dealt with by requiring the public signature assumption that makes agents aware of all the available classes. With this assumption the conservative upgrades rightfully capture the adaptation operators and the results of this paper follow. However, the public signature assumption does not allow ontologies to dynamically evolve. A more dynamic alternative solution could be to consider partial valuations. Partial valuations allow agents to extend their vocabularies upon learning. Compared to the current approach, this approach is a more natural way to model unawareness of agents because it allows for dynamic open models [8]. In addition, introducing partial valuations might well lead to a formal characterization of the features expansion and relaxation as introduced in [6]. These features allow agents to introduce new random correspondences (expansion) or to use shadowed correspondences (relaxation), correspondences that are (possibly incorrect) yet cannot be detected as such because they do not cause

any failure [6]. Further research is required to explore the precise potential of partial valuations for modeling ARG and knowledge representations.

Acknowledgement

References

1. Aberer, K., Cudré-Mauroux, P., Hauswirth, M.: Start making sense: The chatty web approach for global semantic agreements. Web Semantics: Science, Services and Agents on the World Wide Web **1**(1), 89–114 (2003)
2. Baader, F., Calvanese, D., McGuinness, D., Patel-Schneider, P., Nardi, D.: The description logic handbook: Theory, implementation and applications. Cambridge university press (2003)
3. Baader, F., Küsters, R., Wolter, F.: Extensions to description logics. In: The description logic handbook. pp. 219–261. Cambridge University Press (2003)
4. Baltag, A., Moss, L.S., Solecki, S.: The logic of public announcements, common knowledge, and private suspicions. In: Proceedings of the 7th conference on theoretical aspects of rationality and knowledge (TARK 98). pp. 43–56 (1998)
5. Euzenat, J.: First experiments in cultural alignment repair (extended version). In: European Semantic Web Conference. pp. 115–130. Springer (2014)
6. Euzenat, J.: Interaction-based ontology alignment repair with expansion and relaxation. In: 26th International Joint Conference on Artificial Intelligence (IJCAI). pp. 185–191. AAAI Press (2017)
7. Euzenat, J., Shvaiko, P.: Ontology Matching. Springer (2013)
8. Jaspars, J., Thijsse, E.: Fundamentals of partial modal logic. Studies in Logic Language and Information (1996)
9. Jiménez-Ruiz, E., Payne, T.R., Solimando, A., Tamma, V.: Limiting logical violations in ontology alignnment through negotiation. In: Fifteenth International Conference on the Principles of Knowledge Representation and Reasoning (2016)
10. Payne, T.R., Tamma, V.: Negotiating over ontological correspondences with asymmetric and incomplete knowledge. In: Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems. pp. 517–524. International Foundation for Autonomous Agents and Multiagent Systems (2014)
11. Plaza, J.: Logic of public communications. In: et al., M.E. (ed.) Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems. pp. 201–216. North-Holland (1989)
12. Santos, E., Faria, D., Pesquita, C., Couto, F.M.: Ontology alignment repair through modularization and confidence-based heuristics. PloS one **10**(12), e0144807 (2015)
13. Shvaiko, P., Euzenat, J.: Ontology matching: state of the art and future challenges. IEEE Transactions on knowledge and data engineering **25**(1), 158–176 (2013)
14. Steels, L.: Experiments in cultural language evolution, vol. 3. John Benjamins Publishing (2012)
15. Van Benthem, J.: Dynamic logic for belief revision. Journal of applied non-classical logics **17**(2), 129–155 (2007)
16. Van Benthem, J., Van Eijck, J., Kooi, B.: Logics of communication and change. Information and computation **204**(11), 1620–1662 (2006)