# Avoiding Alignment-based Conservativity Violations through Dialogue

Ernesto Jiménez-Ruiz[1], Terry R. Payne[2], Alessandro Solimando[3], and Valentina Tamma[2]

[1] Department of Computer Science, University of Oxford, UK
[2] Department of Computer Science, University of Liverpool, UK
[3] DIBRIS, Università di Genova, Italy

**Abstract.** A number of ontology matching techniques have been proposed that rely on full disclosure of their ontological models prior to the construction of the alignment. However, within open and opportunistic environments, such approaches may not always be pragmatic or even acceptable (due to privacy concerns). Several studies have focussed on collaborative, decentralised approaches to ontology alignment, where agents negotiate the acceptability of correspondences (i.e. mappings between corresponding entities in different ontologies) acquired from past encounters, or try to ascertain novel correspondences on the fly. However, such approaches can lead to logical flaws that may undermine their utility. In this paper, we extend a dialogical approach to correspondence negotiation, whereby agents not only exchange details of possible correspondences, but also identify potential violations to the so-called conservativity principle, where novel but undesirable entailments between named concepts in one of the input ontologies emerge. We present a formal model of the dialogue, and show how conservativity violations can be repaired (using an existing correspondence repair system) during the dialogue through the exchange of *repairs*. We then illustrate how agents negotiate over possible correspondences and repairs by means of a walkthrough example.

## 1 Introduction

For agents to behave appropriately in uncertain or unknown environments, they need an internal representation, or *world model*, of their perceptions of their environment. This representation is often defined within some logical theory (*ontology*) that is not necessarily shared (or only partially shared) with other agents, even though there may be common assumptions regarding how pertinent information and knowledge is modelled, expressed and interpreted. When interoperation between heterogeneous systems is required, an integration phase is necessary whereby different knowledge models are reconciled and possible implicit assumptions are clarified, especially for more dynamic and opportunistic scenarios (*e.g.*, e-commerce, open-data or mobile systems).

Traditionally, the challenge of resolving semantic heterogeneity has been addressed by aligning the agents' ontologies, using one of many existing alignment systems [2,13]; however, most rely on the ontologies to be fully shared with one of the agents, or with a third-party responsible for generating alignments. Furthermore, no single approach is necessarily suitable for all scenarios; and (partial) privacy has become increasingly pertinent, whereby neither agent or knowledge system is prepared to disclose its full ontology [12]; *e.g.*, if the knowledge encoded within an ontology was confidential or commercially sensitive.

The existence of pre-computed alignments has been exploited by several recent alignment approaches, which select and combine the constituent correspondences thus generating a new alignment [1,9,10,17]. However, certain correspondences may be found frequently by different alignment approaches, whereas others could be spurious or erroneous, and only appear rarely, resulting in different levels of confidence or *weight*. Furthermore, different alignment systems may map entities from one ontology to different entities in the other ontology, leading to *ambiguity*. Including such correspondences may be legitimate in certain scenarios; users may be familiar with the notion of synonyms or equivalent labels for certain concepts, and may not want to converge on a single, artificially canonical label. However, there is the danger that integrating such ambiguous correspondences within either ontology can lead to many undesirable logical consequences, and violate the three principles proposed in [7]: *consistency*, *locality*, and *conservativity*, whose satisfaction minimizes the number of potentially unintended consequences. Many alignment systems employ a *brute-force* approach to the selection of one-to-one, or *injective* correspondences to avoid ambiguity, through the identification of a *Matching* from the resulting weighted bipartite graph (obtained by mapping all the entities in the signatures of the two ontologies that are being aligned). This is done by finding either a maximum weighted bipartite matching which can be solved in polynomial time [8], or finding a lower complexity approximation of this based on stability [4]. Although these approaches are effective in eliminating ambiguous correspondences, they can also prune out potentially useful alternatives that satisfy the *conservativity principle*, where correspondences should not introduce new semantic relationships between concepts from one of the input ontologies.

In this paper, we extend an existing dialogical approach [12] to correspondence alignment, whereby agents not only exchange details of possible correspondences, but also identify and eliminate those potential correspondences that could yield conservativity violations by means of a detection and repair mechanism based on LogMap [6,14,15]. The approach assumes that the agents had acquired correspondences from past encounters, or from publicly available alignment systems (that were kept private), and that each agent associated some *weight* to each known correspondence. As this knowledge is *asymmetric* and *incomplete* (*i.e.*, neither agent involved in the dialogue is aware of all of the correspondences, and the weight assigned to each correspondence could vary greatly), the agents engage in the inquiry dialogue to: 1) ascertain the joint acceptability of each correspondence; and to 2) select a set of correspondences which reduce or eliminate the occurrence of possible conservativity violations (from each agent's individual perspective, rather than from a joint perspective). We present a formal model of the dialogue, and show how conservativity violations can be repaired (using an existing correspondence repair system) during the dialogue through the exchange of *repairs*, and present a walkthrough example of a dialogue.

The remainder of this paper is organised as follows: the new dialogue, based on the original *Correspondence Inclusion Dialogue*, is described in Section 2, and the repair mechanism is introduced in Section 3. A complete walkthrough is then provided to illustrate the dialogue (Section 4). Related work is summarised in Section 5, and the paper concludes in Section 6.

## 2 The Dialogue

In [12], the *Correspondence Inclusion Dialogue* (*CID*) was presented which enabled two agents to exchange knowledge about ontological correspondences resulting in an alignment that satisfies the following:

1. each agent is aware of a set of correspondences, each with an associated *weight*;
2. there should be no *ambiguity* with respect to either the source entities in the resulting alignment, or the target entities;
3. if there are alternative choices of correspondences, the selection is based on the combined, or *joint weight* of both agents;
4. that no correspondences should be selected where their joint weight is less than some defined *evidence-threshold*; and
5. the number of correspondences disclosed (*i.e.* whose weight is shared in the dialogue) should be *minimised*.

These conditions were extended in [11] to allow agents to *credulously* or *sceptically* accept correspondences that they had not hitherto encountered, which determined how the joint weight of correspondences was calculated (Subsection 2.3). The rationale behind the dialogue exploited the fact that whilst the agents involved sought to minimise the disclosure of their ontological knowledge (and thus the concepts known), some exchange of ontological knowledge (at least the exchange of a subset of candidate correspondences) was necessary for the determination of a consensual set of correspondences that formed the final alignment. Whilst it was assumed that the agents are inherently self interested, there was also the assumption that the agents were collaborative with respect to determining an alignment that could facilitate communication [5], as it was in the interest of all rational agents involved to be able to communicate successfully.

The dialogue has been significantly modified to retain the ability to negotiate over private beliefs of viable correspondences in such a way as to minimise the number of correspondences disclosed, but to replace the ambiguity mechanism (*i.e.*, removing the *object* move, and the use of argumentation) with a conservativity violation detection and repair mechanism (see Section 3). Furthermore, the formal treatment of correspondences and weights, as well as the syntax for the moves has been improved. The new dialogue is described in detail in the following subsections.

### 2.1 Ontologies, Correspondences and Weights

The agents negotiate over the viability of different correspondences that could be used to align the two agents' ontologies. The dialogue therefore assumes that each agent commits to an *ontology* $\mathcal{O}$, which is an explicit and formally defined vocabulary representing the agent's knowledge about the environment, and its background knowledge (domain knowledge, beliefs, tasks, etc.). $\mathcal{O}$ is modeled as a set of axioms describing

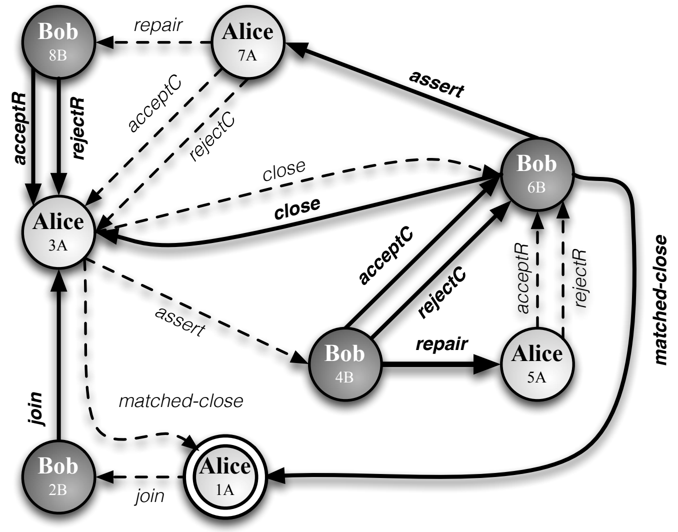


**Fig. 1.** The dialogue as a state diagram. Nodes indicate the agent whose turn it to utter a move. Moves uttered by *Alice* are labelled with a light font / edge, whereas those uttered by *Bob* are labelled with a heavy font / dashed edge. It assumes that *Alice* will always make the first move.

classes and the relations existing between them[4] and $\Sigma$ is the *ontology signature*; *i.e.*, the set of class and property names used in $\mathcal{O}$. To avoid confusion, the sender's ontology is denoted $\mathcal{O}^x$, whereas the recipient's ontology is $\mathcal{O}^{\hat{x}}$. For agents to interoperate in an encounter, they need to determine an *alignment* $\mathcal{A}$ between the two vocabulary fragments $\Sigma^x$ and $\Sigma^{\hat{x}}$ for that encounter. An alignment [3] consists of a set of *correspondences* that establish a logical relationship between the entities (classes, properties or roles, and instances) belonging to each of the two ontologies, and a set of logical relations. The universe of all possible correspondences is denoted $\mathcal{C}$. The aim of the dialogue is to generate an alignment $\mathcal{A} \subseteq \mathcal{C}$, that maps between the entities in $\Sigma^x$ and $\Sigma^{\hat{x}}$, that does not introduce any conservativity violations, and whose joint weight is at least as great as the admissibility threshold $\epsilon$.

***Definition 1:*** *A* **correspondence** *is a triple denoted* $c = \langle e, e', r \rangle$ *such that* $e \in \Sigma^x$, $e' \in \Sigma^{\hat{x}}$, $r \in \{\equiv, \sqsubseteq, \sqsupseteq\}$.

Agents associate a private, static *weight* $\kappa_c$ to a correspondence $c$ (where $0 \leq \kappa_c \leq 1$) that represents the confidence the agent has in the correctness of $c$. Each agent manages a private knowledge base, known as the *Correspondence Store* ($\Delta$), which stores the correspondences and their associated private weights, and a public knowledge base, or *Commitment Store*, $CS$, which contains a trace of all of the moves uttered by each agent [18]. Although each agent maintains its own copy of the $CS$, these will always be identical, and thus we do not distinguish between them. However, we do distinguish between the sender's Correspondence Store, $\Delta^x$, and the recipient's store, $\Delta^{\hat{x}}$.

[4] Here we restrict the ontology definition to classes and roles.

### 2.2 The Inquiry Dialogue Moves

The dialogue consists of a sequence of communicative acts, or *moves*, whereby agents take turns to assert the candidacy of some correspondence $c$ (and its associated *weight*) for inclusion in a mutually acceptable alignment, $\mathcal{A}$, and respond to such assertions by 1) confirming the acceptability of $c$ without the need for any alignment repair; 2) proposing a possible repair to $\mathcal{A}$ to allow $c$ to be added to $\mathcal{A}$ without introducing any conservativity violations (such as weakening or removing an existing correspondence); or 3) rejecting the acceptability of $c$. Each agent discloses its private *belief* regarding some correspondence $c$ and its weight, and the agents negotiate to rationally identify a set of mutually acceptable correspondences, given an *admissibility threshold* $\epsilon$. It assumes that only two agents (referred to as $Alice$ and $Bob$) participate in the dialogue, and that each agent plays a specific role (*i.e.*, an agent is either a *sender* $x$ or *recipient* $\hat{x}$) in any single dialogue move.

The set of possible *moves* $\mathcal{M}$ permitted by the dialogue are summarised in Table 1. The syntax of each move at time $s$ is of the form $m_s = \langle x, \tau, c, \kappa_c, \mathcal{R} \rangle$, where $x$ represents the identity of the agent making the move; $\tau \in \mathcal{M}$ represents the move type; $c$ is the subject of the move, *i.e.*, the correspondence that is being discussed; $\kappa_c$ represents either the personal or joint weight associated with $c$ where $0 \leq \kappa_c \leq 1$; and $\mathcal{R}$ represents a repair for correspondences within the candidate alignment $\mathcal{A}$ or the correspondence $c$ itself (described in Section 3). For some moves, it may not be necessary to specify a correspondence, weight or repair; in which case they will be empty or unspecified (represented as *nil*).

Agents take turns to utter *assert* moves (*i.e.*, to transition from state 3A for *Alice* or 6B for *Bob* in Figure 1). A sender $x$ can also make two consecutive moves in certain circumstances, such as an *accept* or *reject* move (see states labelled 7A for *Alice* and 4B for *Bob* in Figure 1), which enables an agent to accept or reject a disclosed correspondence or repair before making some other move.

### 2.3 Aggregating Weights and the Upper Bound

Within the dialogue, the agents try to ascertain the unambiguous, mutually acceptable correspondences to include in the final alignment $\mathcal{A}$ by selectively sharing those correspondences that are believed to have the highest weight. Once each agent knows of the other agent's weight for a given correspondence $c$, it can calculate $c$'s *joint weight*, and check if it is greater than or equal to the *admissibility threshold*, $\epsilon$. This threshold is used to filter out correspondences with a low weight (*i.e.*, when $\kappa_c < \epsilon$), whilst minimising the number of beliefs disclosed. The function $\mathsf{joint} : \mathcal{C} \mapsto [0, 1]$ returns the *joint weight* for some correspondence $c \in \mathcal{C}$. This results in either: 1) $\kappa_c^{joint}$ calculated based on the weights for both agents (if both weights have been disclosed); or 2) $\kappa_c^{est}$ for a conservative upper estimate, if only one of the two weights is known.

When deciding which correspondences are viable for inclusion in the final alignment $\mathcal{A}$ the agents chose between two opposite mentalistic attitudes for the admission of disclosed correspondences: *sceptical* and *credulous* acceptability [11]. These attitudes reflect whether or not an agent is prepared to accept the viability of new, hitherto unknown correspondences from its peers. A *sceptical* attitude is one where the agent

**Table 1.** The set $\mathcal{M}$ of legal moves permitted by the dialogue.

| Syntax | Description |
|---|---|
| $\langle x, \textsf{\textit{join}}, \text{nil}, \text{nil}, \text{nil}\rangle$ | Agents utter the *join* move to participate within a dialogue. |
| $\langle x, \textsf{\textit{assert}}, c, \kappa_c^x, \mathcal{R}\rangle$ | The agent $x$ will *assert* the correspondence $c$ that is believed to be viable for inclusion into the final alignment $\mathcal{A}$, and is the undisclosed correspondence with highest personal weight $\kappa_c^x$. If $c$ violates conservativity given $\mathcal{A} \cup \mathcal{O}^x$, then $\mathcal{R}$ will contain a repair plan to resolve this violation to be applied either to the correspondences already in $\mathcal{A}$ or to the newly asserted correspondence $c$. |
| $\langle x, \textsf{\textit{rejectC}}, c, \text{nil}, \text{nil}\rangle$ | If the $c$ asserted in the previous move was not viable (*i.e.*, $\kappa_c^{joint} < \epsilon$, or a violation was subsequently found for $x$ where the repair involved removing $c$ from $\mathcal{A}$), then it is rejected. |
| $\langle x, \textsf{\textit{acceptC}}, c, \kappa_c^{joint}, \text{nil}\rangle$ | Given $c$ received in the previous *assert* move and the associated repair $\mathcal{R}$, if $\kappa_c^{joint} \geq \epsilon$ (*i.e.*, the joint weight is above threshold), and no conservativity violation is generated for $\mathcal{A}' \cup \mathcal{O}^x$ (where $\mathcal{A}'$ is the result of applying the repairs in $\mathcal{R}$ to $\mathcal{A}$), then $c$ is accepted and this joint weight is shared.<br>If a further violation occurs due to $\mathcal{A}' \cup \mathcal{O}^x$, an additional repair should be generated and shared using the *repair* move. |
| $\langle x, \textsf{\textit{repair}}, c, \kappa_c^{joint}, \mathcal{R}\rangle$ | If the agent detected a conservativity violation in $\mathcal{A}' \cup \mathcal{O}^x$ (see *acceptC*), it utters a *repair* move to: 1) indicate that $c$ is acceptable if $\hat{x}$ accepts the repair $\mathcal{R}$; and 2) inform $\hat{x}$ of the resulting joint weight $\kappa_c^{joint}$. Note that the previous repair will not be applied to $\mathcal{A}$ at this stage, as it is predicated on $\hat{x}$ accepting the repair. |
| $\langle x, \textsf{\textit{rejectR}}, c, \text{nil}, \text{nil}\rangle$ | If the agent rejects the proposed repair (*e.g.*, it weakens or removes a mapping deemed necessary for its later transaction), then it can reject the repair, which will also result in $c$ being implicitly rejected. |
| $\langle x, \textsf{\textit{acceptR}}, c, \text{nil}, \text{nil}\rangle$ | The agent $x$ accepts the repair $\mathcal{R}$ for the correspondence $c$ and updates $\mathcal{A}$. On receipt of this move, agent $\hat{x}$ also updates its version of $\mathcal{A}$. |
| $\langle x, \textsf{\textit{close}}, \text{nil}, \text{nil}, \text{nil}\rangle$ | An agent utters a *close* move when it has no more undisclosed viable candidate correspondences. The dialogue terminates when both agents utter a sequence of *close* moves (*i.e.*, a *matched-close*). |

may have reservations or doubts about the acceptability of new correspondences; thus a correspondence is only considered within the dialogue if it is known by all agents involved. In contrast, a *credulous* attitude results in an agent accepting the potential viability of any correspondence that is suggested by its peers. These two attitudes affect the way in which the joint weight is calculated (Definition 2, below), and whether or not a correspondence is considered in the dialogue. When the sender $x$ receives an *assert* move from $\hat{x}$ for a correspondence it knowns (*i.e.*, where $c \in \Delta^x$), it can assess the joint weight for $c$ as the average between its own weight and that shared by $\hat{x}$ (**Case 1**). If, however, $x$ has no prior knowledge of $c$ (*i.e.*, $c \notin \Delta^x$), then the acceptability of the correspondence, and its joint weight will depend on the mentalistic attitude of the agent. If it adopts a *sceptical* attitude, then the correspondence will be rejected, and the joint weight for $c$ will be zero (**Case 2a**). However, if a *credulous* attitude is adopted, the joint weight will depend only on $\kappa_c^{\hat{x}}$ (**Case 2b**). Finally, if $x$ holds a belief on $c$ that has not yet been disclosed to $\hat{x}$ ($c \in \Delta^x; c \notin CS$) and if $\kappa_c^{\hat{x}}$ has not been disclosed by $\hat{x}$, then an *upper bound* $\kappa_u^x$ estimate is assumed (**Case 3**). The *upper bound*, $\kappa_u^x$ is explained below.

***Definition 2:*** *The function* joint : $\mathcal{C} \mapsto [0, 1]$ *returns the* joint weight *for* $c \in \mathcal{C}$*:*

$$\mathsf{joint}(c) = \begin{cases} avg(\kappa_c^x, \kappa_c^{\hat{x}}) & \textbf{Case 1: } c \in \Delta^x \cap \Delta^{\hat{x}}, c \in CS \\ 0 & \textbf{Case 2a: (sceptical) } c \notin \Delta^x, c \in CS \\ \kappa_c^{\hat{x}} & \textbf{Case 2b: (credulous) } c \notin \Delta^x, c \in CS \\ avg(\kappa_c^x, \kappa_u^x) & \textbf{Case 3: } c \in \Delta^x, c \notin CS \end{cases}$$

Each agent takes turns to propose a correspondence, and the other participant confirms if the joint weight $\kappa_c^{joint} \geq \epsilon$. Proposals are made by identifying an undisclosed correspondence with the highest weight $\kappa_c^x$. As the dialogue proceeds, each subsequent correspondence asserted will have an equivalent or lower weight than that previously asserted by the same agent.

Whenever a correspondence is asserted, the agent should check that its estimated joint weight $\kappa_c^{est}$ is not less than the *admissibility threshold*, $\epsilon$. Because the estimate is an upper estimate, the final joint weight $\kappa_c^{joint}$ could subsequently be lower, and the correspondence still rejected. Agents determine this upper estimate by exploiting the fact that assertions are always made on the undisclosed correspondence with the highest weight. Thus, if one agent asserts some correspondence, the other agent's weight for that asserted correspondence will never be greater than their own previous assertion. Therefore, each agent maintains an *upper bound*, $\kappa_u^x$, corresponding to the other agents assertions (prior to the dialogue, $\kappa_u^x = 1.0$).

## 3 Repairing the Alignment

The goal of each agent is to extend each ontology so that there exists a set of entities that are common to both ontologies (*i.e.*, those in $\mathcal{A}$). This should subsequently facilitate the meaningful exchange of knowledge between the two agents, provided that it is expressed using entities within $\mathcal{A}$.

The conservativity principle (general notion) states that $\mathcal{O}^x \cup \mathcal{A} \cup \mathcal{O}^{\hat{x}}$ should not induce any change in the concept hierarchies of the input ontologies [7,15]. In [7], a light variant of the conservativity principle was proposed, which required that an alignment $\mathcal{A}$, together with one of the input ontologies $\mathcal{O}$, should not induce new subsumption relationships between concepts from $\mathcal{O}$. In this paper, we reuse this light variant of the conservativity principle since each agent has only full access to its own ontology (*i.e.* $\mathcal{O}^x$) and the ontology of the other agent is seen as private (*i.e.* $\mathcal{O}^{\hat{x}}$).

The violation detection and repair mechanism defined in [14,15] has been adapted to *incrementally* check for violations as new correspondences are proposed for inclusion within $\mathcal{A}$. As the ontologies themselves are considered as immutable, repairs can only occur over the existing set of correspondences $C \subseteq \mathcal{A}$, or with the candidate correspondence $c$. A repair, given $\mathcal{A}$ and a candidate correspondence $c$, is a set of correspondences whose removal from the alignment would eliminate a conservativity violation. We define a *Alignment Repair* as follows:

**Table 2.** The private and joint *weights* for the correspondences in the worked example, and the two ontology fragments assumed by the two agents *Alice* and *Bob* .

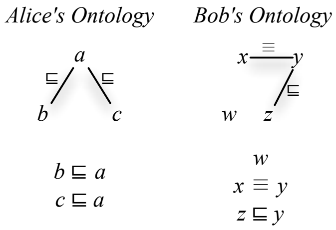


| **c** | $\kappa_c^{Alice}$ | $\kappa_c^{Bob}$ | $\kappa_c^{joint}$ |
|---|---|---|---|
| $\langle a, w, \equiv\rangle$ | 0.25 | 0.35 | **0.3** |
| $\langle a, x, \equiv\rangle$ | 0.9 | 0.8 | **0.85** |
| $\langle b, x, \equiv\rangle$ | 0.55 | 0.45 | **0.5** |
| $\langle b, y, \equiv\rangle$ | 0.4 | 0.7 | **0.55** |
| $\langle b, z, \equiv\rangle$ | 0.6 | 0.55 | **0.575** |

***Definition 3:*** *Let $\mathcal{A}'$ be the new set of correspondences $\mathcal{A} \cup \{c\}$ where $c$ is a candidate correspondence and $\mathcal{A}$ is the current alignment w.r.t. $\mathcal{O}^x$ for which there is a conservativity violation. An alignment $\mathcal{R} \subseteq \mathcal{A}'$ is a repair for $\mathcal{A}'$ w.r.t. $\mathcal{O}^x$ iff there are no such violations in $\mathcal{O}^x \cup \mathcal{A}' \setminus \mathcal{R}$.*

A trivial repair is $\mathcal{R} = \{c\}$, as the removal of the candidate correspondence $c$ that introduces a violation would obviously eliminate that violation. However, the objective is to remove as few (useful) information as possible (*i.e.* a repair may include the weakening of existing equivalence correspondences[5]). Furthermore, in case of multiple options, the correspondence weight will be used as a differentiating factor (*i.e.* a correspondence with a lower weight will be weakened over a correspondence with higher weight). When a correspondence is weakened, it "inherits" the weight of the original correspondence; so that it can be considered for future repairs.

## 4 Inquiry Dialogue Example

We illustrate the dialogue by means of an example. Two agents, $Alice$ and $Bob$, each possess a private ontological fragment, that provides the conceptualisation for the entities that they use to communicate (Table 2). Each agent has acquired a subset of correspondences, with an associated *weight* $\kappa_c$ in the range $[0, 1]$, which is initially private to each agent. These are summarised (with the resulting $\mathsf{joint}(c)$ for each $c$) in Table 2. Note that all of the correspondences are known to both agents, and thus the choice of mentalistic attitude (*i.e.*, *credulous* or *sceptical*) is not relevant in this example. Finally, both agents assume that the *evidence threshold* $\epsilon = 0.45$ to filter out correspondences with a low $\mathsf{joint}(c)$.

The example dialogue between *Alice* and *Bob* is presented in Table 3[6]. The two agents initiate the dialogue by both uttering the *join* move (Moves 1-2 in Table 3), and the turn order is non-deterministic; in this example, *Alice* makes the first *join* move. Each exchange is shown with its move identifier, and the state (taken from Figure 1) from which the move is taken.

[5] As an equivalence correspondence $(a \equiv b) \models (a \sqsubseteq b) \sqcap (b \sqsubseteq a)$, it can be weakened by eliminating one of the two subsumption relations.

[6] The complete trace, obtained through a prototypical implementation, is available at `http://ftp.disi.unige.it/person/SolimandoA/traceOWLED15.txt`.

**Table 3.** The messages exchanged between *Alice* and *Bob* in the example dialogue.

| Move | State | Locution |
|---|---|---|
| 1 | 1A | $\langle$*Alice*, *join*, nil, nil, nil$\rangle$ |
| 2 | 2B | $\langle$*Bob*, *join*, nil, nil, nil$\rangle$ |
| 3 | 3A | $\langle$*Alice*, *assert*, $\langle a, x, \equiv\rangle, 0.9, \varnothing\rangle$ |
| 4 | 4B | $\langle$*Bob*, *acceptC*, $\langle a, x, \equiv\rangle, 0.85$, nil$\rangle$ |
| 5 | 6B | $\langle$*Bob*, *assert*, $\langle b, y, \equiv\rangle, 0.7, \varnothing\rangle$ |
| 6 | 7A | $\langle$*Alice*, *acceptC*, $\langle b, y, \equiv\rangle, 0.55$, nil$\rangle$ |
| 7 | 3A | $\langle$*Alice*, *assert*, $\langle b, z, \equiv\rangle, 0.6, \varnothing\rangle$ |
| 8 | 4B | $\langle$*Bob*, *repair*, $\langle b, z, \equiv\rangle, 0.575, \{\langle b, y, \sqsupseteq\rangle\}\rangle$ |
| 9 | 5A | $\langle$*Alice*, *acceptR*, $\langle b, z, \equiv\rangle$, nil, nil$\rangle$ |
| 10 | 6B | $\langle$*Bob*, *assert*, $\langle b, x, \equiv\rangle, 0.55, \{\langle b, x, \sqsupseteq\rangle\}\rangle$ |
| 11 | 7A | $\langle$*Alice*, *acceptC*, $\langle b, x, \equiv\rangle, 0.5$, nil$\rangle$ |
| 12 | 3A | $\langle$*Alice*, *close*, nil, nil, nil$\rangle$ |
| 13 | 6B | $\langle$*Bob*, *assert*, $\langle a, w, \equiv\rangle, 0.35, \varnothing\rangle$ |
| 14 | 7A | $\langle$*Alice*, *rejectC*, $\langle a, w, \equiv\rangle$, nil, nil$\rangle$ |
| 15 | 3A | $\langle$*Alice*, *close*, nil, nil, nil$\rangle$ |
| 16 | 6B | $\langle$*Bob*, *close*, nil, nil, nil$\rangle$ |

**Move 3:** *Alice* selects one of her undisclosed correspondences with the highest $\kappa_c$; in this case, $\langle a, x, \equiv\rangle$. Initially, *Alice* assumes *Bob's* upper bound $\kappa_u^{Bob} = 1$, and estimates the joint weight for $c$, $\kappa^{est}_{\langle a,x,\equiv\rangle} = \frac{1}{2}(0.9 + 1) = 0.95$. As this is equal to or above threshold ($\epsilon = 0.45$), she asserts $c$. Given that $\mathcal{A} = \varnothing$, the inclusion of $c$ introduces no conservativity violation and thus, no repair is necessary.

**Move 4:** *Bob* confirms that the joint weight $\kappa^{joint}_{\langle a,x,\equiv\rangle} = \frac{1}{2}(0.9 + 0.8) = 0.85$ is above threshold, and checks to see if any repair is needed before accepting the correspondence. As none is needed, he simply accepts $c$, and notifies *Alice* of the joint weight $\kappa^{joint}_{\langle a,x,\equiv\rangle}$. He adds $c$ to $\mathcal{A}$ and updates *Alice's* upper bound $\kappa_u^{Alice} = 0.9$; as if this was *Alice's* highest weighted correspondence, she will have no other undisclosed $c$ where $\kappa_c^{Alice} > \kappa_u^{Alice}$. On receipt of this move, *Alice* adds $c$ to $\mathcal{A}$.

**Moves 5-6:** *Bob* selects his highest privately-weighted undisclosed correspondence $c = \langle b, y, \equiv\rangle$. He estimates the joint weight $\kappa^{est}_{\langle b,y,\equiv\rangle} = \frac{1}{2}(0.7 + \kappa_u^{Alice}) = 0.8$, which is above threshold, and finds that no repairs are necessary if $c$ is added to $\mathcal{A}$. He therefore asserts $c$. *Alice* confirms that $\kappa^{joint}_{\langle b,y,\equiv\rangle} = \frac{1}{2}(0.7 + 0.4) = 0.55$ is above threshold, and from her perspective, no repairs are necessary. She accepts the correspondence, adds $c$ to $\mathcal{A}$ and updates her upper bound $\kappa_u^{Bob} = 0.7$.

At this point, both agents have the alignment $\mathcal{A} = \{\langle a, x, \equiv\rangle, \langle b, y, \equiv\rangle\}$. If we consider $\mathcal{O}^{Alice} \cup \mathcal{A} \cup \mathcal{O}^{Bob}$ then we have a potential conservativity violation, as if we consider *Bob's* axiom $x \equiv y$, then $\mathcal{A} \cup \{x \equiv y\} \models (b \sqsubseteq a) \wedge (a \sqsubseteq b)$. However, as *Bob* only has the axiom $b \sqsubseteq a$, this introduces a new axiom that violates conservativity. It is important to note that this only occurs if both ontologies are known, which is not the case in the current example. If we consider each ontology individually with $\mathcal{A}$, then no violations occur. In the next move, we consider the case where a violation is introduced given a single ontology.

**Table 4.** The status of the dialogue after Move 9.

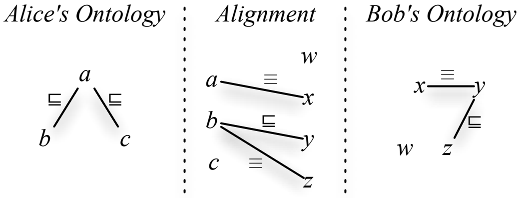


$$\begin{aligned} \mathcal{A} &= \{\langle a,x,\equiv\rangle, \langle b,y,\sqsubseteq\rangle, \\ &\quad \langle b,z,\equiv\rangle\} \\ \kappa_u^{Alice} &= 0.6 \\ \kappa_u^{Bob} &= 0.7 \\ c \in CS &= \{\langle a,x,\equiv\rangle, \langle b,y,\equiv\rangle, \\ &\quad \langle b,z,\equiv\rangle\} \\ c \notin CS &= \{\langle b,x,\equiv\rangle, \langle a,w,\equiv\rangle\} \end{aligned}$$

**Moves 7-9:** *Alice* identifies the next viable correspondence: $\langle b,z,\equiv\rangle$, which is above threshold; *i.e.*, $\kappa^{est}_{\langle b,z,\equiv\rangle} = \frac{1}{2}(0.6 + \kappa_u^{Bob}) = 0.65$. *Bob* calculates $\kappa^{joint}_{\langle b,z,\equiv\rangle} = 0.575$, but also discovers a conservativity violation with this correspondence, as within his ontology, he has the axiom $z \sqsubseteq y$, and the inclusion of both $b \equiv y$ and $b \equiv z$ would also infer the axiom $y \sqsubseteq z$ (similarly for $x \sqsubseteq z$). Given that only correspondences are repaired (rather than ontological axioms), either $b \equiv y$ or $b \equiv z$ should be weakened, based on the joint or estimated weight. As $\kappa^{joint}_{\langle b,y,\equiv\rangle} < \kappa^{joint}_{\langle b,z,\equiv\rangle}$, *Bob* suggests a repair that weakens $b \equiv y$ by removing $b \sqsupseteq y$ and thus leaving the correspondence $b \sqsubseteq y$. *Bob* then updates the upper bound $\kappa_u^{Alice} = 0.6$. As the addition of the "repaired" correspondence doesn't introduce any further violations for *Alice*, she accepts the repair.

The status of the dialogue at this point is illustrated in Table 4.

**Moves 10-11:** *Bob* identifies the next viable correspondence: $\langle b,x,\equiv\rangle$, which is above threshold; *i.e.*, $\kappa^{est}_{\langle b,x,\equiv\rangle} = \frac{1}{2}(0.45 + \kappa_u^{Alice}) = 0.525$. However, the inclusion of this correspondence would introduce a conservativity violation for him, as $(b \equiv x) \wedge (b \equiv z) \models (x \equiv z)$; and yet $\mathcal{O}^{Bob} \models (x \equiv y)$, which also leads to the violation $(z \equiv y)$. The possible repairs are to weaken either $\langle b,x,\equiv\rangle$ or $\langle b,z,\equiv\rangle$. As $\kappa^{est}_{\langle b,x,\equiv\rangle} = 0.525 < \kappa^{joint}_{\langle b,z,\equiv\rangle} = 0.575$, the new correspondence $\langle b,x,\equiv\rangle$ is weakened, resulting in the repair $\mathcal{R} = \{\langle b,x,\sqsupseteq\rangle\}$. *Alice* confirms that the non-weakened version of the asserted correspondence $\kappa^{joint}_{\langle b,x,\equiv\rangle} = 0.5 > \epsilon$, and that no further violations are incurred as a consequence of the repair. As such, she accepts the assertion, and updates the upper bound $\kappa_u^{Bob} = 0.55$

**Moves 12-16:** *Alice* estimates the joint weight of her remaining correspondence, but as $\kappa^{est}_{\langle a,w,\equiv\rangle} = \frac{1}{2}(0.25 + \kappa_u^{Bob}) = 0.4 < \epsilon$, she realises that she has no further viable correspondences to assert, and thus makes a *close* move. *Bob*, unaware that *Alice* is also aware of this correspondence, generates a higher estimate ($\kappa^{est}_{\langle a,w,\equiv\rangle} = 0.5 > \epsilon$), which appears viable. Thus, in move 13, he asserts the correspondence. *Alice* unsurprisingly calculates the actual joint weight, which (as she had originally estimated) was below threshold, and then rejects $c$. Again, as she has no further correspondences, she makes a *close* move. *Bob* has no further correspondences that are viable, so also utters a *close* move, and the dialogue terminates.

The resulting alignment $\mathcal{A} = \{\langle a,x,\equiv\rangle, \langle b,x,\sqsubseteq\rangle, \langle b,y,\sqsubseteq\rangle, \langle b,z,\equiv\rangle\}$ is illustrated in Figure 2, together with the resulting alignment that would have been generated if the original *Correspondence Inclusion Dialogue* had been used on this example [12]. The

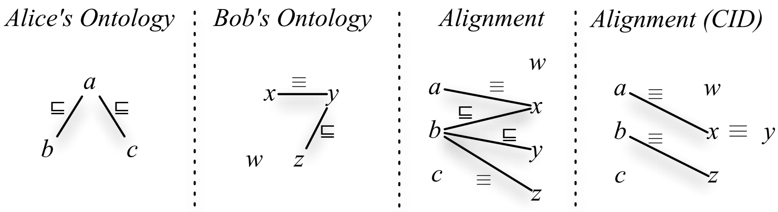


**Fig. 2.** The final alignment $\mathcal{A}$, with the two ontologies. The second alignment (Alignment CID) represents the result if the example had been evaluated using the original *Correspondence Inclusion Dialogue* [12].

dialogue presented here results in two further correspondences, $\langle b, x, \sqsubseteq \rangle$ and $\langle b, y, \sqsubseteq \rangle$, both of which are consistent with the fact that: 1) from *Alice's* perspective she has the correspondence $\langle a, x, \equiv \rangle$ and that $b \sqsubseteq a$; or 2) from *Bob's* perspective he has the correspondence $\langle b, z, \equiv \rangle$ and that $z \sqsubseteq y$.

## 5 Related Work

A number of different approaches have addressed the reconciliation of heterogeneous ontologies by using some form of rational reasoning. Argumentation has been used as a rational means for agents to select ontology correspondences based on the notion of partial-order preferences over their different properties (*e.g.*, structural vs terminological) [9]. A variant was also proposed [17] which represented ontology mappings as disjunctive queries in Description Logics. Typically, these approaches have used a course-grained decision metric based on the *type* of correspondence, rather than whether or not each correspondence was *acceptable* to each agent (given other mutually accepted correspondences), and do not consider the notion of private, or asymmetric knowledge (the correspondences are assumed to be publicly accessible). [16] used a Max-Sum algorithm for synthesising ontology alignment methods whilst maximising social welfare in a group of interacting agents. Although similar to the aims of our study, [16] assumes that all agents have knowledge of the ontologies to align, and each agent is associated with an alignment method with its own preferences on the assessed relation, and quantified by a degree of confidence.

## 6 Conclusions

This paper presents a novel inquiry dialogue that significantly extends the *Correspondence Inclusion Dialogue* described in [12,11]. The dialogue facilitates negotiation over asymmetric and incomplete knowledge of ontological correspondences. Our dialogue enables two agents to selectively disclose private correspondences given their perceived utility. Ambiguous correspondences are only permitted when they do not introduce conservativity violations for each agent's ontology in isolation. A walkthrough example was also presented, that illustrates how the conservativity violation repairs are shared and applied.

## Acknowledgements

This work was partially funded by the EU project Optique (FP7-ICT-318338), and the EPSRC projects $MaSI^3$, Score!, and DBOnto.

## References

1. Doran, P., Tamma, V., Payne, T.R., Palmisano, I.: Dynamic selection of ontological alignments: a space reduction mechanism. In: Proc. IJCAI. pp. 2028–2033 (2009)
2. Dragisic, Z., Eckert, K., Euzenat, J., Faria, D., Ferrara, A., Granada, R., Ivanova, V., Jiménez-Ruiz, E., Kempf, A.O., Lambrix, P., Montanelli, S., Paulheim, H., Ritze, D., Shvaiko, P., Solimando, A., dos Santos, C.T., Zamazal, O., Cuenca Grau, B.: Results of the Ontology Alignment Evaluation Initiative 2014. In: Proceedings of the 9th International Workshop on Ontology Matching, Riva del Garda, Trentino, Italy, October 20, 2014. pp. 61–104 (2014)
3. Euzenat, J., Shvaiko, P.: Ontology Matching. Springer-Verlag (2007)
4. Gale, D., Shapley, L.S.: College admissions and the stability of marriage. The American Mathematical Monthly 69(1), 9–15 (1962)
5. Grice, H.P.: Logic and conversation. In: Cole, P., Morgan, J.L. (eds.) Syntax and Semantics: Vol. 3: Speech Acts, pp. 41–58. Academic Press, San Diego, CA (1975)
6. Jiménez-Ruiz, E., Cuenca Grau, B.: LogMap: Logic-based and scalable ontology matching. In: Proc. of the International Semantic Web Conference (ISWC), Lecture Notes in Computer Science, vol. 7031, pp. 273–288. Springer Berlin Heidelberg (2011)
7. Jiménez-Ruiz, E., Cuenca Grau, B., Horrocks, I., Berlanga, R.: Logic-based assessment of the compatibility of UMLS ontology sources. J. Biomedical Semantics 2(1) (2011)
8. Kuhn, H.W.: The hungarian method for the assignment problem. Naval research logistics quarterly 2(1-2), 83–97 (1955)
9. Laera, L., Blacoe, I., Tamma, V., Payne, T., Euzenat, J., Bench-Capon, T.: Argumentation over ontology correspondences in MAS. In: Proc. AAMAS 2007. pp. 1285–1292 (2007)
10. Meilicke, C.: Alignment Incoherence in Ontology Matching. Ph.D. thesis, Dissertation, Universität Mannheim, Mannheim (2011)
11. Payne, T.R., Tamma, V.: A dialectical approach to selectively reusing ontological correspondences. In: Knowledge Engineering and Knowledge Management - 19th International Conference, EKAW 2014, Linköping, Sweden, November 24-28, 2014. Proceedings. pp. 397–412 (2014)
12. Payne, T.R., Tamma, V.: Negotiating over ontological correspondences with asymmetric and incomplete knowledge. In: Proc. AAMAS 2014. pp. 517–524 (2014)
13. Shvaiko, P., Euzenat, J.: Ontology matching: State of the art and future challenges. IEEE Trans. Knowl. Data Eng. 25(1), 158–176 (2013)
14. Solimando, A., Jiménez-Ruiz, E., Guerrini, G.: A Multi-strategy Approach for Detecting and Correcting Conservativity Principle Violations in Ontology Alignments. In: OWLED Workshop. vol. CEUR-WS 1265, pp. 13–24 (2014)
15. Solimando, A., Jiménez-Ruiz, E., Guerrini, G.: Detecting and correcting conservativity principle violations in ontology-to-ontology mappings. In: Proc. ISWC 2014. vol. LNCS 8797, pp. 1–16 (2014)
16. Spiliopoulos, V., Vouros, G.A.: Synthesizing ontology alignment methods using the max-sum algorithm. IEEE TKDE 24(5), 940–951 (2012)
17. Trojahn dos Santos, C., Quaresma, P., Vieira, R.: Conjunctive queries for ontology based agent communication in MAS. In: Proc. AAMAS 2008. pp. 829–836 (2008)
18. Walton, D., Krabbe, E.: Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning. SUNY series in Logic and Language, State University of New York Press (1995)