Dialogue Based Meaning Negotiation

Gabrielle Santos, Valentina Tamma, Terry R. Payne, and Floriana Grasso

Department of Computer Science, University of Liverpool,
Liverpool L69 3BX, United Kingdom
{sggsantos,V.Tamma,T.R.Payne,floriana}@liverpool.ac.uk

Abstract. Recent advances in technology have caused a proliferation of data and knowledge sources on a global scale. The ability to access and integrate these knowledge sources is crucial for critical decision making, and to facilitate this, knowledge-based intelligent applications (*agents*) need to resolve the differences between their knowledge models (*ontologies*). Research into ontology alignment has typically relied on the use of an external system such as an oracle that takes as input two ontologies and produces the most plausible alignment between the entities (classes, properties and instances), given some pre-defined similarity models. However, such approaches are used irrespectively of whether the intended models underlying the ontologies overlap, and hence without any indication as to whether an alignment representing this overlap can be meaningfully computed. Furthermore, traditional alignment methods usually require the disclosure of the full ontological model, even in those situations in which only a few concepts or a module is needed. In this paper, we present preliminary work that allows two agents to jointly determine a single correspondence between two concepts in their respective ontologies. The agents engage in a dialogue that permits the participants to exchange information about the concepts to support the assertion or rejection of a correspondence. The agents reason over the plausibility of the correspondence by considering information related to the ontological context of a concept (expressed in terms of properties) and the partial knowledge acquired during the dialogue about the other agent's ontology. Thus, the approach does not require the full disclosure of both ontologies prior to the reconciliation phase, but supports knowledge sharing in an incremental and opportunistic fashion.

1 Introduction

Recent advances in technology have resulted in a proliferation of data and knowledge sources on a global scale, with low-cost sensors and personal devices being the drivers of some of the fastest growing sources and consumers of data. With the huge diversity and volume of data expected to further increase in the near future, collaboration and coordination between different systems is essential. This has led to a pressing demand for more flexible knowledge-based intelligent applications (*agents*) that can exploit machine-readable domain knowledge (*ontologies*), and perform tasks that require the integration of disparate knowledge sources. Ontologies are machine readable specifications of a conceptualisation of some given domain knowledge, through the definition of entities and their relationship with each other. However it is often the case that the agents differ in the vocabulary (and ontologies) they assume, thus compromising seamless *semantic interoperability* between dynamic and evolving systems.

Ontology alignment [6] (the creation of sets of mappings between the corresponding entities within a pair of ontologies) provide a basis for semantic interoperability between the knowledge bases of the respective agents, and thus are an essential component for agent communication [13]. However, even in similar domains, ontologies can be modelled differently using a variety of modelling languages and contrasting assumptions, which can make translating one ontology into another increasingly difficult. In order for two systems to accurately and successfully communicate over their vocabularies, the heterogeneity within ontologies needs to be resolved. The ontology alignment community has proposed many diverse approaches that *align* ontologies in order to find these sets of correspondences; however, most rely on the ontologies to be fully shared [6] and attempt to find correspondences between entities without having any indication that the intended models underlying the ontologies are actually compatible. A recent review of ontology matching research [15] highlighted that whilst considerable progress had been made in recent years, the performance of the matchers across different tasks is still an issue, and can vary greatly. Thus it can be problematic to determine which approach would produce the best alignment between two particular ontologies for a specific task [13]. Furthermore, it may not always be necessary to align all of the ontological entities, particularly in those scenarios where ontologies have grown considerably due to the inclusion of several different ontological models, or where the disclosure of certain modules in the ontologies may be problematic or undesirable, as they contain definitions that are private or commercially sensitive [8].

In this paper we propose a dialogue-based approach that allows two agents, with different ontologies, to dynamically form a consensus over candidate correspondences that align their ontologies. The approach we proposed is represented by an *Inquiry Dialogue* [16] as a *cognitive-inspired* mechanism that allows two agents to iteratively query, propose and verify details regarding the entities in their respective ontologies and to collaboratively explore the viability of potential correspondences by implementing an incremental comparison of the structural semantics of the entities in question. The agents deliberate over the assertion of some correspondence only when they can provide (and are provided with) sufficient evidence of its plausibility. No assumptions are made regarding prior access or disclosure of both ontologies by either agent, and that the agents may chose to keep certain entities private. The agents take a *lazy* approach by only sharing entities on a needs-basis, such that the number of entities shared is minimised. This contrasts with traditional approaches that make the assumption that all ontological knowledge can be disclosed by the agents, and that the disclosure merely facilitates the alignment process; *i.e.* they *greedily* disclosed ontological elements whether or not such a disclosure is relevant or useful. The approach described here presents a *dialogical fragment* for determining a correspondence between concepts in two ontologies, which can be used repeatedly by the agents to determine as many correspondences as are needed, covering the whole ontology or just an ontological fragment (or module).

This paper is organised as follows: Section 2 presents a cognitive justification for the use of dialogue to facilitate ontology alignment, and the formal model itself is presented in Section 3. A walkthrough example illustrates how the model should work, and is discussed in Section 4. A survey of related work appears (Section 5) before the paper concludes in Section 6.

2 Cognitive Phases

Ontology alignment has traditionally been viewed as a centralised process, whereby two ontologies are submitted to a central oracle that identifies mappings between corresponding entities. Such approaches try to maximise the coverage, *i.e.* the number of correspondences created (given some objective function), often disregarding the reason why the alignment was being generated (i.e. to facilitate some task), or other knowledge possessed by the ontology owner (e.g. an agent). The dialogue based alignment mechanism proposed here is based on the notion of *conversations* as social constructs where utterances are exchanged in order to achieve some *joint activity* or task [3]; and on the cognitive mechanisms for communication and coordination of activities [14, 10]. The motivation behind this lies in the observation that often alignments are generated without considering whether the intended model underlying the ontologies are actually compatible. Thus, generic alignment mechanisms do not offer any guarantee that even if an alignment can be found, this will actually support the representation of a joint task.

Furthermore, the increasing interest in sound ontology development has elevated the reuse of (fragments of) well-established ontologies as part of the ontology engineering process. Reuse in OWL ontologies is facilitated by the *import* mechanism: one ontology can import another ontology fragment (a coherent set of axioms) and the resulting ontology is logically equivalent to the union of the two theories. Applications need access to both sets of axioms [8], and thus generic alignment approaches match entities defined in both of these theories as a result, irrespectively of whether they are relevant to the intended joint activity. Likewise, they may require access to axioms that represent proprietary or commercially sensitive knowledge that an agent may otherwise be unwilling to share. Therefore, there is a need for alignment approaches that only share knowledge pertinent to some joint task. By structuring the alignment process as a conversation, agents can themselves determine what axioms they are prepared to share.

For a conversation to be successful, it must result in some contribution; *i.e.* the parties involved cooperate in order to achieve the mutual belief that the *listener* has understood the *speaker* well enough to finalise the conversation, thus *grounding the meaning* within the scope of that specific conversation [3]. The dialogue we propose determines whether there is a *common ground* [4] for establishing the alignment. An underlying assumption is that it satisfies the principle of least collaborative effort, where participants try to minimise the total effort spent on a conversation, as typically the fewer exchanges required to clarify references, the better the common ground. The dialogue also obeys Grice's *cooperative principle* [9] by assuming that: the participating agents are truthful; that they make informative contributions as required; and that they keep their interactions terse and thus do not provide more information than necessary.

From a cognitive speech perspective, a conversation comprises three phases: *Open*, *Information Exchange*, and *Close*[1]. The dialogue fragment illustrated in Section 3 negotiates the existence of a specific correspondence, and forms part of a larger dialogical structure that iterates over all of the ontological entities relevant to a task[2]. This conversational structure has been extended to four phases within our dialogue: *Open*; *Propose* and *Confirm* (corresponding to the *Information Exchange* phase); and *Close*.

[1] This structure appears in many other cognitive-based models of dialogues [3, 10, inter alia].

[2] The description of the conditions supporting these iterations is outside the scope of this paper.

- During the *Open* phase, the initiating agent states the aim of the dialogue fragment by uttering the name of the concept for which a correspondence needs to be found. The interlocutor agent then decides whether or not to collaborate in determining a correspondence for that given concept. The agents then cycle through a sequence of *Propose* and *Confirm* phases to establish the candidacy of the correspondence.
- The *Propose* phase allows one agent to gather *supporting evidence (support)* in favour of a given correspondence by determining the other's *context* of the matched concept and comparing it with that of its own, expressed as the relationships originating from that concept. It can also choose to reject the correspondence, resulting in a return to the *Open* phase whereby the interlocutor can suggest an alternate concept, or allow the dialogue fragment to fail. Once the support for the correspondence is deemed satisfactory, the agent can assert it as a viable candidate.
- The *Confirm* phase allows the other agent to verify that the *support* received satisfactorily matches the internal structure of its own ontology. If so, it can then accept the correspondence. Otherwise, it can either reject the support itself, or try to augment the support with additional evidence. This results in a transition back to the *Propose* phase, only this time, driven by the other agent.
- The *Close* phase terminates the dialogue fragment.

3 Dialogue Module System

The dialogue fragment proposed in this paper allows two agents to exchange details over the existence of ontological entities and their perceived similarity to support a candidate correspondence. We assume 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* a or *recipient* $\hat{a}$) in any single dialogue move. 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 modelled as a set of axioms describing classes and the relations existing between them. Given $\mathsf{N_C}$, a set of concept names, and $\mathsf{N_R}$ a set of role names, then $\Sigma = \mathsf{N_C} \cup \mathsf{N_R}$ 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 $\mathcal{O}^a$, whereas the recipient's ontology is $\mathcal{O}^{\hat{a}}$. For agents to interoperate in an encounter, they need to determine an *alignment* between the two vocabulary fragments Σ^a and $\Sigma^{\hat{a}}$ for that encounter. An alignment [6] consists of a set of *correspondences* that establish the logical relationship between the entities (classes, properties or roles, and instances) belonging to each of the two ontologies, respectively.

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

In this paper we focus our attention on finding concept correspondences, hence we only consider aligning entities in $\mathsf{N_C}^a$ and $\mathsf{N_C}^{\hat{a}}$. An ontology can be represented as a *directed labelled graph* $G = (V, E)$, where V denotes the set of vertices (concepts) and the E denotes the set of directed edges representing the properties (object and datatype) in the ontology. The vertices V are labelled with concept names, $\mathsf{N_C}$ while the directed edges in E are labelled with the roles names, $\mathsf{N_R}$. An edge p denotes a relation from the vertex s to the target vertex o, and are referred to as *subject-predicate-object* triples, where the *subject* s is a member of the domain of p, the *object* o is a member of the range of p, and the *predicate* p is also known as the *property* of s.

Definition 2: *A* **subject-predicate-object** *triple is denoted* $\varpi = \langle s, p, o \rangle$ *such that* $s, p, o \in \Sigma$*, where* $s,\ o \in \mathsf{N_C}$ *are vertices and* $p \in \mathsf{N_R}$ *is an edge that relates* s *to* o*.*

Both agents manage a public knowledge base, or *Commitment Store*, CS, which contains a trace of all of the moves uttered by each agent [16]. Although the agents maintain individual copies of the CS, these will always be identical, and thus we do not distinguish between them. Each agent manages its own private knowledge base, known as the *Gamma Store*[3]($\mathit{\Gamma}$), which stores private knowledge garnered throughout the dialogue regarding the ontological entities acquired from the interlocutor. Each of the gamma stores contains a partially connected graph, that is comprised:

- either by an independent vertex $v_i \in \mathsf{N_C}$ representing a candidate concept from the interlocutor's ontology for inclusion in a correspondence;
- or by the *neighbourhood* of the concept v_i, i.e. the subgraph originating from the vertex v_i constructed through the exchange of ϖ triples to support its candidacy.

Within the dialogue, the agents try to ascertain a similarity between the shared entities to determine whether or not there is sufficient evidence to justify proposing or accepting a candidate correspondence. Many approaches for determining similarity have been proposed, or evaluated in the ontology matching literature [2, 6, 7, 15, inter alia]. The dialogue described here assumes that each agent has the capability of determining lexical (i.e. string-based) similarity matching, which could be based on a plethora of different approaches[4], and is defined formally as:

Definition 3: *The* **lexical similarity metric** *is the function* $\sigma_l : \mathsf{N_C} \times \mathsf{N_C} \mapsto [0, 1]$ *which returns the lexical similarity between the labels of two entities* $e, e' \in \mathsf{N_C}$*, such that* $\sigma_l(e, e') = 1$ *iff lexically* $e \equiv e'$ *and 0 if there is no lexical similarity between the two.*

This function is primarily used in the *Open* phase of the dialogue to discover a seed entity in agent a_2's signature that could lexically match an entity in agent a_1's signature.

An important component of the dialogue is how the agents share structural details about the ontology in the neighbourhood of the entity under consideration. As further information becomes available about the *subject-predicate-object* triples used by one agent, the second agent should try to identify similar localised structures in its own ontology. This may be based purely on the triples themselves, or may also take into account other information that has so far been ascertained or inferred. As with the σ_l function, we make no assumptions about how such a similarity function is defined, but simply that there is some function defined formally as:

Definition 4: *The* **structural similarity metric** *is the function* $\sigma_s : \mathit{\Pi} \times \mathit{\Pi} \mapsto [0, 1]$ *which returns the structural similarity between two triples* $\varpi, \varpi' \in \mathit{\Pi}$*, such that* $\sigma_s(\varpi, \varpi') = 1$ *if the two triples are perceived to be equivalent, and 0 if there is no structural similarity between them.*

This function is used in the *Propose* and *Confirm* phases of the dialogue to determine if the neighbourhood of the concepts in the candidate correspondence are similar.

[3] We distinguish between the sender's Gamma Store, $\mathit{\Gamma}^{a}$, and the recipient's store, $\mathit{\Gamma}^{\hat{a}}$.

[4] See [2] for a good survey and evaluation of different string similarity metrics.

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

Syntax	**Description**
$\langle a, \textit{initiate}, e, \text{nil}, \text{nil}\rangle$	Agent a states the source entity e of the candidate correspondence.
$\langle a, \textit{propose}, e, e', \text{nil}\rangle$	Agent a proposes a candidate target entity e' that e may map to.
$\langle a, \textit{justify}, e, e', \text{nil}\rangle$	Agent a requests a new (i.e not previously disclosed) ϖ triple that could potentially support the hypothesis that e could map to e'.
$\langle a, \textit{testify}, e, e', \varpi\rangle$	Agent a responds to a *justify* move by providing a new ϖ triple that bests supports the hypothesis that e maps to e' (w.r.t. $\mathcal{O}^x$ and the facts stored so far in Γ^a). If no such triple exists, then $\varpi = \text{nil}$.
$\langle a, \textit{assert}, e, e', \alpha\rangle$	Given the ϖ pairs in Γ^x and given $\alpha = (S, c)$, a identifies a subset S of these whose local mean similarity supports the candidacy of c.
$\langle a, \textit{accept}, e, e', \alpha\rangle$	Agent a verifies that it accepts S (*i.e.* the support for α), given its own similarity metrics, and thus accepts the correspondence c.
$\langle a, \textit{rejectS}, e, e', \alpha\rangle$	Agent a rejects the validity of S (*i.e.* the support for α), given its own similarity metrics.
$\langle a, \textit{rejectC}, e, e', \text{nil}\rangle$	Agent a is unable to find any further evidence to support the candidacy of a correspondence c between e and e', and thus rejects c.
$\langle a, \textit{fail}, e, \text{nil}, \text{nil}\rangle$	Agent a fails to find a potential entity that could map to e.

As the candidacy of a correspondence c is established, an agent can *assert* an *argument* in favour of c, by proposing evidence or *support* based on the exchanged triples garnered in its *Gamma* store, resulting in matched ϖ pairs $s = (\varpi, \varpi')$, where $\varpi \in \Sigma^a \cap \Gamma^a$ and $\varpi' \in \Sigma^{\hat{a}} \cap \Gamma^{\hat{a}}$. Thus, the *support* for the argument S is a set of ϖ pairs that form a pair of subgraphs of the exchanged ontological fragments stored in the *Gamma* stores. We denote an argument for a correspondence c as a pair $\alpha = (S, c)$. The mean similarity $\bar{\sigma}_s(S)$ for the support S is simply calculated as the arithmetic mean of all of the ϖ pairs in the support. This is used by the agents to determine whether or not to accept the argument for a correspondence c.

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 $m_s = \langle a, \tau, e, e', \varphi\rangle$, where x represents the identity of the agent making the move; $\tau \in \mathcal{M}$ represents the move type; e represents the source entity being discussed (which is defined within the *initiate* move); e' is the candidate target entity (*i.e.* the entity that could be mapped to from e); and $\varphi \in \{\varpi, \alpha\}$. For some moves, it may not be necessary to specify an entity or φ; in which case they will be empty or unspecified (represented as *nil*).

4 Walkthrough Example

We illustrate how two agents utilise the dialogue fragment to find a correspondence c between two concepts 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. The ontologies are represented by their signature $\Sigma = \mathsf{N_C} \cup \mathsf{N_R}$, where $\mathsf{N_C}$ is the set of *concept* names and $\mathsf{N_R}$ is the set of *roles* or properties. In this example, *Alice* possesses an ontology whose signature $\Sigma^{\mathit{Alice}} = \{d, e, f, g, h\}_{\mathsf{C}} \cup \{k, l, m, n\}_{\mathsf{R}}$, whereas *Bob*'s ontology has the signature $\Sigma^{\mathit{Bob}} = \{w, x, y, z\}_{\mathsf{C}} \cup \{r, s, t\}_{\mathsf{R}}$. The concept-property relationships of the two ontologies appear in Table 2, which also lists the most similar *subject-predicate-object* (ϖ) triple pairs. For example, the structural similarity σ_s between the triple $\langle d, k, e\rangle$ and $\langle w, r, z\rangle$ for *Alice*, $\sigma_s^{\mathit{Alice}} = 0.70$. As

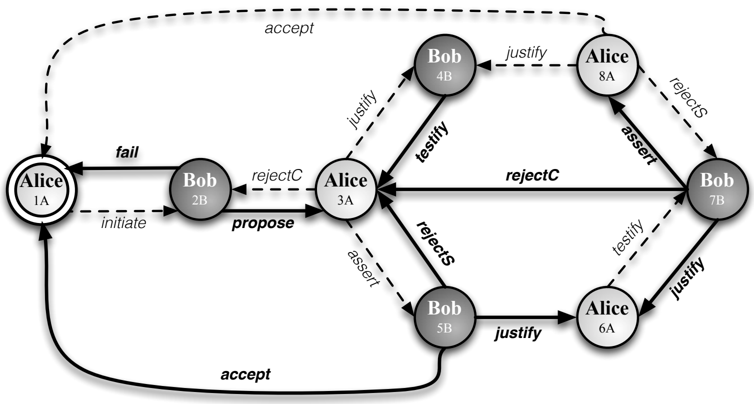

Fig. 1. The dialogue as a state diagram. Nodes indicate the agent whose turn it to utter a move. Moves uttered by *Alice* are labelled with a light font / dashed edge, whereas those uttered by *Bob* are labelled with a heavy font / solid edge. It assumes *Alice* will always make the first move.

we assume that the similarity metrics[5] used by the two agents may differ, the similarity for *Bob* is $\sigma_s^{Bob} = 0.68$. Although other similarity pairs have be calculated, these do not appear in the dialogue example below (for example, because the distance is lower than those explicitly stated), and thus have not been given for the sake of brevity. Note that these similarity pairs are not generated a priori, but are calculated during the dialogue.

Both agents generate a strict pre-ordering[6] of the properties for each concept e, denoted $rank_e$. In this case, the orderings for each concept owned by each agent are: $\mathsf{rank}_d^{Alice} = \{k, m, n, l\}$, and $\mathsf{rank}_w^{Bob} = \{t, r, s\}$. A threshold function $\varepsilon(x)$ is also used to determine if the similarity between the neighbourhoods is acceptable, based on some neighbourhood size x. For the purposes of this example, a function was chosen that reduces the threshold asymptotically as the size of x (i.e. number of ϖ pairs within the support S of an argument) increases; *i.e.* $\varepsilon(x) = \frac{1}{2(x+1)} + 0.5$.

The example dialogue is presented in Table 3. The dialogue fragment starts with move 1 where *Alice* uttering a *initiate* move (state 1A in Figure 1), by stating that she wants to discuss a possible match for concept d (corresponding to the *Open* phase).

Move 2: *Bob*, who is interested in finding correspondences between entities in his and *Alice's* ontologies, identifies w as lexically the most similar concept in his ontology to the symbol d, with a lexical similarity σ_l of 0.82. As this is above threshold[7], he responds with the move $\langle Bob, \mathsf{propose}, d, w, \text{nil}\rangle$.

Move 3: *Alice* now knows that $\langle d, w, \equiv\rangle$ is a potential correspondence c (based on *Bob's* lexical similarity claim). She verifies that her lexical similarity for the concept pair is above threshold (in this case $\sigma_l^{Alice}(d, w) = 0.79$), and then initiates the *Propose* phase (moves 3-7) at state 3A, where she asks *Bob* to provide some evidence to justify the candidacy of c. At this point, neither agents has support for c; i.e. $S = \varnothing$.

[5] Some metrics may exploit the lexical similarity of the labels for the concepts and roles. However, as we are not prescriptive w.r.t. the similarity metric used within this paper, the concept and role names have been replaced with single character identifiers for brevity.

[6] The way in which this ranking is determined is out of scope of this paper.

[7] For simplicity, we use the threshold function $\varepsilon(1)$; *i.e.* $\sigma_l^{Bob}(d, w) \geq \varepsilon(1) = 0.82 \geq 0.75$.

Table 2. The structural similarities of possible corresponding triples between *Alice* & *Bob*'s ontologies. Though not exhaustive, it shows the most similar triples between the two ontologies.

Alice's ϖ	*Bob*'s ϖ	σ_s^{Alice}	σ_s^{Bob}
$\langle d, k, e\rangle$	$\langle w, r, z\rangle$	0.70	0.68
$\langle d, m, f\rangle$	$\langle w, s, x\rangle$	0.65	0.61
$\langle d, n, h\rangle$	$\langle w, r, z\rangle$	0.68	0.84
$\langle d, l, g\rangle$	$\langle w, t, y\rangle$	0.66	0.60

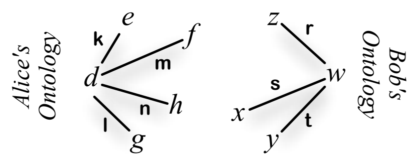

Move 4: *Bob* (state 4B) determines the next property that has w as its *domain* (*i.e.* for which the concept is the subject) that has not yet been disclosed and thus has not yet appeared in the commitment store CS. Given that his ranking for the concept w is $\mathsf{rank}_w^{Bob} = \{t, r, s\}$ and that so far, none of these properties have been disclosed, he shares the fact that the highest ranked property t relates the two concepts w and y.

Move 5: *Alice* tries to determine if there is sufficient support for c. She realises that $\langle d, l, g\rangle$ in her ontology is the most similar triple to the one *Bob* disclosed in move 4, with a similarity $\sigma_s^{Alice} = 0.66$ (Table 2). She adds this to the support, so that $S = \{(\langle d, l, g\rangle, \langle w, t, y\rangle)\}$ with a mean similarity between the triple pairs as $\bar{\sigma}_s^{Alice}(S) = 0.66$. She will only *assert* an argument for c if the mean similarity for S is greater than the threshold given the number of pairs in S. As $\varepsilon(|S|) = 0.75$, the mean similarity $\bar{\sigma}_s^{Alice}(S)$ is below threshold, and thus she requests additional evidence to justify c.

Move 6: *Bob's* next highest ranked property that has not been disclosed (*i.e.* does not appear in CS) whose domain is w is r. Therefore he shares the triple $\langle w, r, z\rangle$.

Move 7: *Alice* verifies if one of her triples is similar to that disclosed by *Bob* in move 6. Although she has two triples that share their highest similarity with *Bob's* disclosed triple, she chooses $\langle d, k, e\rangle$ as the similarity is higher than $\langle d, n, h\rangle$. She adds this to S, and calculates $\bar{\sigma}_s^{Alice}(S) = (0.66 + 0.7)/2 = 0.68$. As the threshold for two pairs of triples is now $\varepsilon(|S|) = \frac{1}{2\times(2+1)} + 0.5 = 0.\dot{6}$, the mean similarity for S (from *Alice's* perspective) is above threshold. Therefore she proposes the argument α for the correspondence $c = \langle d, w, \equiv\rangle$, given that $S = \{(\langle d, l, g\rangle, \langle w, t, y\rangle), (\langle d, k, e\rangle, \langle w, r, z\rangle)\}$.

Move 8: Now that *Bob* (state 5B) has to reflect on the asserted argument α for the correspondence c from the previous move, the dialogue transitions to the *Confirm* phase. He can make one of three possible moves:

- *accept* the argument α if $\bar{\sigma}_s^{Bob}(S)$ is above threshold, and transition to state 1A;
- *rejectS* the support for α if no further support is available (*e.g. Bob* has no undisclosed properties that could support α), and return to state 3A;
- *justify* the candidacy of c by requesting further support (if other undisclosed properties exist), by initiating a new *Propose* phase (state 6A).

In this case, *Bob* calculates that $\bar{\sigma}_s^{Bob}(S) = (0.6 + 0.68)/2 = 0.64$ which is below threshold. However, *Bob* is aware of other properties of the concept w that don't appear in S, and thus asks *Alice* if she could provide some further evidence to justify c.

Move 9: *Alice* shares the triple $\langle d, m, f\rangle$ as m is her highest ranked, non-disclosed property for the domain concept d (property k was ranked higher but was disclosed).

Move 10: *Bob* recalculates the mean similarity for the new support (inclusive of the triple shared by *Alice* in Move 9): $\bar{\sigma}_s^{Bob}(S) = (0.60 + 0.68 + 0.61)/3 = 0.63$. Given that there are three triple pairs in S, the threshold is now $\varepsilon(3) = \frac{1}{2\times(3+1)} + 0.5 = 0.625$, and thus the mean similarity is above threshold. If, however, *Alice* had not provided any

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

Move	Locution
1	$\langle Alice, \textsf{initiate}, d, \text{nil}, \text{nil}\rangle$
2	$\langle Bob, \textsf{propose}, d, w, \text{nil}\rangle$
3	$\langle Alice, \textsf{justify}, d, w, \text{nil}\rangle$
4	$\langle Bob, \textsf{testify}, d, w, \langle w, t, y\rangle\rangle$
5	$\langle Alice, \textsf{justify}, d, w, \text{nil}\rangle$
6	$\langle Bob, \textsf{testify}, d, w, \langle w, r, z\rangle\rangle$
7	$\langle Alice, \textsf{assert}, d, w, (\{(\langle d, l, g\rangle, \langle w, t, y\rangle), (\langle d, k, e\rangle, \langle w, r, z\rangle)\}, \langle d, w, \equiv\rangle)\rangle$
8	$\langle Bob, \textsf{justify}, d, w, \text{nil}\rangle$
9	$\langle Alice, \textsf{testify}, d, w, \langle d, m, f\rangle\rangle$
10	$\langle Bob, \textsf{assert}, d, w, (\{(\langle w, t, y\rangle, \langle d, l, g\rangle), (\langle w, r, z\rangle, \langle d, k, e\rangle), (\langle w, s, x\rangle, \langle d, m, f\rangle)\},$ $\langle d, w, \equiv\rangle)\rangle$
11	$\langle Alice, \textsf{accept}, d, w, \langle d, w, \equiv\rangle)$

further evidence (*i.e.* if $\varpi = \text{nil}$), he would have rejected the support using the *rejectC* move (returning to state 3A). As it is, *Bob* is now happy to accept the candidacy of c. It is now his turn to *assert* the new argument for c given the new support S; signalling another transition from the *Propose* to the *Confirm* phase.

Move 11: *Alice* confirms that from her perspective, $\bar{\sigma}_s^{Alice}(S) = (0.66 + 0.7 + 0.65)/3 = 0.67 \geq 0.625$, and in accepting the argument, transitions to the *End* phase.

At this point, through co-operation, the agents were able to engage in the joint activity of determining a correspondence between two concepts based on the similarity of the local neighbourhood of the concepts. Although all of *Bob's* ϖ triples were disclosed, *Alice* was able to to reach the consensus without revealing knowledge of one of her triples: $\langle d, n, h\rangle$, even though from *Bob's* perspective, it was actually more similar to *Bob's* triple $\langle w, r, z\rangle$ than $\langle d, k, e\rangle$. If in move 7, *Alice* had found that the triple with the highest similarity to $\langle w, r, z\rangle$ was actually $\langle d, n, h\rangle$, then *Bob* would have accepted the support in move 8 (as $\bar{\sigma}_s^{Bob}(S) = (0.6 + 0.84)/2 = 0.67$ which was above threshold for $\varepsilon(|2|) = 0.\dot{6}$), and fewer properties would have been disclosed.

5 Related Work

A number of different approaches have addressed the reconciliation of heterogeneous ontologies by using some form of rational reasoning. In [1] the notion of *ontology negotiation* as a communication protocol was presented that allowed agents to exchange ontological fragments by successively specifying the meaning of given entities. This is done on a per-needs basis and the further specification is only applied when the communication becomes ineffective. This idea formed the basis of the Anemone approach [5], which advocated a lazy, minimal protocol whereby agents exchange logical definitions in the attempt to define a minimal shared ontology with no information loss. However, it makes the assumption that the agents had perfect knowledge over the instances of their ontological models (i.e. the underlying approach was grounded through an extensional model), which could be used to induce a class description covering certain instances. Whilst these approaches resolve semantic interoperability through negation to achieve semantic *homogeneity*, other approaches attempt to align the heterogeneous ontologies through negotiation [11] or argumentation [12]. In [11], agents selectively exchange details of a priori privately known correspondences (unlike our approach that induces the correspondences collaboratively through repeated questions / answers), and propose repairs to address any emergent conservativity violations, resulting in alignments that are mutually acceptable to both agents without disclosing the full ontological model.

6 Conclusions

In this paper, we present initial work on a cognitive-speech based dialogue fragment designed to allow two agents to exchange knowledge about the ontological neighbourhood of a pair of concepts to determine the candidacy of a correspondence. An initiator asks questions about a potential match for its own concept to ascertain if there is sufficient evidence to support it, and the interlocutor through introspection accepts, rejects or seeks further or more compelling evidence to support the claim. This work is an iterative component within ongoing work on a larger dialogical framework that identifies a number of (potentially synergistic) correspondences between two ontologies that form an alignment targeted at specific joint task (*i.e.* the rationale behind the transaction between the two agents). The choice of suitable similarity metrics, as well as determining the depth of exploration of the neighbourhood, is currently under investigation.

References

1. Bailin, S., Truszkowski, W.: Ontology negotiation: How agents can really get to know each other. In: Innovative Concepts for Agent-Based Systems, LNCS, vol. 2564, pp. 320–334. Springer (2003)
2. Cheatham, M., Hitzler, P.: String similarity metrics for ontology alignment. In: The Semantic Web–ISWC 2013, pp. 294–309. Springer (2013)
3. Clark, H.H.: Using language. Cambridge university press (1996)
4. Clark, H.H., Schaefer, E.F.: Contributing to discourse. Cognitive science 13(2), 259–294 (1989)
5. van Diggelen, J., Beun, R.J., Dignum, F., van Eijk, R.M., Meyer, J.J.C.: Anemone: an effective minimal ontology negotiation environment. In: Proc. AAMAS '06. pp. 899–906 (2006)
6. Euzenat, J., Shvaiko, P.: Ontology Matching. Springer-Verlag (2007)
7. Grau, B.C., Dragisic, Z., Eckert, K., Euzenat, J., et al.: Results of the ontology alignment evaluation initiative 2013. In: Proc. 8th ISWC workshop on ontology matching (OM). pp. 61–100 (2013)
8. Grau, B.C., Motik, B.: Reasoning over ontologies with hidden content: The import-by-query approach. Journal of Artificial Intelligence Research (JAIR) 45, 197–255 (October 2012)
9. 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 (1975)
10. Hulstijn, J.: Dialogue models for inquiry and transaction. University of Twente (2000)
11. Jiménez-Ruiz, E., Payne, T.R., Solimando, A., Tamma, V.: Avoiding alignment-based conservativity violations through dialogue. In: Proc. OWLED '15 (2015)
12. Laera, L., Blacoe, I., Tamma, V., Payne, T., Euzenat, J., Bench-Capon, T.: Argumentation over ontology correspondences in MAS. In: Proc. AAMAS '07. pp. 1285–1292 (2007)
13. Payne, T.R., Tamma, V.: Negotiating over ontological correspondences with asymmetric and incomplete knowledge. In: Proc. AAMAS '14. pp. 517–524 (2014)
14. Searle, J.R.: Speech acts: An essay in the philosophy of language, vol. 626. Cambridge university press (1969)
15. Shvaiko, P., Euzenat, J.: Ontology matching: state of the art and future challenges. Knowledge and Data Engineering, IEEE Transactions on 25(1), 158–176 (2013)
16. 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)