

A Dialogue Protocol to Support Meaning Negotiation

(Extended Abstract)

Gabrielle Santos, Valentina Tamma, Terry R. Payne, Floriana Grasso
University of Liverpool, Liverpool L69 3BX, UK
{G.S.J.Santos, V.Tamma, T.R.Payne, floriana}@liverpool.ac.uk

ABSTRACT

Despite numerous efforts, the problem of dynamically reconciling heterogeneity within open distributed multi-agent systems is far from solved. As different systems often use their own vocabularies to express the content of communication messages (*ontologies*), semantic reconciliation requires some form of agreement over a shared model, obtained through an alignment whereby concepts in the requester's ontology are *mapped* (translated) into concepts in the respondent's one. This paper presents a dialogue that allows agents to reach an agreement over a correspondence between two entities in their respective ontologies in a decentralised fashion, without requiring prior knowledge over their ontological models.

Keywords

Ontology Matching & Alignment; Meaning Negotiation; Dialogues

1. INTRODUCTION

The emergence of annotated data and sophisticated mechanisms for representing formal data models has promoted the proliferation of agents that typically commit to their own knowledge model (*ontologies*). The success of agent interactions therefore depends on an agent's ability to successfully make sense of the intended meaning of the messages exchanged with other agents. To overcome the resulting ontological heterogeneity, semantic reconciliation is needed to facilitate seamless interaction, which relies on finding an agreement on the choice of *mappings* that translate concepts from a source ontology to a target ontology. This problem has been investigated by numerous research efforts [1, 2]; however these approaches require sharing the ontological model in part or fully, without guaranteeing if such mappings could be determined. If two vocabularies have no overlapping intended models, the mappings calculated are unlikely to support the meaningful translation of requests between agents, potentially hindering interoperability.

We address the reconciliation problem using negotiation, exploring how dialogue protocols can be used to determine mappings that satisfy the agents requirements and strategies. The use of dialogical models allows the agents to state their position over the correctness of a mapping in an asynchronous and distributed fashion, whilst maintaining control over the type of knowledge (class labels vs. ontological axioms) disclosed. We investigate reaching an agreement that facilitates the translation of one term to another

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.
Copyright © 2016, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

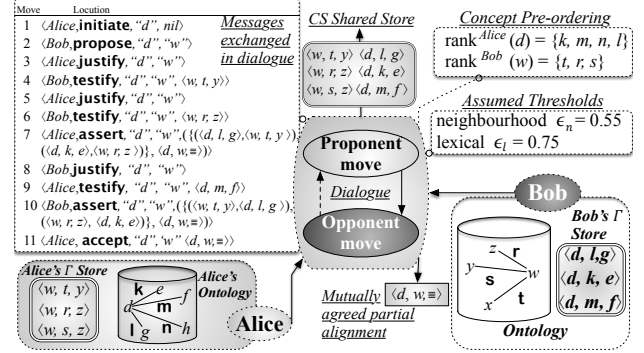


Figure 1: Dialogue Architecture using Alice and Bob described in the walkthrough example (Section 3).

in different vocabularies. These translations are computed opportunistically (*anytime*), and limit the information exchange only to what is pertinent, thus supporting a specific translation. Our main contribution is a dialogue based negotiation protocol that allows agents to propose a viable lexical mapping, and support this proposal with evidence in the form of ontological fragments. These are shared on a per-need basis, and hence the mechanism is purely opportunistic.

2. DIALOGUE MODULE SYSTEM

The proposed dialogue consists of a sequence of *moves*, \mathcal{M} , between the participating agents [3], where a move is a message exchanged between two agents x (the sender, or agent making the move), and \hat{x} (the recipient) at time s , and $m_s = \langle x, \tau, e, e' \rangle$. We denote τ as the move type s.t $\tau \in \mathcal{T}$, where $\mathcal{T} = \{\text{initiate}, \text{propose}, \text{assert}, \text{accept}, \text{reject}, \text{testify}, \text{justify}, \text{fail}\}$; e is the source entity being discussed (defined within the *initiate* move); e' is the candidate target entity (*i.e.* the entity that could be mapped to from e) or nil where there is no entity specified (in the case of the *initiate* and *fail* moves). The moves are presented in Table 1, complete with a description of their informal semantics. Both agents manage a public, joint knowledge base, or *Commitment Store*, *CS*, containing a trace of the moves uttered [3]; whereas, each agent manages its own private knowledge base, (*Gamma Store*, Γ), which stores knowledge regarding the ontological entities acquired from the other agent.

3. WALKTHROUGH EXAMPLE

In this section, we present an example dialogue between two agents, *Alice* and *Bob*. Each agent possesses a private ontological fragment (see Fig. 1), that models the entities they use to com-

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

Move	Pre-condition	Post-condition
$\langle x, \text{initiate}, e, \text{nil} \rangle$	<i>initiate</i> is uttered when a correspondence is desired for e , which must be in x 's ontology; and has not previously appeared in an <i>initiate</i> move.	The only permissible moves are <i>propose</i> (i.e. the opponent has a potential match for the entity e), or <i>fail</i> (i.e. no potential match for e can be found).
$\langle x, \text{propose}, e, e' \rangle$	The <i>propose</i> move is uttered when the sender x has some previously undisclosed entity e' in its ontology which is a <i>lexically viable</i> match for e .	Permissible moves: <i>justify</i> - the opponent seeks evidence to support the candidacy of $\langle e, e', \equiv \rangle$; <i>assert</i> - the opponent has sufficient evidence for the candidacy; or <i>reject</i> - there is no evidence to support the candidacy.
$\langle x, \text{assert}, e, e', A \rangle$	The sender has a candidate correspondence between e and e' ; if i) the two entities are believed to be a <i>lexically viable</i> match; and ii) if the premise of the previously undisclosed argument A is considered <i>acceptable</i> ; i.e. its aggregate similarity score is greater than or equal to the threshold ϵ_n .	Permissible moves: <i>accept</i> if the opponent accepts the correspondence in the claim given its own assessment of the support; or <i>justify</i> when the opponent rejects the support, and thus seeks additional or alternative evidence to support the claim.
$\langle x, \text{justify}, e, e' \rangle$	The sender x believes the opponent/proponent has undisclosed knowledge supporting the candidacy of $\langle e, e', \equiv \rangle$, and that a <i>lexically viable</i> match between e and e' exists, but does not have sufficient evidence to assert it.	The only permissible move is <i>testify</i> .
$\langle x, \text{testify}, e, e', \varphi \rangle$	The <i>testify</i> move is uttered in response to a <i>justify</i> move requesting evidence to support the candidacy of some correspondence between e and e' which is believed to be a <i>lexically viable</i> match.	The value of φ will be a new triple supporting $\langle e, e', \equiv \rangle$ (if this exists), otherwise no triple (i.e. \emptyset) will be returned. The triple is added to the recipient's (Γ). Permissible moves: <i>justify</i> , <i>assert</i> or <i>reject</i> .
$\langle x, \text{accept}, e, e', A \rangle$	The sender x believes a candidate correspondence $\langle e, e', \equiv \rangle$ is <i>lexically viable</i> and there is sufficient evidence to support it. The argument A is the one appearing in the previous move (which should be an <i>assert</i> move).	Once the move has been made, the argument is added to the commitment store. The only move that is permissible following a <i>accept</i> move is an <i>initiate</i> move, whereby a new source entity is considered.
$\langle x, \text{reject}, e, e' \rangle$	Although x believes that there could be a <i>lexically viable</i> match between the two entities e and e' , it does not have sufficient evidence to assert a correspondence; and it no longer believes the other agent has undisclosed knowledge that could support its candidacy.	Permissible moves: If x is the opponent, the proponent can either <i>assert</i> an alternate argument for the correspondence, or seek further evidence using <i>justify</i> . If x is the proponent, it can utter <i>propose</i> if another entity could match e , or <i>fail</i> if it is no longer possible to find a correspondence for e .
$\langle x, \text{fail}, e, \text{nil} \rangle$	Sender x can't find an undisclosed target entity e' in its ontology that is a <i>lexically viable</i> match for e . Thus no possible match for e could be found.	The only move that is permissible following a <i>fail</i> move is an <i>initiate</i> move, whereby a new source entity is considered.

municate. We assume they use the same metric, σ_n , to assess the similarity of the neighbourhood of a concept, expressed in terms of triples, and the same lexical¹ similarity σ_l . Agents generate a pre-ordering of properties for each concept e , using some private function $\text{rank}()$. Each agent attempts to identify a *premise*, Pr , which consists of pairs of triples believed to be structurally similar. A neighbourhood similarity metric $\sigma_n(Pr)$ was chosen that calculates the average structural similarity $\bar{\sigma}_s$ of the triple pairs in some premise Pr , with a coefficient increasing asymptotically as the cardinality of Pr increases; i.e. $\sigma_n(Pr) = \bar{\sigma}_s \times (1 - \frac{1}{2(|Pr|+1)})$. We assume a neighbourhood threshold $\epsilon_n = 0.55$ and a lexical threshold $\epsilon_l = 0.75$. This example dialogue is presented in (Fig. 1).

Move 1: *Alice* utters an *initiate* move stating she wants to discuss a possible match for concept d . **Move 2:** *Bob* agrees to cooperate by identifying w as the most (lexically) similar concept to d in his ontology, with a lexical similarity $\sigma_l^{Bob}(\langle d, w, \equiv \rangle, \langle w, r, s \rangle) = 0.82$, which is above threshold ϵ_l . He therefore responds with $\langle \text{Bob}, \text{propose}, \langle d, w, \equiv \rangle, \langle w, r, s \rangle \rangle$. **Move 3:** *Alice* now knows that $\langle d, w, \equiv \rangle$ is a potential correspondence c (based on *Bob*'s lexical similarity claim). She checks that the lexical similarity is above threshold ($\sigma_l^{Alice}(d, w) = 0.79$), and asks *Bob* to justify the candidacy of c . So far, no agents has support for c ; i.e. $Pr = \emptyset$. **Move 4:** *Bob* determines the next property with w as the subject that has not yet been disclosed, and thus is not in the CS . As his ranking for w is $\text{rank}^{Bob}(w) = \{t, r, s\}$, *Bob* shares the highest ranked property t relating w to y . **Move 5:** *Alice* determines if there is sufficient support for c . She realises that $\langle d, l, g \rangle$ is the most similar triple she has to the one *Bob* disclosed in move 4, where $\sigma_s^{Alice} = 0.66$. She calculates that the premise $Pr = \{(\langle d, l, g \rangle, \langle w, t, y \rangle)\}$ has a neighbourhood similarity $\epsilon_n = 0.495$. She will only *assert* an argument for c if this is above the threshold $\epsilon_n = 0.55$. However, as this not the case, she requests additional evidence to justify c .

Move 6: *Bob*'s next highest undisclosed property with the subject w is r , and so he shares $\langle w, r, z \rangle$. **Move 7:** *Alice* checks to see if she has a similar triple to this; and finds two candidates. She selects $\langle d, k, e \rangle$ as the similarity is higher than $\langle d, n, h \rangle$. She adds

this to Pr and calculates the neighbourhood similarity which is $\sigma_n^{Alice} = 0.56$, which (from *Alice*'s perspective) is above threshold. She therefore *asserts* the correspondence $c = \langle d, w, \equiv \rangle$, given the premise $Pr = \{(\langle d, l, g \rangle, \langle w, t, y \rangle), (\langle d, k, e \rangle, \langle w, r, z \rangle)\}$.

Move 8: *Bob* can either *accept* this premise if $\sigma_n^{Bob}(Pr)$ is above threshold or request further evidence for c . He calculates the neighbourhood similarity (from his perspective), and finds that $\sigma_n^{Bob} = 0.53$, which is below threshold. *Bob* is also aware other triples for the concept w that do not appear in Pr , and asks *Alice* for further evidence. **Move 9:** *Alice* shares $\langle d, m, f \rangle$ as m is her highest ranked, non-disclosed property for the domain concept d (property k was ranked higher but has already been disclosed). **Move 10:** *Bob* recalculates the mean similarity for the new support (inclusive of the triple shared in Move 9): $\sigma_n^{Bob} = 0.551$, now above threshold, *Bob* accepts the candidacy of c . He can now *assert* the new argument for c given the new premise Pr . **Move 11:** *Alice* confirms from her perspective, $\sigma_n^{Alice} = 0.59$, which is also above threshold, and accepts the candidacy of the correspondence $\langle d, w, \equiv \rangle$.

4. CONCLUSIONS

We present our initial work on a dialogue based mechanism that allows agents to reach an agreement in a decentralised fashion over a correspondence between two entities in their respective ontologies, without having any prior information over the ontological structures that the agents commit to. This work is part of an ongoing effort over a larger dialogical framework that identifies a number of (potentially synergistic) correspondences between two ontologies, thus dynamically forming agreement on an alignment.

REFERENCES

- [1] S. Bailin and W. Truszkowski. Ontology negotiation: How agents can really get to know each other. In *Innovative Concepts for Agent-Based Systems*, pages 320–334. 2003.
- [2] E. Jiménez-Ruiz, T. R. Payne, A. Solimando, and V. Tamma. Limiting logical violations in ontology alignment through negotiation. In *Proc. KR'16*, 2016.
- [3] D. Walton and E. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY series in Logic and Language. State University of New York Press, 1995.

¹The metrics σ_n and σ_l may exploit lexical similarity of labels for concepts and roles, but are not defined here. In addition, the concept and role names have replaced by single character identifiers.