I-SSA: Interaction-Situated Semantic Alignment

Manuel Atencia and Marco Schorlemmer

IIIA, Artificial Intelligence Research Institute
CSIC, Spanish National Research Council
Bellaterra (Barcelona), Catalonia, Spain

Abstract. We tackle semantic heterogeneity in multi-agent communication by looking at semantics related to interaction in order to avoid dependency on a priori semantic agreements. Our underlying claim is that semantic alignment is often relative to the particular interaction in which agents are engaged, and that in such cases the interaction state should be taken into account and brought into the alignment mechanism. We provide a formal foundation along with an alignment mechanism and protocol based on it.

1 Introduction

In multi-agent communication one usually assumes that all agents make use of a shared terminology with the same meaning for message passing. If agents, nevertheless, are engineered separately one has to foresee that, when they interact, they will most likely use different terminology in their respective messages, and that, if some terms coincide, they may not have the same meaning for all agents participating in an interaction. This is essentially the problem of semantic heterogeneity.

One early solution goes with agreeing upon a common ontology for the domain in which interoperability has to take place [1]. Current state-of-the-art approaches tackling semantic heterogeneity, however, no more seek to agree on one shared global ontology, but instead attempt to establish correspondences between varying terminologies [2,3]. In these approaches, matching is generally performed outside the context of the integration. Moreover, most current ontology matching techniques follow a classical functional approach, taking two or more ontologies as input and producing a semantic alignment of ontological entities as output. This involves several drawbacks. On the one hand, it limits the dynamism and openness; on the other, it keeps matching out of the context of the interaction, since semantic similarity of terms is established in an interaction-independent fashion. But the meaning of certain terms are often very interaction-specific. For example, the semantic similarity that exists, in the context of an auction, between the Spanish term "remate" and the English expression "winning bid" is difficult to establish if we are left to rely solely on syntactic-based or structural matching techniques, or even on external sources such as dictionaries and thesauri. The term "remate" may have many different senses, and none of them may hint at its meaning as "winning bid". However, it actually has this very precise meaning when uttered at a particular moment of the interaction happening during an auction.

In this paper, we attempt to overcome the mentioned drawbacks and present a very parsimonious approach to the problem of semantic heterogeneity in multi-agent communication with the aim of complementing the previous solutions applied so far. Our

R. Meersman and Z. Tari (Eds.): OTM 2008, Part I, LNCS 5331, pp. 445–455, 2008.

claim is that semantic alignment is often also relative to the particular interaction in which agents are engaged in, and, more specifically, to the particular state of the interaction. In such cases the interaction state should be taken into account and brought into the alignment mechanism. We shall address the case in which agents need to establish the semantic relationships with terminologies of other agents on the grounds of their communication within a specific interaction. We call this approach *Interaction-Situated Semantic Alignment* (or I-SSA, in short).

An Example: a Sealed-Bid Auction. In a sealed-bid auction, after the auctioneer has announced the start of a round for auctioning a particular good, bidders are given a period of time to submit their bids (without other bidders knowing it). After that period the auctioneer announces the winner, namely the bidder that submitted the highest bid. In certain cases, the auctioneer may decide to withdraw a good instead (for example, if no bids where submitted). Hence the interaction that unfolds is as follows. In the initial state of the interaction, bidders wait for the auctioneer to send a message announcing the *start of round* for a particular good GID at a reserve price RP with bidding time BT. This message passing causes a state transition in the interaction to a state in which bidders are allowed to send their *bids* O for good GID. From the auctioneer viewpoint, the interaction remains in this state until the bidding time BT has elapsed, in which case the interaction moves to a state in which bidding messages are no more expected and in which the auctioneer is supposed to either send a message informing the bidders that the good GID has been *sold* to bidder W for the price P, or to send a message informing that good GID has been *withdrawn*. Either of these messages makes the interaction state change to the initial state that is also the final state in this case. From the bidders viewpoint, however, if they have submitted a *bid* O, they consider the interaction to have changed to a state in which they cannot send bids any more, but where they wait for a message from the auctioneer informing about the outcome of the round. Alternatively, they may also assume this state transition without themselves having submitted a bid.

The above interaction model for a sealed-bid auction can be formally specified in numerous ways. One way is by means of finite state automata (e.g., electronic institutions [4]). Figures 1 and 2 illustrate the message-passing behaviour of an agent in the role of an auctioneer and of a bidder, respectively. Transitions between states are labelled by means of illocutions, which are tuples consisting of an illocutionary particle, the identifier of the sender together with the role it is playing, the identifier of the receiver together with the role it is playing, and the content of the message uttered. We may label transitions also with *timeout* or with λ denoting state transitions not caused by message passing. Variables in messages are written in uppercase letters and get their values in those illocutions in which they occur preceded by a question mark (?), and these values are subsequently used in those illocutions in which the corresponding variable occurs preceded by an exclamation mark (!).

When auctioneers and bidders interact by message passing, an interaction unfolds implicitly that contains more detail than the ones specified in Figures 1 or 2. These capture only a partial view of the actual *global interaction*, which matches together all messages occurring in illocutions with the same illocutionary particle, sender, and receiver, and that trigger the same state transition (see Figure 3). In addition, each actual state of the global interaction should have a corresponding state in each of the role

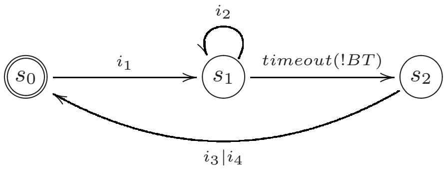

$$i_1 = \langle inform, (?A : auctioneer), (?B : bidder), start(?GID, ?BT, ?RP)\rangle$$
$$i_2 = \langle commit, (!B : bidder), (!A : auctioneer), bid(!GID, ?O)\rangle$$
$$i_3 = \langle inform, (!A : auctioneer), (!B : bidder), sold(!GID, ?P, ?W)\rangle$$
$$i_4 = \langle inform, (!A : auctioneer), (!B : bidder), withdrawn(!GID)\rangle$$

Fig. 1. Interaction model for the auctioneer role

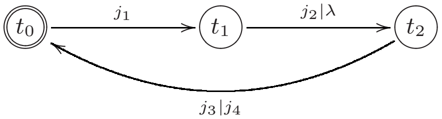

$$j_1 = \langle inform, (?A : auctioneer), (?B : bidder), start(?GID, ?BT, ?RP)\rangle$$
$$j_2 = \langle commit, (!B : bidder), (!A : auctioneer), bid(!GID, ?O)\rangle$$
$$j_3 = \langle inform, (!A : auctioneer), (!B : bidder), sold(!GID, ?P, ?W)\rangle$$
$$j_4 = \langle inform, (!A : auctioneer), (!B : bidder), withdrawn(!GID)\rangle$$

Fig. 2. Interaction model for the bidder role

interaction models. This means that the states of the interaction models in Figures 1 and 2 are projections of states of a global interaction model, i.e., the view from the perspective of an auctioneer and of a bidder.

Imagine now the interaction model of Figure 2, but for a Spanish-speaking bidder, with exactly the same illocutions but this time “start”, “bid”, “sold” and “withdrawn” turn to “ronda”, “postura”, “remate” and “sin_ganador”, respectively. The Spanish-speaking bidder initially expects a “ronda” message from the auctioneer. The English-speaking auctioneer initially is supposed to broadcast a “start” message to bidders. When this illocution is uttered the Spanish-speaking bidder may safely assume that “start” means “ronda”, which makes the interaction change to the state in which the English-speaking auctioneer expects “bid” messages from buyers and the Spanish-speaking bidder is supposed to either send a “postura” or change state without sending or receiving any message. Consequently, if “postura” is uttered the English-speaking auctioneer can safely assume that “postura” means “bid”. Notice that these equivalences stem from the assumption that the local states of the auctioneer and bidder are always projections from the same state of the actual global interaction and follow projections the same state transition when a illocution is uttered.

In Section 2, we provide a formal definition of “global interaction model” through the idea of a product of interaction models, representing all compatible state transitions, and from it we define the notion of semantic equivalence that arises from compatible interactions. We shall treat messages as propositions, however, i.e., as grounded atomic sentences, leaving the generalisation to first-order sentences for future work. All these notions are succinctly presented due to size limitations but the reader can find further explanations in [5]. In Section 3 we thoroughly describe an alignment mechanism, as well as an alignment protocol that agents can follow in practice.

2 Formalising Interaction Models and Their Relations

We model a multi-agent system as a set MAS of agents. Each agent in MAS has a unique identifier and may take one (or more) roles in the context of an interaction. Let Role be the set of roles and Id the set of agent identifiers. We write $(id : r)$, with $r \in \mathit{Role}$ and $id \in \mathit{Id}$, for the agent in MAS with identifier id playing role r.

Each agent is able to communicate by sending messages from a set M, which is local to the agent. We assume that a set $\mathfrak{I}_{\mathrm{P}}$ of *illocutionary particles* (such as "inform", "ask", "advertise", etc.) is shared by all agents (see KQML [6] or FIPA ACL [7]).

Definition 1. *Given a non-empty set M of messages, the set of* illocutions generated by *M, denoted by $\mathfrak{I}(M)$, is the set of all tuples $\langle \iota, (id : r), (id' : r'), m \rangle$ with $\iota \in \mathfrak{I}_{\mathrm{P}}$, $m \in M$, and $(id : r), (id' : r')$ agents such that $id \neq id'$. If $i = \langle \iota, (id : r), (id' : r'), m \rangle$ is an illocution then $(id : r)$ is the* sender *of i and $(id' : r')$ is the* receiver *of i. In addition, $\langle \iota, (id : r), (id' : r') \rangle$ and m are called the* head *and* content *of i, respectively.*

2.1 Interaction Models

We model an interaction model as a (partial) deterministic finite-state automaton whose transitions are labelled either with illocutions, or with special transitions such as, for instance, timeouts or null transitions (also λ-transitions):

Definition 2. *An* interaction model *is a tuple $\mathrm{IM} = \langle Q, q^0, F, M, C, \delta \rangle$ where:*

- *Q is a finite set of* states*,*
- *q^0 is a distinguished element of Q named the* initial state*,*
- *F is a non-empty subset of Q which elements are called* final states*,*
- *M is a finite non-empty set of* messages*,*
- *C is a finite set of* special transitions*, and*
- *δ is a partial function from $Q \times (\mathfrak{I}(M) \cup C)$ to Q called the* transition function.

Every interaction model is related with an automaton in a natural way. The notion of history associated to an interaction model presented bellow is very similar to a string accepted for an automaton. The clear difference is that the former one takes into account the states explicitly.

Definition 3. *Let* IM *be an interaction model, where $\mathrm{IM} = \langle Q, q^0, F, M, C, \delta \rangle$. An* IM-history *or* history *associated with* IM *is a finite sequence:*

$$h = q^0, \sigma^1, q^1, \ldots, q^{k-1}, \sigma^k, \ldots, q^{n-1}, \sigma^n, q^n$$

where $q^n \in F$ and for each k: $q^k \in Q$, $\sigma^k \in \mathfrak{I}(M) \cup C$ and $\delta(q^{k-1}, \sigma^k) = q^k$.

2.2 The Communication Product

As hinted in the previous section, we shall use the algebraic product of two interaction models to capture all possible interactions between agents. In general, a product of two objects is the natural algebraic construction that represents all possible behaviours of the combination of those two objects. The *communication product* (CP) defined below, thus, captures the global interaction with respect to the message-passing behaviour of agents of two interaction models. It is not an unconstrained product, since it takes into account the compatibility of illocutions and special transitions in terms of illocutionary particles, senders, and receivers.

Definition 4. *Let* IM_1 *and* IM_2 *be two interaction models,* $\mathrm{IM}_i = \langle Q_i, q_i^0, F_i, M_i, C_i, \delta_i\rangle$ $(i = 1, 2)$. *The* communication product *of* IM_1 *and* IM_2*, denoted by* $\mathrm{IM}_1 \otimes \mathrm{IM}_2$*, is the interaction model* $\langle Q, q^0, F, M, C, \delta\rangle$ *where:*

- Q *is the Cartesian product of* Q_1 *and* Q_2*; specifically,* Q *states are all possible ordered pairs* $\langle q_1, q_2\rangle$ *with* $q_1 \in Q_1$ *and* $q_2 \in Q_2$*,*
- *the initial state* q^0 *is the pair* $\langle q_1^0, q_2^0\rangle$*,*
- F *is the Cartesian product of* F_1 *and* F_2*,*
- M *the Cartesian product of* M_1 *and* M_2*,*
- $C = (C_1 \times C_2) \cup (C_1 \times \{\epsilon\}) \cup (\{\epsilon\} \times C_2)$*; and finally*
- δ *is defined as follows:* $\langle q_1', q_2'\rangle = \delta(\langle q_1, q_2\rangle, \sigma)$ *if*
  - σ *is an illocution* $\langle \iota, (id : r), (id' : r'), \langle m_1, m_2\rangle\rangle$ *and* $q_i' = \delta_i(q_i, \langle \iota, (id : r), (id' : r'), m_i\rangle)$ *for every* i*,*
  - $\sigma = (c_1, c_2)$ *and* $q_i' = \delta_i(q_i, c_i)$ *for every* i*,*
  - $\sigma = (c_1, \epsilon)$ *and* $q_1' = \delta_1(q_1, c_1)$ *and* $q_2' = q_2$*,*
  - $\sigma = (\epsilon, c_2)$ *and* $q_2' = \delta_2(q_2, c_2)$ *and* $q_1' = q_1$*.*

Notice that ϵ symbols are paired with agent special transitions. They capture the idea that though the global interaction state may change, this could not be the case for one agent interaction model.

Example 1. The communication product of interaction models for the auctioneer and bidder roles is depicted in Figure 3. We only write those illocutions that take part in the language generated by the automaton (k_1 to k_6).

2.3 Semantic Alignment through the Communication Product

In order to establish the relations among messages, we look at the whole set of histories associated with the communication product. Messages of different interaction models are semantically related if they are paired in illocutions whose utterance make the interaction reach a final state (i.e., make the interaction succeed) according to the global interaction determined by the communication product. This is formally given below.

Definition 5. *Let* IM_1 *and* IM_2 *be two interaction models,* $\mathrm{IM}_i = \langle Q_i, q_i^0, F_i, M_i, C_i, \delta_i\rangle$ $(i = 1, 2)$. *Let* $m \in M_1$ *and* $m^1, \ldots, m^n \in M_2$*. We write:*

$$\langle 1, m\rangle \sqsubseteq \langle 2, m^1\rangle \sqcup \cdots \sqcup \langle 2, m^n\rangle$$

if for all histories h *of the communication product* $\mathrm{IM}_1 \otimes \mathrm{IM}_2$*, if the illocution* $\langle \iota, (id : r), (id' : r'), \langle m, m' \rangle \rangle$ *appears in* h *then there exists* k $(k = 1, \dots, n)$ *with* $m' = m^k$*. Similar considerations apply to* $\langle 2, m \rangle \sqsubseteq \langle 1, m^1 \rangle \sqcup \cdots \sqcup \langle 1, m^n \rangle$.

Example 2. We have the following relationships among messages (without pairing them with natural numbers): $start \equiv ronda$, $bid \equiv postura$, $sold \sqsubseteq remate \sqcup sin\_ganador$, $withdrawn \sqsubseteq remate \sqcup sin\_ganador$, $remate \sqsubseteq sold \sqcup withdrawn$, and $sin\_ganador \sqsubseteq sold \sqcup withdrawn$.

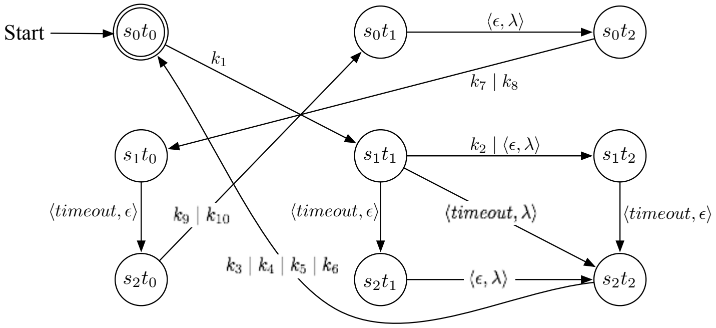

$k_1 = \langle inform, (a : auctioneer), (b : bidder), \langle start, ronda \rangle \rangle$
$k_2 = \langle commit, (b : bidder), (a : auctioneer), \langle bid, postura \rangle \rangle$
$k_3 = \langle inform, (a : auctioneer), (b : bidder), \langle sold, remate \rangle \rangle$
$k_4 = \langle inform, (a : auctioneer), (b : bidder), \langle sold, sin\_ganador \rangle \rangle$
$k_5 = \langle inform, (a : auctioneer), (b : bidder), \langle withdrawn, remate \rangle \rangle$
$k_6 = \langle inform, (a : auctioneer), (b : bidder), \langle withdrawn, sin\_ganador \rangle \rangle$

Fig. 3. The communication product

3 Aligning While Interacting

As said before, interaction models specify the space of interactions that are allowed, and its communication product captures the entire space of actual interactions when combining particular ones. The above semantic relationships are, thus, those justified by the entire space of actual interactions. This product, however, may be no accessible to agents. This is the case when, for example, interaction models are not completely open for inspection, so agents are only aware of their local ones. Furthermore, interaction models could be of a size that the product computation became unfeasible.

It is therefore necessary to provide agents with a mechanism to discover the above semantic relationship while interactions unfold —in the sort of manner as intuitively described for our example in Section 1— assuming that for all agents participating in the interaction, the state of the interaction they perceive stems from the actual global state (i.e., their locally managed states are projections of the actual global state), and this throughout the entire interaction.

3.1 The Alignment Protocol

With agents knowing that they follow different interaction models and that semantic mismatches are likely to occur, communication requires to be processed in another level. For this reason, we define an *alignment protocol* that links interaction models. This protocol is seen as a meta-protocol through which the communication is carried out: any communication act regarding the lower level becomes ineffective and has an effective counterpart according to the meta-level. The alignment protocol (from here on AP) is depicted in Figure 4. Let us explain it in detail.

There are two final states by name of letters s and u. If the state s is reached, then the interaction is considered *successful*, otherwise is considered *unsuccessful*. In this sense, we distinguish for the moment only two sorts of interactions. Regarding transitions, all of them are listed below the figure except one that has a special status. Notice that agents can adopt only one role, namely, the 'aligner' role, or $algn$ in short. There are two kind of messages: `failure` and `final_state`. Moreover, the former one can be tagged with the illocutionary particle $inform$, and the later one with $inform$, $confirm$ and $deny$.

The reader can check that AP is not an interaction model as defined in Definition 2. The reason is that the following illocution contains variables:

$$\langle utter, (?A : algn), (?B : algn), I\rangle \tag{1}$$

Recall that our simplified definition of an illocution does not consider variables, but nonetheless we will consider (1) also to be an illocution. A and B are identifier variables and I is an illocution variable. So (1) can be considered a meta-illocution since its content, in turn, is an illocution. We state that it can be instantiated with expressions of the form $\langle utter, (id_i : algn), (id_j : algn), \mu\rangle$ where $\mu = \langle \iota, (id_i : r), (id_j : r'), m\rangle$ is an illocution of IM_i. Therefore, μ sender and receiver must be equal to A's instance and B's instance, respectively. Furthermore, let us stress that μ has to come from A's

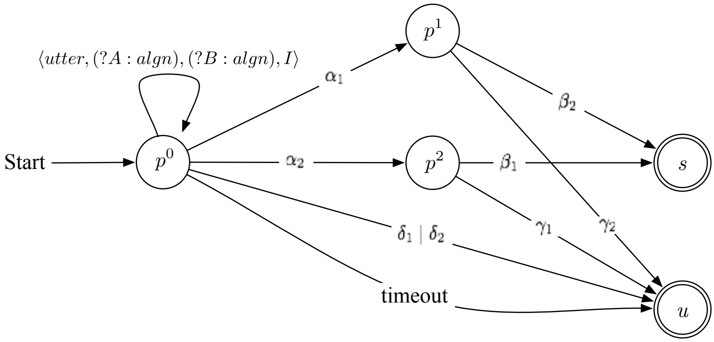

$$\begin{aligned}
\alpha_i &= \langle inform, (id_i : algn), (id_j : algn), \texttt{final\_state}\rangle\\
\beta_i &= \langle confirm, (id_i : algn), (id_j : algn), \texttt{final\_state}\rangle\\
\gamma_i &= \langle deny, (id_i : algn), (id_j : algn), \texttt{final\_state}\rangle\\
\delta_i &= \langle inform, (id_i : algn), (id_j : algn), \texttt{failure}\rangle
\end{aligned}$$

Fig. 4. The alignment protocol

instance interaction model. Consequently, the choice of 'utter' as illocutionary particle seems natural. This performative expresses the sender attitude with respect to its own interaction model. If A_j receives $\langle utter, (id_i : algn), (id_j : algn), \mu\rangle$ then it can assume that A_i has decided to utter μ according to IM_i. At this point the alignment protocol dynamics and alignment mechanism come into play.

3.2 The Alignment Protocol Dynamics

Each agent follows both the alignment protocol and its own interaction model, whilst effective communication is done in accord to AP as it is emphasised above.

When agents agree to initiate an interaction, both of them are in state p^0 wrt AP. In addition, agent A_i ($i = 1, 2$) is in state q_i^0 wrt IM_i.

Imagine agent A_i is in state q_i, where q_i is an arbitrary element of Q_i. There can be several possibilities.

1. A_i decides to utter $\mu = \langle \iota, (id_i : r), (id_j : r'), m\rangle\rangle$ in accord with IM_i, where $\mu \in \delta_i(q_i, \cdot)$.[1] The communication act must be carried out via AP so agent A_i sends illocution $\langle utter, (id_i : algn), (id_j : algn), \mu\rangle$ to A_j. Therefore, the state remains the same in the AP context, whereas q_i turns to $q_i' = \delta_i(q_i, \mu)$ in the IM_i context.
2. A_i prompts a state change by a special transition $c_i \in C_i$ in the IM_i context. Thus q_i turns to $q_i' = \delta_i(q_i, c_i)$. This action is not reflected in AP since it does not entail any communication act.
3. A_i receives $\langle utter, (id_j : algn), (id_i : algn), \mu\rangle$ where $\mu = \langle \iota, (id_j : r), (id_i : r'), m)\rangle$ with regard to AP. Recall that from A_i's viewpoint, m is a foreign message so it is considered semantically different from all local messages.

   Now, the **key issue** is that m is to be mapped with one of those messages that A_i expects to receive at state q_i in the IM_i context. Furthermore, we can make a selection and just consider those messages encased in illocutions which head is equal to that of μ. In this way, A_i is to choose an element of the following set:

   $$\mathcal{R} = \{a \mid \langle \iota, (id_j : r), (id_i : r'), a\rangle \in dom(\delta_i(q_i, \cdot))\}$$

   There can be two possibilites: $\mathcal{R}$ is empty or not.

   3.1 As long as $\mathcal{R}$ is not empty, A_i can select an element a of $\mathcal{R}$ making use of the **alignment mechanism** explained further below. So q_i turns to $q_i' = \delta_i(q_i, \nu)$ where $\nu = \langle \iota, (id_j : r), (id_i : r'), a)\rangle$.

   3.2 In case $\mathcal{R}$ is empty, then no mapping is possible. The interaction is considered unsuccessful. In order to state it, A_i sends a failure message to A_j by uttering $\delta_i = \langle inform, (id_i : algn), (id_j : algn), \texttt{failure}\rangle$. Thus p_0 turns to u in the AP context.
4. If q_i is a final state and A_i considers the interaction finished, it can send illocution $\alpha_i = \langle inform, (id_i : algn), (id_j : algn), \texttt{final\_state}\rangle$ to A_j. In this case, p_0 turns to p_i and A_i expects to receive illocutions β_j or γ_j ($j \neq i$), either confirming or denying the interaction end, respectively. If it receives β_j, then p_i turns to s

---

[1] $\delta_i(q_i, \cdot)$ is the function defined from $\Sigma_i = \Im(M_i) \cup C_i$ to Q_i in a natural way.

and the interaction is considered successful; if it receives γ_j, p_i turns to u and the interaction is considered unsuccessful.

5. Finally, we have to take into account the possibility of a deadlock. This is the case when, for example, successive mappings have led the agents to states where both of them only can receive. In order to avoid deadlocks, the special transition *timeout* is linked to the initial state p_0 in AP. When a specific period of time is exceeded, this transition leads agents to finish the interaction considered unsuccessful.

3.3 The Alignment Mechanism

As it is mentioned above, the alignment mechanism is called whenever a message is received. In a nutshell, the alignment mechanism (henceforth AM) is based on the three following assertions:

- every foreign message is associated with a categorical variable ranging over local messages, likewise a variable assignment represents a mapping element,
- AM computes frequency distributions of all these variables on the basis of past successful interactions,
- agents mapping choices are determined by virtue of these distributions.

Let us give the details of AM. Assume agent A_i tackles a situation like the one described above in case 3.1. Message m is associated with a variable X that takes values in M_i. The equality $X = a$ represents a *mapping element* (the fact that m is mapped to a), also written $[m/a]$. If there is no past experience, $[m/a]$ is chosen with probability $p = \frac{1}{n}$, where n is the cardinality of $\mathcal{R}$.

Now, things are different as long as agents have interacted successfully in the past. In order to reason about past experiences, agents have to keep track of these ones. A *history* is a sequence of the form:

$$h = q_i^0, \sigma_i^1, q_i^1, \ldots, q_i^{k-1}, \sigma_i^k, \ldots, q_i^{n-1}, \sigma_i^n, q_i^n$$

computed recursively as follows:

- q_i^0 is the initial state of IM$_i$, and
- if A_i is in case 1, then $[\imath, q_i']$ is queued in h,
- if A_i is in case 2, then $[c_i, q_i']$ is queued in h,
- if A_i is in case 3.1, $[\langle \iota, (id_j : r), (id_i : r'), [m/a]) \rangle, q_i']$ is then queued in h,
- q_i^n is a final state of IM$_i$.

Notice that unsuccessful interactions are not considered.

Agents resort to all past histories in order to calculate the frequency distributions. Remember foreign messages do not occur in isolation: each message is the content of a specific illocution which is received at a particular state. To capture this dependency two more variables are considered: Q and H. Q takes values in the set of states Q_i and H can be instantiated with heads of illocutions.

So coming back to a situation like the one described in 3.1, agent A_i wonders whether $X = a$, where a varies in M_i, given that m is the content of an illocution with head $H = \langle \iota, (id_j : r), (id_j : r') \rangle$ that has been received at state $Q = q_i$. Using the corresponding frequency distribution:

$$f_r[X = a \mid Q = q_i, H = \langle \iota, (id_i : r), (id_j : r') \rangle] = \frac{v}{w} \in \mathbb{Q}$$

and $[m/a]$ is chosen with probability $p = \frac{v}{w}$, or alternatively $p = \frac{1+v}{n+w}$. Note that the former option prevents the agent from discovering new mapping elements, whereas the later one makes it possible.

3.4 Semantic Alignment through the Alignment Mechanism

Hence the alignment mechanism described above helps agents to interact successfully. Note that agent messages are related as more interactions are completed. Now, in what follows, we pin down these semantic relationships in a logical fashion and we finally compare them with the ones deduced from the communication product as stated in Definition 5.

Let us assume that agent A_i $(i = 1, 2)$ has generated a family $\mathcal{F} = \{\mathcal{F}_X \mid X \in \mathfrak{X}\}$ of frequency distributions using AM. Recall that each variable X represents a foreign message so $\mathfrak{X}$ stands for the whole set of received messages. If $f_r[X = m^k] \neq 0$ where $m^k \in M_i$ for $k = 1, \ldots, n$ and $f_r[X = m'] = 0$ for $m' \in M_i \setminus \{m^1, \ldots, m^n\}$, then the following holds:

$$\langle j, m \rangle \sqsubseteq \langle i, m^1 \rangle \sqcup \ldots \sqcup \langle i, m^n \rangle$$

where m is the foreign message represented by X. The idea behind this subsumption is that according to $\mathcal{F}_X$ only the mapping elements $[m/m^1], \ldots, [m/m^n]$ have triggered states transitions making the interaction reach a final state eventually. It is also possible to discriminate between the disjunction members on the basis of $\mathcal{F}_X$, if $f_r[X = m'] = t \in (0, 1]$, it holds:

$$\langle j, m \rangle \sqsubseteq \langle i, m' \rangle [t]$$

The real number t expresses the *confidence degree* of the mapping element $[m/m']$. Finally, the *semantic alignment* is made up of the set of all these expressions.

Example 3. With the help of the alignment mechanism, the auctioneer can deduce (Spanish to English): $\mathit{postura} \sqsubseteq \mathit{bid}\,[1.0]$; whereas the bidder can deduce (English to Spanish): $\mathit{start} \sqsubseteq \mathit{ronda}\,[1.0]$, $\mathit{sold} \sqsubseteq \mathit{remate}\,[t_1]$, $\mathit{sold} \sqsubseteq \mathit{sin\_ganador}\,[t_2]$, $\mathit{withdrawn} \sqsubseteq \mathit{remate}\,[s_1]$, and $\mathit{withdrawn} \sqsubseteq \mathit{sin\_ganador}\,[s_2]$, with $t_1 + t_2 = 1$ and $s_1 + s_2 = 1$.

The following theorem illustrates the relationship between the semantic alignment computed via the alignment mechanism and the one drawn from the communication product. The proof is straightforward. It suffices to show that every history computed by the alignment mechanism comes from a history associated with the communication product (the former one is essentially a projection of the latter one).

Theorem 1. *Assume that $\langle j, m \rangle \sqsubseteq \langle i, m^1 \rangle \sqcup \ldots \sqcup \langle i, m^n \rangle$ is a subsumption expression computed using the alignment mechanism. Then if $\langle j, m \rangle \sqsubseteq \langle i, w^1 \rangle \sqcup \ldots \sqcup \langle i, w^r \rangle$ is deduced from the communication product, we can assure that $n \leq r$ and there exist indices $s_1, \ldots, s_n$ with $1 \leq s_k \leq r$ such that $m^k = w^{s_k}$ for $k = 1, \ldots, n$.*

4 Conclusions

In [8] we described an alignment process by which two agents can establish the semantic relationship between their vocabularies based on the assumption that mismatching terms describe a partial perspective of a shared physical environment state not entirely accessible to any of the two agents. In this paper, however, agents do not share a physical environment but the same interaction. Hence their "environment" is captured by the coproduct interaction model that captures the entire space of actual interactions, but which are not accessible to agents in general. An uttered illocution provides a "description" of the interaction state, because its utterance "means" that the illocution was allowed in the current interaction state according to the partial perspective of the uttering agent. An agent receiving the illocution can now compute a semantic alignment based on the assumption that both agents where sharing the same interaction state. One advantage of this approach is that it takes into account meaning that is very interaction-specific and cannot be derived from sources that are external to the interaction. In this sense we see it as a complement to current semantic alignment techniques as it may provide valuable information for pruning the search space or disambiguating the results of candidate semantic alignments. Our work shares with that of Besana and Robertson [9] the insight that semantics is often interaction-specific. Unlike in our work, however, Besana and Robertson aim at reducing the search space of possible a priori mappings between ontological entities (in a classical sense).

Acknowledgements. This work is supported under EU STREP OpenKnowledge (FP6-027253), Spanish projects IEA (TIN2006-15662-C02-01) and Agreement Technologies (CSD2007-0022), and the Generalitat de Catalunya (2005-SGR-00093). Schorlemmer is supported by a *Ramón y Cajal* Fellowship from Spain's Ministry of Science and Innovation, and the European Social Fund.

References

1. Gruber, T.R.: A translation approach to portable ontology specification. Knowledge Acquisition 5(2), 199–220 (1993)
2. Kalfoglou, Y., Schorlemmer, M.: Ontology mapping: the state of the art. The Knowledge Engineering Review 18(1), 1–31 (2003)
3. Euzenat, J., Shvaiko, P.: Ontology Matching. Springer, Heidelberg (2007)
4. Arcos, J.L., Esteva, M., Noriega, P., Rodríguez-Aguilar, J.A., Sierra, C.: Engineering open evironments with electronic institutions. Engineering Applications of Artificial Intelligence 18(2), 191–204 (2005)
5. Atencia, M., Schorlemmer, M.: Formalising Interaction-Situated Semantic Alignment: the communication product. In: ISAIM 2008. Proceedings (2008)
6. Finin, T., Labrou, Y., Mayfield, J.: KQML as an agent communication language. In: Bradshaw, J. (ed.) Software Agents, AAAI/MIT Press (1997)
7. O'Brien, P.D., Nicol, R.C.: FIPA - towards a standard for software agents. BT Technology Journal 16(3), 51–59 (1998)
8. Atencia, M., Schorlemmer, M.: A formal model for situated semantic alignment. In: AAMAS 2007. Proceedings, pp. 1270–1277 (2007)
9. Besana, P., Robertson, D.: How service choreography statistics reduce the ontology mapping problem. In: Aberer, K., Choi, K.-S., Noy, N., Allemang, D., Lee, K.-I., Nixon, L., Golbeck, J., Mika, P., Maynard, D., Mizoguchi, R., Schreiber, G., Cudré-Mauroux, P. (eds.) ASWC 2007 and ISWC 2007. LNCS, vol. 4825, pp. 44–57. Springer, Heidelberg (2007)