

ANEMONE: An Effective Minimal Ontology Negotiation Environment

Jurriaan van Diggelen, Robbert-Jan Beun, Frank Dignum, Rogier M. van Eijk, John-Jules Meyer
Institute of Information and Computing Sciences
Utrecht University, the Netherlands
{jurriaan, rj, dignum, rogie, jj}@cs.uu.nl

ABSTRACT

Communication in open heterogeneous multi agent systems is hampered by lack of shared ontologies. To overcome these problems, we propose a layered communication protocol which incorporates techniques for ontology exchange. Using this protocol, the agents gradually build towards a semantically integrated system by establishing minimal and effective shared ontologies. We tested our approach, called ANEMONE, on a number of heterogeneous news agents. We show how these agents successfully exchange information on news articles, despite initial difficulties raised by heterogeneous ontologies.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*multiagent systems, coherence and coordination*;
D.2.12 [Software Engineering]: Interoperability—*data mapping*

General Terms

Languages, Theory

Keywords

Ontology negotiation, Semantic interoperability, Ontology alignment

1. INTRODUCTION

A fundamental communication problem in open multi agent systems is caused by the heterogeneity of the agent's knowledge sources, or more specifically of the underlying *ontologies*. Although ontologies are often advocated as a complete solution for knowledge sharing between agents, this is only true when all agents have knowledge about each others' ontology. The most straightforward way to establish this would be to develop one common ontology which is used by all agents. However, this scenario would be very

unlikely in open multi agent systems, as those on the internet, because it would require all involved system developers to reach consensus on which ontology to use. Moreover, a common ontology forces an agent to abandon its own world view and adopt one that is not specifically designed for its task [7]. This may result in a suboptimal situation.

In an open multi agent system, communication problems that arise from heterogeneous ontologies should be *solved*, rather than avoided. In this paper, we present a suitable solution, called ANEMONE: AN Effective Minimal Ontology Negotiation Environment. We illustrate our approach with a case study that involves some semantic integration problems that are typical for open systems. We consider an open community of agents that periodically download news articles on different topics from RSS news feeds¹. Different news providers are represented by different agents; the taxonomy of news topics supplied by a provider forms the agent's ontology. The fact that different news providers categorize their news differently gives rise to a proliferation of heterogeneous ontologies. The resulting semantic integration problems crop up once the agents start to exchange news articles with each other. For example, consider the agent Ag-M that represents the news provider *Moreover* [5], and the agent Ag-Y that represents the news provider *Yahoo* [3]. Driven by a user's request to Ag-M for more articles on the topic *Basketball*, Ag-M decides to ask Ag-Y for *Basketball* articles. Although Ag-Y has relevant articles on this topic, in Ag-Y's ontology they are classified under *NBA*². Because Ag-Y does not know *Basketball* and Ag-M does not know *NBA* their communication fails.

Traditional approaches, such as standardization and *Ontology alignment* [16], are not suited for these kinds of semantic integration problems. Standardization efforts in this domain are almost doomed to fail, because the news providers deliberately distinguish themselves from others by using different ontologies. Ontology alignment has been proposed as a technique that enables agents to keep their individual ontologies by making use of *mappings* between the different ontologies. Although this is a step in the right direction, it assumes that the mappings can be pre-defined before the agents start interacting. In our case, it is not known beforehand which ontology mappings are needed due to the openness of the system and the fact that ontologies change from time to time. For example, one of Yahoo's news topics, *Asia Tsunami Disaster*, is clearly a temporary topic.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

¹RSS is a popular XML format for the syndication of news content on the internet

²"NBA" is the National Basketball Association

In ANEMONE, the agents exchange ontological information on an as-needed basis. We hereby adhere to the emerging paradigm of ontology negotiation [6]. Agents first try to cope with the situation as is; when communication fails to be effective, the agents seek a minimal solution which solves their communication problem at hand. Drawing on related work on ontology negotiation and on our experiences with the news agent case, we adopt a standpoint on issues as what kind of shared ontology should be built up, and when and where this should occur. These standpoints are incorporated in ANEMONE and make it a novel approach for ontology negotiation. By implementing ANEMONE and applying it to the news agent case, we show that ANEMONE provides a solution for some semantic integration problems which, to the best of our knowledge, cannot be solved using existing techniques.

In the following section we compare our approach with related work. The third section provides a broad outline of our approach. The communication mechanisms of ANEMONE are discussed in section four. Section five describes implementation details and experimental results. We conclude and give directions for future research in section six.

2. RELATED WORK

Recently, researchers have turned their attention to more flexible solutions for semantic integration problems. An ontology agent [1] has been proposed to facilitate agent communication by registering ontologies and performing services such as translating between ontologies. This would allow agents to reconcile heterogeneous ontologies at interaction time, *on the condition* that every agent knows and trusts the ontology agent and that the ontology agent knows a mapping between every ontology in the system. In our case, such a centralized service is unsuitable as it is not known beforehand which inter-ontology mappings are needed.

Most recently, a few approaches appeared in the literature which tackled the problem in a fully *decentralized* way. W. Truszkowski and S. Bailin have coined the term *Ontology Negotiation* to refer to such approaches [6]. In their paper, the authors present a communication mechanism which enables agents to exchange parts of their ontology in a pattern of successive clarifications. The DOGGIE approach [20] makes similar assumptions, but focuses mainly on the machine learning aspects of ontology exchange, namely the problem of teaching the meaning of a concept to another agent. In this context, we also mention the approach by Soh and Chen [17], where agents exchange ontological knowledge when they believe it would improve operational efficiency. Ontology negotiation can be considered as the most ambitious approach for semantic integration [19]. Because this area is also relatively new, much work remains to be done.

The main differences between ANEMONE and the approaches mentioned above are:

- Contrary to [6], [20] and [17], ANEMONE is not restricted to equivalence mappings (mappings between two equivalent concepts in two ontologies), but deals with arbitrary mappings, such as subset mappings and disjointness mappings. For example, a mapping might state that one agent's concept *NBA* is a subset of another agent's concept *Basketball*. This has some repercussions in the overall communication mechanism regarding the prevention of overgeneralization.




Figure 1: Example ontologies

- Before the agents use resource consuming machine learning techniques, they try to exchange a meaning of a concept by stating the relations with other concepts. This is done by defining the unknown concept in terms of other concepts. The receiver decides whether this definition sufficiently explains the concept.
- Whereas the approach by Soh and Chen uses improved operational efficiency as a main criterium to start ontology exchange, we use a formal, and thus more precise, notion of lossless communication [10]. As lossless communication depends on the ontologies of both agents, both agents are involved in assessing this property.

3. APPROACH

3.1 Ontologies

An ontology consists of a set of concepts and a set of concept relations. We avoid naming conflicts by ensuring that two concepts in different ontologies with the same name also have the same meaning. This can be easily achieved by using namespaces. A concept relation is one of the following: $\sqsubset$ (strict subconcept relation), $\sqsupset$ (strict superconcept relation), $\equiv$ (equivalence), $\perp$ (disjointness), $\oplus$ (overlap). The symbol $\perp$ is also used to denote the bottom concept. Figure 1 presents a graphical representation of two example ontologies. An arrow between two concepts represents a subconcept relation (and against the flow a superconcept relation). Two concepts in two different branches in the ontology are disjoint. Concepts that are equivalent or overlap are connected with a line with the $\equiv$ or $\oplus$ symbol in it. For readability, we have left out concept relations that are derivable from other relations.

Besides concepts and their relations, some other aspects of ontologies are important in the context of ontology negotiation. Firstly, there is a difference between a *native* concept and an *acquired* concept. In Figure 1, the former is represented as non-shaded and the latter as shaded. An agent that has not yet exchanged any ontological knowledge with other agents, has only native concepts in its ontology. The native ontology of an agent is implemented by its system developer to enable it to store and reason with operational knowledge (knowledge that is used by the agent to carry out its task). Because the operational knowledge base is defined in terms of the native ontology, this ontology is constrained to be static. During ontology negotiation, the agents occasionally teach each other new concepts. These are the agent's acquired concepts and are defined additionally to its native concepts. Acquired concepts are not used for storing knowledge but only serve as vehicles for communication.

This allows the acquired ontology to be dynamic.

Another important issue in ontology negotiation is keeping track of which concepts are shared with other agents. In Figure 1, a concept that is shared between Ag-1 and Ag-2 is represented by a box around the concept. When we speak of a concept that is shared between two agents, we actually mean that the concept is common knowledge between the agents. In ontology negotiation, a concept may also be *unknowingly shared*. In this case, both agents know the concept, but do not know this of each other. Examples of unknowingly shared concepts are n and m in Figure 1. Such a situation might arise when two agents learn the same concept from a third agent.

The following terminology is useful for discussing the communication mechanisms of ANEMONE:

- γ_1 is a *particularization* w.r.t. γ_2 iff $\gamma_1 \sqsubset \gamma_2$ or $\gamma_1 \oplus \gamma_2$
- γ_1 is an *implied concept* w.r.t. γ_2 iff $\gamma_1 \sqsupset \gamma_2$ or $\gamma_1 \equiv \gamma_2$

3.2 Communication

ANEMONE does not make any initial demands on agents with regard to their amount of shared concepts. We will discuss an initial situation which is the most difficult, i.e. the agents do not share any native concepts besides the $\top$ concept. At first, these agents lack any means to cooperate. As they participate in conversations, they build up something which enables them to reach a desired level of coordination. An overview of the communication protocol that provides for this is presented in Figure 2. The upper layer of the pro-




Figure 2: Layered communication protocol

tol (NCP) deals with normal agent communication, i.e. the kind of social interaction which agents normally exhibit when no ontology problems exist in the system. To deal with ontology problems, two layers are added to the protocol: a protocol for exchanging concept definitions (CDP), and a protocol for teaching concepts to each other using machine learning techniques (CEP).

This modular structure of ANEMONE's communication mechanism enables us to make alterations in one layer of the protocol without affecting the other layers. For example, in our news agent domain we have implemented CEP using text classification techniques. In another domain, which deals with individuals other than text documents, other concept learning techniques may be used in CEP.

Our focus in this section is on the domain independent design choices of the protocol, for example, on how the agents traverse through the protocol layers and on general techniques that can be used in the individual layers. These design choices are motivated by a number of standpoints on issues as: *what* should be built up during ontology negotiation and *where* and *when* is this built up?

- *What: Minimal & Effective* Agents should build up a minimal and effective shared ontology. An ontology that is minimal in size is not bulky and slowly processable due to superfluous concepts. An ontology that is minimal in use enables agents to keep their conversations short. An effective ontology is expressive enough to convey sufficient information in a sound manner. Particularly, communication should be *lossless*, meaning that no information is lost in communication that could have been avoided by spending more effort on learning each others' ontologies. A formal definition of lossless communication is presented in [10].
- *When: Lazy* Agents should build up a shared ontology only when strictly necessary. Ontology exchange is not a goal in itself; it is a means of enabling communication. The agents leave NCP when communication fails to be effective and try to return as soon as possible. Defining a concept in CDP is less resource consuming than learning a concept in CEP. Therefore an agent explicates instead of defines a concept only when strictly necessary. In other words, the agents try to stay as high up as possible in the layered communication protocol.
- *Where: Decentralized* In line with the ontology negotiation paradigm, no central location exists where a shared ontology is built up. Every agent increments its own ontology with the necessary concepts. In this way, they collectively address the semantic integration problem.

4. COMMUNICATION MECHANISM

In the following three sections we discuss and motivate the communication mechanisms that underly NCP, CDP and CEP. An overview of the messages that are exchanged while following the protocol is presented in figure 6.

4.1 Normal Communication Protocol

Normal communication protocols have been extensively studied in the agent communication literature. The focus of this section is on the adjustments that are required to deal with the partially shared ontologies described in section 3.1, and the criteria upon which the agents base their decision to resort to CDP. This confronts us with the task of finding a proper way to deal with *message composition*, *message interpretation*, and the decision of when to *switch to CDP*. To discuss these issues, we assume a simple setting in which a speaker intends to convey one of its native concepts to a hearer. *Message composition* is the speaker's task of translating the native concept it intends to convey to a concept it will actually speak to the hearer. *Message interpretation* is the hearer's task of translating the concept in the message to one of its native concepts to store the information in its knowledge base. Either the speaker or the hearer may decide to *switch to CDP* when it believes that more concepts should be shared between them. We begin our discussion with a very simple way to implement message composition, interpretation, and a switch to CDP. As we evaluate this protocol against the quality criteria of minimality, effectivity, and laziness (discussed in section 3.2), we detect shortcomings and adjust the protocol accordingly. Subsequent evaluations and refinements eventually give rise

to a communication mechanism of considerable complexity which is described in section 4.1.5.

4.1.1 Equivalence mappings

Our first proposal for NCP is a simple communication mechanism that exploits some features of partially shared ontologies introduced in section 3.1, namely the distinction between native and acquired concepts, equivalence mappings, and the notion of shared concepts. These features can be incorporated in a relatively straightforward manner. In this proposal, message composition involves translating the native concept the speaker intends to convey to an *equivalent shared concept*, and to use this in a message. Message interpretation involves translating the concept used in the message to an *equivalent native concept*. The speaker switches to CDP when no equivalent shared concept is available in its ontology. Some examples w.r.t. Figure 1: Ag-1 intends to convey *d*, speaks *d* which Ag-2 interprets as *k*; Ag-1 intends to convey *e*, can not do so, and switches to CEP to teach the concept to Ag-2. Note that in this scenario, a shared concept enables communication in *both* directions. For example, when Ag-2 intends to convey *k*, it may speak *d* which Ag-1 interprets as *d*. The speech act that allows the speaker to translate what it intends to convey to an equivalent shared concept is called *ExactInform*. It is clear that this communication mechanism does not give rise to information loss, i.e. it enables lossless communication.

4.1.2 Beyond equivalence mappings

The previous proposal for NCP does not describe the handling of acquired concepts that are mapped in an agent's native ontology as sub- and superconcepts. For example, Ag-2's acquired concept *c* is mapped between its native concepts *j* and *i*. One measure we must take to overcome this shortcoming, is to make message interpretation less strict; from now on, message interpretation involves translating the concept used in the message to the *most specific implied native concept* (instead of to an equivalent native concept). For example, Ag-1 intends to convey *c*, exactinforms *c* which Ag-2 interprets as *i*. Note that although this version of *ExactInform* may introduce some information loss, no information is lost in the communication process that could have been prevented by sharing extra concepts. Therefore, *ExactInform* still guarantees lossless communication.

Besides this refinement to message interpretation, the message composition part in the protocol also requires revision. This is motivated in the following example. If message composition is defined as strictly as in the previous proposal, Ag-2 would never use concept *c* in its message because none of Ag-2's native concepts is equivalent to *c*. To comply with minimality, shared concepts should allow for communication in both directions. For example, Ag-2 should be allowed to communicate *j* using the shared concept *c*. For this reason, we loosen the conditions for message composition such that the speaker can translate the native concept it intends to convey to a shared implied concept. The speech act that is associated with this version of message composition is called *Inform*. For example: Ag-2 intends to convey *j*, informs *c* which Ag-1 interprets as *c*; Ag-1 intends to convey *e*, informs *d*, which Ag-2 interprets as *k*.

4.1.3 Towards effectivity

The previous proposal for NCP threatens the effectiveness of communication by allowing the agents to become overly general. In the extreme case, an agent that intends to convey one of its native concepts, speaks the shared concept $\top$. Apparently, *inform* does not guarantee lossless communication. We prevent overgeneralization by enabling the hearer to recognize when communication has been lossless and when not. When the hearer is not certain that communication was lossless, it requests specification. When the speaker cannot convey more specific information due to the lack of shared concepts, the agent switches to CDP. Stated in this way, the decision to switch to CDP has to do with recognizing overgeneralization.

The hearer recognizes overgeneralization by reasoning as follows. Upon hearing that an individual belongs to some concept, it knows that the individual is a member of every implied concept and that the individual is not a member of all concepts that are disjoint with the concept. This knowledge cannot be a symptom of overgeneralization. However, the hearer remains ignorant about the particularizations of the concept. This ignorance may be a symptom of overgeneralization. The rule for recognizing overgeneralization we propose here is: if there are native particularizations of the concept used in the message, communication may not have been lossless. For example: Ag-2 intends to convey *k*, informs *c*, Ag-1 requests specification (because of concept *d*), Ag-2 exactinforms *d* and the conversation finishes; Ag-1 intends to convey *e*, informs *d*, Ag-2 interprets this as *k*, recognizes lossless communication and the conversation finishes; Ag-2 intends to convey *l*, informs *f*, after which Ag-1 requests specification (because of *g*), after which Ag-2 switches to CEP to teach *l* to Ag-1.

4.1.4 Towards efficiency

The proposal for NCP described in the previous section eventually guarantees lossless communication. However, the communication mechanism is not minimal in use. For example, every time that Ag-2 informs Ag-1 about *j*, Ag-1 requests specification (due to the presence of *d*), upon which Ag-2 answers that it does not know whether *d* holds. Especially when larger ontologies are involved, every time the speaker intends to convey a general concept, a cumbersome dialogue follows where the hearer requests for specification many times. To make communication more efficient, we make the following adjustments to the *Inform* speech act. In message composition, the agent is obliged to translate to the *most specific* shared implied concept. This not only prevents the speaker from becoming more general than necessary, it also enables the hearer to form a belief about what the speaker intended to convey and what it did not. In philosophy of language, such a derivation is known as a *conversational implicature* [12]. In ANEMONE, it works as follows: upon receiving a concept, the hearer knows that the speaker did *not* intend to convey the shared subconcepts of that concept, otherwise it would have spoken differently. It therefore considers it useless to request these particularizations. For example, when Ag-2 speaks *c* (with the intention to convey *j*), Ag-1 does not request whether *d* holds, because it knows that Ag-2 could not have intended to convey *d*. But the hearer can form an even stronger belief than that. The hearer also knows that the speaker did *not* intend to convey any non-shared subconcepts of the shared subconcept of the spoken concept. For example, when Ag-2 informs

c (with the intention to convey j), Ag-1 does not request whether e holds, because it knows that Ag-2 could not have intended to convey e , otherwise it would have informed d . This communication mechanism has been proven to yield lossless communication [8].

4.1.5 Dealing with unknown concepts

Until now, we have discussed how agents should deal with shared concepts that are common knowledge. As argued before, ontology negotiation also gives rise to unknowingly shared concepts. A minimal ontology negotiation protocol should exploit this property. Therefore, in message composition, we also allow an agent to translate to a non-shared concept (which may turn out to be an unknowingly shared concept). The rule is that an agent may still translate what it intends to convey to the most specific shared implied concept, but if there are more specific non-shared implied concepts present, it may also choose one of these. Once an agent has used an unknowingly shared concept in the message, this concept becomes shared. Because an agent is not allowed to speak non-shared implied concepts that are more general than a shared implied concept, the efficiency measures that make use of conversational implicatures are still applicable. The choice to translate to a shared concept with the risk of being too general, or to translate to a more specific non-shared concept, with the risk of being not understood, is left to the agent. Communication strategies that resolve non-determinacy in the communication protocol are investigated in [9].

The final version of NCP is summarized in Figure 3. The next example shows that it may be beneficial to speak an unknowingly shared concept rather than a more general shared concept. Ag-2 intends to convey m , exactinforms the non-shared concept m which is understood by Ag-1 as h . Ag-2 could also have informed the shared concept $\top$ with the intention to convey m , which would have been answered by Ag-1 with a request for specificity. The next example shows that it may be beneficial to speak a knowingly shared concept, rather than a more specific non-shared concept. Suppose that Ag-1 intends to convey e , exactinforms e which does not get understood by Ag-2. If Ag-1 would have informed d instead, Ag-2 would have understood the message and would have recognized lossless communication.

4.2 Concept Definition Protocol

In the concept definition protocol, the speaker tries to convey the meaning of a concept by stating the relations with other concepts, i.e. it speaks a number of concept definitions. If these definitions enable the hearer to derive the *complete meaning* of the concept, the hearer switches back to NCP. An agent considers the meaning of an acquired concept complete, if it knows the relation with every other concept in its ontology. If there are not sufficient shared concepts available to convey the complete meaning, the agents switch to CEP. For example, Ag-2 intends to define k , speaks “ $k \equiv d$ ” after which Ag-1 knows the complete meaning of k and switches back to NCP. Suppose Ag-1 intends to define h and speaks “ $m \sqsubset h \sqsubset n$ ”. This definition enables Ag-2 to derive that h is disjoint with f, l, i, c, j, k, d , that h is subset of $\top$ and n , and that h is a superset $\perp$ and m . It therefore regards the meaning of h complete and switches back to NCP. Suppose Ag-1 intends to convey the meaning of g and speaks “ $\perp \sqsubset g \sqsubset f$ ”. This definition leaves Ag-2

Composition

- The agent translates what it intends to convey to:
 - the most specific shared implied concept, or
 - an implied concept that is more specific than the most specific shared implied concept

Interpretation

- The agent translates the concept in the message to:
 - the most specific implied native concept

Switch to CDP

- The hearer switches to CDP when the speaker cannot convey what it intends to convey in a message that gets understood by the hearer and which is not overgeneralized. The hearer considers an inform message which mentions a concept γ overgeneralized if native particularizations of γ exist that are not
 - a shared subconcept of γ , or
 - a subconcept of a shared subconcept of γ
-

Figure 3: Normal Communication Protocol

ignorant about the relation between g and l . It does not consider this meaning complete, and switches to CEP.

Composition

- The agent states the relations of what it intends to define with every shared concept.

Interpretation

- The agent adds the defined concept as an acquired concept to its ontology.
- The agent adds the relations in the message to its ontology, and derives relations that follow from this.

Switch to NCP

- When the hearer knows the relation of the defined concept with every concept in its ontology.

Switch to CEP

- When the speaker does not know any shared concepts to define concept relations.
 - When the hearer does not know the relation of the defined concept with every concept in its ontology.
-

Figure 4: Concept Definition Protocol

4.3 Concept Explication Protocol

The purpose of CEP is to convey the meaning of a concept when no satisfactory definition of the concept in terms of other concepts can be given. A commonly used technique for generating a concept mapping is to invoke a dictionary or a thesaurus, such as WordNet [11]. However, this turns out to have its limitations in our domain. Firstly, the technique is incomplete, e.g. Yahoo’s concept *NBA* is not defined in WordNet. Secondly, the technique may be incorrect. For example, the *International* concept of the American company Reuters is really a superset of the *UK* concept of the (British) BBC. From the perspective of the BBC agent, WordNet would suggest differently. For these reasons, we do not take the name of a concept into account while finding out its meaning.

Following Luc Steels [18], we assume that the meaning of a concept can be conveyed to another agent by pointing to instances. A basic requirement for this approach is a *classi-*

for each concept in an agent's ontology. The speaking agent, upon explicating a concept, communicates a number of positive and a number of negative examples of the concept. The hearer classifies these examples using the concept classifiers from its own ontology. We use the following terms that are common in machine learning [15]:

- TPR: True positive rate: the number of positively classified positive examples divided by the total number of positive examples.
- TNR: True negative rate: the number of negatively classified negative examples divided by the total number of negative examples.

For each concept, the agent calculates TPR and TNR and uses this information to assess the concept relation. In the ideal situation, when the agent's classifiers are perfect and every element in the domain of discourse is used as an example, TPR and TNR offer strict criteria for determining the correct concept relation. Consider an ideal situation, where TPR and TNR result from applying a classifier of concept γ_1 to examples of concept γ_2 . The relation between concept γ_1 and γ_2 can be determined as follows: $\gamma_1 \equiv \gamma_2$ iff $\text{TPR} = 1 \wedge \text{TNR} = 1$; $\gamma_1 \perp \gamma_2$ iff $\text{TPR} = 0$; $\gamma_1 \sqsubset \gamma_2$ iff $0 < \text{TPR} < 1 \wedge \text{TNR} = 1$; $\gamma_1 \sqsupset \gamma_2$ iff $\text{TPR} = 1 \wedge \text{TNR} < 1$; $\gamma_1 \oplus \gamma_2$ iff $0 < \text{TPR} < 1 \wedge \text{TNR} < 1$. In most practical situations, the agents have noisy classifiers and the domain of discourse may be infinite (as in our case study). We deal with this by loosening the conditions for determining concept relations. There are actually many ways to do this. We follow a simple approach by introducing some threshold parameters θ_1 , θ_2 and θ_3 . The way the concept relation is assessed is described in Figure 5.

Composition

- The agent sends positive and negative examples. The positive examples are individuals belonging to the concept to be explicated. The negative examples are individuals that belong to concepts that are disjoint with the concept to be explicated.

Interpretation

- For every native concept γ_1 in its ontology, the agent computes TPR and TNR using the examples in the message.
- The agent assesses the relation of the explicated concept γ_2 with every native concept γ_1 as follows:
 - $\gamma_1 \equiv \gamma_2$ if $\text{TPR} \geq \theta_1 \wedge \text{TNR} \geq \theta_2$
 - $\gamma_1 \perp \gamma_2$ if $\text{TPR} \leq \theta_3$
 - $\gamma_1 \sqsubset \gamma_2$ if $\text{TPR} < \theta_1 \wedge \text{TNR} \geq \theta_2$
 - $\gamma_1 \sqsupset \gamma_2$ if $\text{TPR} \geq \theta_1 \wedge \text{TNR} < \theta_2$
 - $\gamma_1 \oplus \gamma_2$ if none of the previous conditions holds

Switch to CDP

- After the hearer has processed the examples.

Figure 5: Concept Explication Protocol

5. RESULTS

5.1 Implementation

We have implemented ANEMONE as a multi-threaded Delphi application. A single news agent corresponds to an RSS




Figure 6: Message protocol of ANEMONE

reader. It periodically checks the internet for new news articles that belong to one of the concepts in its ontology. If so, it downloads these articles and uses an XML-parser to extract the title, summary, date, and URL of the article. It stores this information in a local database. The agents actively exchange information with each other about new news articles and their topics. NCP provides for this communication. A query for news articles of a topic, is interpreted as a request to the other agent for an inform speech act, after which NCP deals with the subsequent communication. The non-determinacy in message-composition is resolved using the following simple communication strategy. If the speaker intends to convey γ and the only translation of γ to a shared concept is $\top$, the speaker prefers to speak the unshared concept γ ; in all other cases, the speaker translates to the most specific shared implied concept of γ .

The concept explication protocol incorporates some application specific design choices. Because our case deals with text documents, the classifiers are implemented using text classification techniques. We use a support vector machine [14] to classify the *description* fields of the news articles of a topic. The texts are preprocessed using common techniques in natural language processing [13], such as *stemming* and *stopword removal*; subsequently, the texts are transformed in term frequency vectors. These vectors are used for training the agent's concept classifiers. When the agent has sufficiently trained its concept classifiers on its own news articles, it is capable of learning concepts from other agents by following CEP. In CEP, the teacher sends news articles belonging to the concept to be explicated as positive examples, and news articles belonging to the concepts that are disjoint with the concept to be explicated as negative examples. As described in the previous section, the hearing agent computes TPR and TNR values for each concept in its ontology and derives the concept relations. Because in our application, a small amount of misclassifications is acceptable, we used relatively tolerant criteria to assess concept relations, i.e. $\theta_1 = 0.75$, $\theta_2 = 0.75$, $\theta_3 = 0.4$.

5.2 Case study

We demonstrate the system using the four relatively simple agents depicted in Figure 7. The agents represent news publishers BBC [2], Moreover [5], Reuters [4] and Yahoo [3]. The agent's ontologies consist of subsets of the news feeds provided by the news publishers. After the agents have col-




Figure 7: Initial situation

lected news articles for a period of two months, their knowledge bases were filled with approximately 200 news articles per topic. This enabled them to train their classifiers and participate in ontology negotiation. The following examples present a series of successive dialogues which were generated by the agents. For demonstration purposes, the agents make some of their internal workings public; these are preceded by a *. Due to space limitations, we abbreviated long path names of news articles with [], and summarized some of the conversations. The agents begin in a situation where they do not share any concepts with each other (Figure 7).

Ag-Y intends to inform Ag-R about Science News.

```
1 Ag-Y:<Ag-R,<ExactInform,Science News,[]>>
2 Ag-R:<Ag-Y,<StartCDP>>
3 Ag-Y:<Ag-R,<ProposeStartCEP>>
4 Ag-R:<Ag-Y,<StartCEP>>
5 Ag-Y:<Ag-R,<Explicate,Science News,[],[]>>
* Ag-R Business News: TPR=0.15 TNR=0.86
* Ag-R Science: TPR=0.76 TNR=1
* Ag-R Sports: TPR=0 TNR=0.24
* Ag-R equivalent(Science News,Science)
6 Ag-R:<Ag-Y,<ExitCDP>>
7 Ag-Y:<Ag-R,<ExactInform,Science News,[]>>
* Ag-R translates to Science
```

Ag-R intends to inform Ag-Y about Science

```
8 Ag-R:<Ag-Y,<ExactInform,Science News,[]>>
```

Ag-Y intends to inform Ag-M about Science News

```
9 Ag-Y:<Ag-M,<ExactInform,Science News,[]>>
10 Ag-M:<Ag-Y,<StartCDP>>
11 Ag-Y:<Ag-M,<ProposeStartCEP>>
12 Ag-M:<Ag-Y,<StartCEP>>
13 Ag-Y:<Ag-M,<Explicate,Science News,[],[]>>
* Ag-M Basketball: TPR=0.29 TNR=0.32
* Ag-M Tennis: TPR=0.21 TNR=0.84
* Ag-M Boxing: TPR=0.26 TNR=1
* Ag-M disjoint(Science News,Basketball;Tennis;Boxing)
14 Ag-M:<Ag-Y,<ExitCDP>>
15 Ag-Y:<Ag-M,<ExactInform,Science News,[]>>
* Ag-M translates to Article
```

Ag-R intends to inform Ag-M about Science

```
16 Ag-R:<Ag-M,<ExactInform,Science News,[]>>
* Ag-R translates to article
```




Figure 8: Situation after 9 conversations

Ag-R intends to inform Ag-Y about Sports

```
17 Ag-R:<Ag-Y,<ExactInform,Sports,[]>>
18 Ag-Y:<Ag-R,<StartCDP>>
19 Ag-R:<Ag-Y,<Define Sports,
    <disjoint Science News>>>
20 Ag-Y:<Ag-R,<StartCEP>>
21-27 In CEP, Ag-Y learns that Sports is
    equivalent with Sports News
```

Ag-Y intends to inform Ag-R about NBA

```
28 Ag-Y:<Ag-R,<Inform,Sports,[]>>
29 Ag-R:<Ag-Y,<OK>>
```

Ag-R intends to inform Ag-Y about Business News

```
30 Ag-R:<Ag-Y,<ExactInform,Business News,[]>>
31 Ag-Y:<Ag-R,<StartCDP>>
32 Ag-R:<Ag-Y,<Define,Business News,
    <disjoint Sports,disjoint Science News>>>
33 Ag-Y:<Ag-R,<ExitCDP>>
34 Ag-R:<Ag-Y,<ExactInform,Business News,[]>>
* Ag-Y translates to Article
```

Ag-Y intends to inform Ag-M about Sports News

35-41 In CEP, Ag-M derives that Sports overlaps with basketball. Ag-M translates Sports to Article

Ag-Y intends to inform Ag-M about NBA

```
42 Ag-Y:<Ag-M,<Inform,Sports,[]>>
* Ag-M translates to Article
43 Ag-M:<Ag-Y,<ReqSpec>>
44 Ag-Y:<Ag-M,<ExactInform,NBA>>
45 Ag-M:<Ag-Y,<StartCDP>>
46-51 In CEP Ag-M learns that NBA is a subset of basketball
    and translates to basketball
```

After this conversation, the agents have built up knowledge about each others' ontologies, as shown in Figure 7. The agents have built up a shared ontology that enabled them to convey what they intended to convey (effectivity), but have not made the shared ontology larger than required (minimality). The acquired concepts are not only mapped to equivalent native concepts but also to native superconcepts (e.g. Ag-M's *NBA* and *Basketball*) and to disjoint native concepts (e.g. Ag-M's *Science News*). Because the approaches mentioned in section 2 only deal with equivalence mappings, these approaches would have failed to solve

the semantic integration problems of the news agents.

The conversation shows that many of the ontology negotiation techniques we have discussed in this paper are useful in a realistic domain. We mention the following: the use of mappings other than equivalence mappings (9-15, 42-51), the use of generalization in message composition (28-29), detecting information loss (42-43), the use of speaking in *unknowingly shared* concepts (16), using acquired concepts for communication in *both* directions (8,28), the use of communicating concept definitions (30-34), detecting inadequate concept definitions (19-20).

In this case, it took around 10 seconds (on a Pentium 4, 2.5 GHz) for the agents to finish concept explication. Although the CEP layer worked well enough for this relatively simple case, in a scenario with more complex ontologies or where the correctness of concept mappings is of critical importance, the CEP layer needs improvements. Most likely, these improved versions will be even more time consuming. This strengthens our motivation for lazy ontology negotiation to reduce the occurrence of concept explication to a minimum.

6. CONCLUSION

In this paper, we have presented the ANEMONE approach for solving semantic integration problems. Instead of trying to solve all ontology problems at one stretch at design time, ANEMONE provides agents with the tools to overcome ontology problems at agent interaction time. The layered communication mechanism tackles semantic integration problems when needed and only when needed.

We have applied ANEMONE to the domain of news articles. Despite the fact that agents used different ontologies to classify news articles, they were able to overcome their communication problems and successfully exchanged news articles with each other. The performance of ANEMONE can be improved by incorporating more sophisticated machine learning techniques in the lower layer of the protocol. For our application, where the agents use relatively simple ontologies, the concept explication layer worked sufficiently well. The focus in this paper, however, has been on the communication mechanisms which serve to embed the machine learning techniques in the overall communication architecture, i.e. the upper two layers in the protocol. Even in our relatively simple case, these communication mechanisms proved to be useful. The agents refrain from applying resource consuming machine learning techniques as much as possible. When necessary, they gradually contribute to a shared ontology which is minimal in size and enables them to convey sufficient information. There is no need for a central coordinating agent; the agents find out by themselves if their communication is not satisfactory and solve the problem between themselves. These features of ANEMONE established that the news article agents, after having participated in only a few conversations, became part of a reasonably semantically integrated system.

In the future, we plan to apply ANEMONE to more complex domains. This may require us to replace the current concept explication protocol with one that is better suited for that domain. Furthermore, the ontologies in that domain may be more complex than the concept hierarchies we have focussed on in this paper. This requires an extension of the other layers in the communication protocol to deal with, for example, the alignment of concept attributes.

7. REFERENCES

- [1] FIPA Ontology Service Specification. <http://www.fipa.org/specs/fipa00086/>.
- [2] <http://news.bbc.co.uk/1/hi/help/3223484.stm>.
- [3] <http://news.yahoo.com/rss>.
- [4] <http://today.reuters.com/rss/newsrss.aspx>.
- [5] http://w.moreover.com/site/other/categories_rss.html?
- [6] S. Bailin and W. Truszkowski. Ontology negotiation between intelligent information agents. *Knowledge Engineering Review*, 17(1):7–19, 2002.
- [7] T. Bylander and B. Chandrasekaran. Generic tasks for knowledge-based reasoning: the right level of abstraction for knowledge acquisition. *Int. J. Man-Mach. Stud.*, 26(2):231–243, 1987.
- [8] J. v. Diggelen, R. Beun, F. Dignum, R. v. Eijk, and J.-J. Meyer. Combining normal communication with ontology alignment. *Proceedings of the International Workshop on Agent Communication (AC'05)*.
- [9] J. v. Diggelen, R. Beun, F. Dignum, R. v. Eijk, and J.-J. Meyer. A decentralized approach for establishing a shared communication vocabulary. *Proceedings of the International Workshop on Agent Mediated Knowledge Management (AMKM'05) held with AAMAS 2005*.
- [10] J. v. Diggelen, R. Beun, F. Dignum, R. v. Eijk, and J.-J. Meyer. Optimal communication vocabularies and heterogeneous ontologies. In *Developments in Agent Communication*, LNAI 3396. Springer Verlag, 2004.
- [11] C. Fellbaum. *WordNet: An Electronic Lexical Database*. MIT Press, 1998.
- [12] H. P. Grice. Logic and conversation. *Cole, P., and J.L. Morgan, eds. Speech Acts*. New York: Academic Press, 4158, 1975.
- [13] P. Jackson and I. Moulinier. *Natural Language Processing for Online Applications: Text retrieval, extraction, and categorization*. John Benjamins Publishing, 2002.
- [14] T. Joachims. Making large-scale svm learning practical. *Advances in Kernel Methods - Support Vector Learning*, 1999.
- [15] R. Kohavi and F. Provost. *Machine Learning*, (30):271–274, 1998.
- [16] N. F. Noy and M. A. Musen. Prompt: Algorithm and tool for automated ontology merging and alignment. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2000.
- [17] L.-K. Soh and C. Chen. Balancing ontological and operational factors in refining multiagent neighborhoods. *Proceedings of the Fourth International Conference on Autonomous Agents and Multi-Agent Systems*, 2005.
- [18] L. Steels. Language as a complex adaptive system. In *Proceedings of PPSN VI*, Lecture Notes in Computer Science, Berlin, September 2000. Springer-Verlag.
- [19] M. Uschold and M. Gruninger. Creating semantically integrated communities on the world wide web. *Semantic Web Workshop Co-located with WWW 2002 Honolulu*, 2002.
- [20] A. Williams. Learning to share meaning in a multi-agent system. *Autonomous Agents and Multi-Agent Systems*, 8(2):165–193, 2004.