

Introducing Fuzzy Trust for Managing Belief Conflict over Semantic Web Data

Miklos Nagy¹, Maria Vargas-Vera², and Enrico Motta¹

¹ Knowledge Media Institute (KM_i)

The Open University

Walton Hall, Milton Keynes

MK7 6AA, United Kingdom

`mn2336@student.open.ac.uk, e.motta@open.ac.uk`

² Computing Department

The Open University

Walton Hall, Milton Keynes

MK7 6AA, United Kingdom

`m.vargas-vera@open.ac.uk`

Abstract. Interpreting Semantic Web Data by different human experts can end up in scenarios, where each expert comes up with different and conflicting ideas what a concept can mean and how they relate to other concepts. Software agents that operate on the Semantic Web have to deal with similar scenarios where the interpretation of Semantic Web data that describes the heterogeneous sources becomes contradicting. One such application area of the Semantic Web is ontology mapping where different similarities have to be combined into a more reliable and coherent view, which might easily become unreliable if the conflicting beliefs in similarities are not managed effectively between the different agents. In this paper we propose a solution for managing this conflict by introducing trust between the mapping agents based on the fuzzy voting model.

1 Introduction

Assessing the performance and quality of different ontology mapping algorithms, which operate in the Semantic Web environment has gradually been evolved during the recent years. One remarkable effort is the Ontology Alignment Evaluation Initiative ³, which provides a possibility to evaluate and compare the mapping quality of different systems. However it also points out the difficulty of evaluating ontologies with large number of concepts i.e. the library track where due to the size of the vocabulary only a sample evaluation is carried out by a number of domain experts. Once each expert has assessed the correctness of the sampled mappings their assessment is discussed and they produce a final assessment, which reflects their collective judgment. Our ontology mapping algorithm DSSim [1] tries to mimic the aforementioned process, using different software

³ <http://oaei.ontologymatching.org/>

agents as experts to evaluate and use beliefs over similarities of different concepts in the source ontologies. Our mapping agents use WordNet as background knowledge to create a conceptual context for the words that are extracted from the ontologies and employ different syntactic and semantic similarities to create their subjective beliefs over the correctness of the mapping. DSSim addresses the uncertain nature of the ontology mapping by considering different similarity measures as subjective probability for the correctness of the mapping. It employs the Dempster-Shafer theory of evidence in order to create and combine beliefs that has been produced by the different similarity algorithms. For the detailed description of the DSSim algorithm one can refer to [2]. Using belief combination has their advantages compared to other combination methods . However the belief combination has received a verifiable criticism from the research community. There is a problem with the belief combination if agents have conflicting beliefs over the solution. The main contribution of this paper is a novel trust management approach for resolving conflict between beliefs in similarities, which is the core component of the DSSim ontology mapping system.

The paper is organized as follows. Section 2 provides the description of the problem and its context. Section 3 describes the voting model and how it is applied for determining trust during the ontology mapping. In section 4 we present our experiments that have been carried out with the benchmarks of the Ontology Alignment Initiative. Section 5 gives an overview of the related work. Finally, section 6 describes our future work.

2 Problem description

In the context of the Semantic Web trust can have different meaning therefore before we describe the problem let us define the basic notions of our argument.

Definition 1 *Trust: One mapping agent's measurable belief in the competence of the other agents' belief over the established similarities.*

Definition 2 *Content related trust: Dynamic trust measure that is dependent on the actual vocabulary of the mappings, which has been extracted from the ontologies and can change from mapping to mapping.*

Definition 3 *Belief: The state in which a software agent holds a proposition or premise over a possible mapping of selected concept pair combination to be true. Numerical representation of belief can be assigned to a value between [0..1].*

If we assume that in the Semantic Web environment it is not possible to deduct an absolute truth from the available sources then we need to evaluate content dependent trust levels by each application that processes the information on the Semantic Web e.g. how a particular information coming from one source compares the same or similar information that is coming from other sources.

Dominantly the existing approaches that address the problem of the trust-worthiness of the available data on the Semantic Web are reputation based e.g.

using digital signatures that would state who the publisher of the ontology is. However another and probably most challenging aspect of trust appears when we process the available information on the Semantic Web and we discover contradictory information from the evidences. Consider an example from ontology mapping. When we assess similarity between two terms, ontology mapping can use different linguistic and semantic[3] information in order to determine the similarity level e.g. background knowledge or concept hierarchy. In practice any similarity algorithm will produce good and bad mappings for the same domain depending of the actual interpretation of the terms in the ontologies e.g. using different background knowledge descriptions or class hierarchy. In order to overcome this shortcoming the combination of different similarity measures are required. During the recent years a number of methods and strategies have been proposed[3] to combine these similarities. In practice considering the overall results these combination methods will perform well under different circumstances except when contradictory evidence occurs during the combination process.

In our ontology mapping framework different agents assess similarities and their beliefs on the similarities need to be combined into a more coherent result. However these individual beliefs in practice are often conflicting. A conflict between two beliefs in Dempster-Shafer theory can be interpreted qualitatively as one source strongly supports one hypothesis and the other strongly supports another hypothesis, where the two hypotheses are not compatible. In this scenario applying Dempster's combination rule to conflicting beliefs can lead to an almost impossible choice, because the combination rule strongly emphasizes the agreement between multiple sources and ignores all the conflicting evidences.

We argue that the problem of contradictions can only be handled from case to case by introducing trust for the similarity measures, which is applied only for the selected mapping and can change from mapping to mapping during the process depending on the available evidences. We propose evaluating trust in the different beliefs that does not depend on the credentials of the ontology owner but it purely represents the trust in a proposed subjective belief that has been established by using different similarity algorithms.

3 Fuzzy trust management for conflicting belief combination

In ontology mapping the conflicting results of the different beliefs in similarity can be resolved if the mapping algorithm can produce an agreed solution, even though the individual opinions about the available alternatives may vary. We propose a solution for reaching this agreement by evaluating fuzzy trust between established beliefs through voting, which is a general method of reconciling differences. Voting is a mechanism where the opinions from a set of votes are evaluated in order to select the alternatives that best represent the collective preferences. Unfortunately deriving binary trust like trustful or not trustful from the difference of belief functions is not so straightforward since the different voters express their opinion as subjective probability over the similarities. For a

particular mapping this always involves a certain degree of vagueness hence the threshold between the trust and distrust cannot be set definitely for all cases that can occur during the process. Additionally there is no clear transition between characterising a particular belief highly or less trustful.

Fuzzy model is based on the concept of linguistic or "fuzzy" variables. These variables correspond to linguistic objects or words, rather than numbers e.g. trust or belief conflict . The fuzzy variables themselves are adjectives that modify the variable (e.g. "high" trust, "small" trust). The membership function is a graphical representation of the magnitude of participation of each input. It associates a weighting with each of the inputs that are processed, define functional overlap between inputs, and ultimately determines an output response. The membership function can be defined differently and can take different shapes depending on the problem it has to represent. Typical membership functions are trapezoidal, triangle or exponential. The selection of our membership function is not arbitrary but can be derived directly from fact that our input the belief difference has to produce the trust level as an output. Each input has to produce output, which requires a trapezoidal and overlapping membership function. Therefore our argument is that the trust membership value, which is expressed by different voters, can be modelled properly by using fuzzy representation as depicted on Fig. 1.

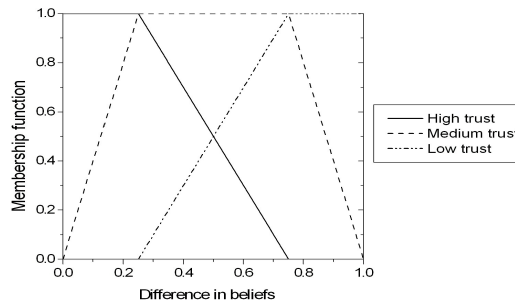


Fig. 1. Trust representation

Imagine the scenario where before each agent evaluates the trust in other agent's belief over the correctness of the mapping it calculates the difference between its own and the other agent's belief. The belief functions for each agent are derived from different similarity measures therefore the actual value might differ from agent to agent. Depending on the difference it can choose the available trust levels e.g. one agent's measurable belief over the similarity is 0.85 and an another agent's belief is 0.65 then the difference in beliefs is 0.2 which can lead to high and medium trust levels. We model these trust levels as fuzzy membership functions.

In fuzzy logic the membership function $\mu(x)$ is defined on the universe of discourse U and represents a particular input value as a member of the fuzzy set i.e. $\mu(x)$ is a curve that defines how each point in the U is mapped to a membership value (or degree of membership) between 0 and 1.

For representing trust in beliefs over similarities we have defined three overlapping trapezoidal membership functions, which represents high, medium and low trust in the beliefs over concept and property similarities in our ontology mapping system.

3.1 Fuzzy voting model

The fuzzy voting model was developed by Baldwin [4] and has been used in Fuzzy logic applications. However, to our knowledge it has not been introduced in the context of trust management on the Semantic Web. In this section, we will briefly introduce the fuzzy voting model theory using a simple example of 10 voters voting against or in favour of the trustfulness of an another agent's belief over the correctness of mapping. In our ontology mapping framework each mapping agent can request a number of voting agents to help assessing how trustful the other mapping agent's belief is.

According to Baldwin [4] a linguistic variable is a quintuple $(L, T(L), U, G, \mu)$ in which L is the name of the variable, $T(L)$ is the term set of labels or words (i.e. the linguistic values), U is a universe of discourse, G is a syntactic rule and μ is a semantic rule or membership function. We also assume for this work that G corresponds to a null syntactic rule so that $T(L)$ consists of a finite set of words. A formalization of the fuzzy voting model can be found in [5].

Consider the set of words $\{ \text{Low_trust } (L_t), \text{ Medium_trust } (M_t) \text{ and High_trust } (H_t) \}$ as labels of a linguistic variable trust with values in $U = [0, 1]$. Given a set "m" of voters where each voter is asked to provide the subset of words from the finite set $T(L)$, which are appropriate as labels for the value u . The membership value $\chi_{\mu(w)(u)}$ is taking the proportion of voters who include u in their set of labels which is represented by w .

We need to introduce more opinions to the system i.e. we need to add the opinion of the other agents in order to vote for the best possible outcome. Therefore we assume for the purpose of our example that we have 10 voters (agents). Formally, let us define

$$\begin{aligned} V &= A1, A2, A3, A4, A5, A6, A7, A8, A9, A10 \\ \Theta &= L_t, M_t, H_t \end{aligned} \tag{1}$$

The number of voters can differ however assuming 10 voters can ensure that

1. The overlap between the membership functions can proportionally be distributed on the possible scale of the belief difference [0..1]
2. The work load of the voters does not slow the mapping process down

Let us start illustrating the previous ideas with a small example - By definition consider our linguistic variable L as TRUST and $T(L)$ the set of linguistic values as $T(L) = (Low_trust, Medium_trust, High_trust)$. The universe of discourse is U , which is defined as $U = [0, 1]$. Then, we define the fuzzy sets $\mu(Low_trust)$, $\mu(Medium_trust)$ and $\mu(High_trust)$ for the voters where each voter has different overlapping trapezoidal membership functions as described on Table 1.

Table 1. Fuzzy set definitions

Voters	$\mu(Low_trust)$	$\mu(Medium_trust)$	$\mu(High_trust)$
A1	[0.25:0,0.75:1,1:1]	[0:0,0.25:1,0.75:1,1:0]	[0:1,0.25:1,0.75:0]
A2	[0.25:0,0.70:1,1:1]	[0:0,0.30:1,0.70:1,1:0]	[0:1,0.30:1,0.75:0]
A3	[0.25:0,0.65:1,1:1]	[0:0,0.35:1,0.65:1,1:0]	[0:1,0.35:1,0.75:0]
A4	[0.25:0,0.60:1,1:1]	[0:0,0.40:1,0.60:1,1:0]	[0:1,0.40:1,0.75:0]
A5	[0.25:0,0.55:1,1:1]	[0:0,0.45:1,0.55:1,1:0]	[0:1,0.45:1,0.75:0]
A6	[0.25:0,0.50:1,1:1]	[0:0,0.50:1,0.50:1,1:0]	[0:1,0.50:1,0.75:0]
A7	[0.30:0,0.75:1,1:1]	[0.5:0,0.50:1,0.50:1,0.95:0]	[0:1,0.25:1,0.70:0]
A8	[0.35:0,0.75:1,1:1]	[0.10:0,0.50:1,0.50:1,0.90:0]	[0:1,0.25:1,0.65:0]
A9	[0.40:0,0.75:1,1:1]	[0.15:0,0.50:1,0.50:1,0.85:0]	[0:1,0.25:1,0.60:0]
A10	[0.45:0,0.75:1,1:1]	[0.20:0,0.50:1,0.50:1,0.80:0]	[0:1,0.25:1,0.55:0]

The data in Table 1 are demonstrative only for the purpose of an example, which is presented in this paper. The difference in the membership functions represented by the different vertices of the trapezoid in Table 1 ensures that voters can introduce different opinions as they pick the possible trust levels for the same difference in belief.

The possible set of trust levels $L = TRUST$ is defined by the Table 2. Note that in the table we use a short notation L_t means Low_trust, M_t means Medium_trust and H_t means High_trust. Once the fuzzy sets (membership functions) have been defined the system is ready to assess the trust memberships for the input values. Based on the difference of beliefs in similarities the different voters will select the words they view as appropriate for the difference of belief. Assuming that the difference in beliefs(x) is 0.67(one agent's belief over similarities is 0.85 and an another agent's belief is 0.18) the voters will select the labels representing the trust level as described in Table 2. Note that each voter has its own membership

Table 2. Possible values for the voting

A1	A2	A3	A4	A5	A6	A7	A8	A9	A10
L_t	L_t	L_t	L_t	L_t	L_t	L_t	L_t	L_t	L_t
M_t	M_t	M_t	M_t	M_t	M_t				
H_t	H_t	H_t							

function where the level of overlap is different for each voter. As an example the belief difference 0.67 can represent high, medium and low trust level for the first voter(A1) and it can only represent low trust for the last voter(A10). Then we compute the membership value for each of the elements on set $T(L)$.

$$\chi_{\mu(Low_trust)}(u) = 1 \quad (2)$$

$$\chi_{\mu(Medium_trust)}(u) = 0.6 \quad (3)$$

$$\chi_{\mu(High_trust)}(u) = 0.3 \quad (4)$$

and

$$L = \frac{Low_trust}{1} + \frac{Medium_trust}{0.6} + \frac{High_trust}{0.3} \quad (5)$$

A value x (actual belief difference between two agents) is presented and voters randomly pick exactly one word from a finite set to label x as depicted in Table 3. The number of voters will ensure that a realistic overall response will prevail during the process.

Table 3. Voting

A1	A2	A3	A4	A5	A6	A7	A8	A9	A10
H _t	M _t	L _t	L _t	M _t	M _t	L _t	L _t	L _t	L _t

Taken as a function of x these probabilities form probability functions. They should therefore satisfy:

$$\sum_{w \in T(L)} Pr(L = w|x) = 1 \quad (6)$$

which gives a probability distribution on words:

$$\sum Pr(L = Low_trust|x) = 0.6 \quad (7)$$

$$\sum Pr(L = Medium_trust|x) = 0.3 \quad (8)$$

$$\sum Pr(L = High_trust|x) = 0.1 \quad (9)$$

As a result of voting we can conclude that given the difference in belief $x = 0.67$ the combination should not consider this belief in the similarity function since based on its difference compared to another beliefs it turns out to be a distrustful assessment. The before mentioned process is then repeated as many times as many different beliefs we have for the similarity i.e. as many as different similarity measures exist in the ontology mapping system.

3.2 Introducing trust into ontology mapping

The problem of trustworthiness in the context of ontology mapping can be represented in different ways. In general, trust issues on the Semantic Web are associated with the source of the information i.e. who said what and when and what credentials they had to say it. From this point of view the publisher of the ontology could greatly influence the outcome of the trust evaluation and the mapping process can prefer mappings that came from a more “trustful” source.

However we believe that in order to evaluate trust it is better to look into our processes that map these ontologies, because from the similarity point of view it is more important to see how the information in the ontologies are “conceived” by our algorithms than who have created them e.g. do our algorithms exploit all the available information in the ontologies or just part of it. The reason why we propose such trust evaluation is because ontologies of the Semantic Web usually represent a particular domain and support a specific need. Therefore even if two ontologies describe the same concepts and properties their relation to each other can differ depending on the conceptualisation of their creators, which is independent from the organisation where they belong. In our ontology mapping method we propose that the trust in the provided similarity measures, which is assessed between the ontology entities are associated to the actual understanding of the mapping entities, which differs from case to case e.g. a similarity measure can be trusted in one case but not trustful in an another case during the same process. Our mapping algorithm that incorporates trust management into the process is described by Algorithm 1.

<pre> Input: Similarity belief matrixes $S_{n \times m} = \{S_1, \dots, S_k\}$ Output: Mapping candidates 1 for $i=1$ to n do 2 BeliefVectors BeliefVectors \leftarrow GetBeliefVectors($S[i, 1 - m]$) ; 3 Concepts \leftarrow GetBestBeliefs(BeliefVectors BeliefVectors) ; 4 Scenario \leftarrow CreateScenario(Concepts) ; 5 for $j=1$ to $size(Concepts)$ do 6 Scenario \leftarrow AddEvidences (Concepts) ; 7 end 8 if Evidences are contradictory then 9 for $count=1$ to $numberOf(Experts)$ do 10 Voters \leftarrow CreateVoters(10) ; 11 TrustValues \leftarrow VoteTrustMembership(Evidences) ; 12 ProbabilityDistribution \leftarrow CalculateTrustProbability(TrustValues) ; 13 Evidences \leftarrow SelectTrustedEvidences(ProbabilityDistribution) ; 14 end 15 end 16 Scenario \leftarrow CombineBeliefs(Evidences) ; 17 MappingList \leftarrow GetMappings(Scenario) ; 18 end </pre>
--

Algorithm 1: Belief combination with trust

Our mapping algorithm receives the similarity matrixes(both syntactic and semantic) as an input and produces the possible mappings as an output. The similarity matrixes represent the assigned similarities between all concepts in ontology 1 and 2. Our mapping algorithm iterates through all concepts in ontology 1 and selects the best possible candidate terms from ontology 2 which is represented as a vector of best beliefs(step 2). Once we have selected the best

beliefs we get the terms that corresponds to these beliefs and create a mapping scenario. This scenario contains all possible mapping pairs between the selected term in ontology 1 and the possible terms from ontology 2(step 3 and 4). Once we have build our mapping scenario we start adding evidences from the similarity matrixes(step 6). These evidences might contradict because different similarity algorithms can assign different similarity measure for the same mapping candidates. In these evidences are contradictory we need to evaluate which measure i.e. mapping agent's belief we trust in this particular scenario(step 8-15). The trust evaluation(see details in section 3.1) is invoked which invalidates the evidences(agent beliefs) which cannot be trusted in this scenario. Once the conflict resolution routine is finished, the valid beliefs can be combined and the possible mapping candidates can be selected from the scenario.

The advantage of our proposed solution is that the evaluated trust is independent from the source ontologies themselves and can change depending on the available information in the context.

4 Empirical evaluation

The evaluation was measured with recall and precision, which are useful measures that have a fixed range and meaningful from the mapping point of view. Before we present our evaluation let us discuss what improvements one can expect considering the mapping precision or recall. Most people would expect that if the results can be doubled i.e. increased by 100% then this is a remarkable achievement. This might be the case for anything but ontology mapping. In reality researchers are trying to push the limits of the existing matching algorithms and anything between 10% and 30% is considered a good improvement. The objective is always to make improvement in preferably both in precision and recall

We have carried out experiments with the benchmark ontologies of the Ontology Alignment Evaluation Initiative(OAEI),⁴ which is an international initiative that has been set up for evaluating ontology matching algorithms. The experiments were carried out to assess how trust management influences results of our mapping algorithm. Our main objective was to evaluate the impact of establishing trust before combining beliefs in similarities between concepts and properties in the ontology. The OAEI benchmark contains tests, which were systematically generated starting from some reference ontology and discarding a number of information in order to evaluate how the algorithm behave when this information is lacking. The bibliographic reference ontology (different classifications of publications) contained 33 named classes, 24 object properties, 40 data properties. Further each generated ontology was aligned with the reference ontology. The benchmark tests were created and grouped by the following criteria:

- Group 1xx: simple tests such as comparing the reference ontology with itself, with another irrelevant ontology or the same ontology in its restriction to OWL-Lite

⁴ <http://oaei.ontologymatching.org/>

- Group 2xx: systematic tests that were obtained by discarding some features from some reference ontology e.g. name of entities replaced by random strings or synonyms
- Group 3xx: four real-life ontologies of bibliographic references that were found on the web e.g. BibTeX/MIT, BibTeX/UMBC

As a basic comparison we have modified our algorithm (without trust), which does not evaluate trust before conflicting belief combination just combine them using Dempster’s combination rule. The recall and precision graphs for the algorithm with trust and without trust over the whole benchmarks are depicted on Fig. 2. Experiments have proved that with establishing trust one can reach higher average precision and recall rate.




Fig. 2. Recall and Precision graphs

Figure 2 shows the improvement in recall and precision that we have achieved by applying our trust model for combining contradictory evidences. From the precision point of view the increased recall values have not impacted the results significantly, which is good because the objective is always the improvement of both recall and precision together. We have measured the average improvement for the whole benchmark test set that contains 51 ontologies. Based on the experiments the average recall has increased by 12% and the precision is by 16%. The relative high increase in precision compared to recall is attributed to the fact that in some cases the precision has been increased by 100% as a consequence of a small recall increase of 1%. This is perfectly normal because if the recall increases from 0 to 1% and the returned mappings are all correct (which is possible since the number of mappings are small) then the precision is increases from 0 to 100%. Further the increase in recall and precision greatly varies from test to test. Surprisingly the precision have decreased in some cases(5 out of 51). The maximum decrease in precision was 7% and maximum increase was 100%. The recalls have never decreased in any of the tests and the minimum increase was 0.02% whereas the maximum increase was 37%.

As mentioned in our scenario in our ontology mapping algorithm there are number of mapping agents that carry out similarity assessments hence create belief mass assignments for the evidence. Before the belief mass function is combined each mapping agent need to calculate dynamically a trust value, which describes how confident the particular mapping agent is about the other mapping agent's assessment. This dynamic trust assessment is based on the fuzzy voting model and depending on its own and other agents' belief mass function. In our ontology mapping framework we assess trust between the mapping agents' beliefs and determine which agent's belief cannot be trusted, rejecting the one, which is as the result of trust assessment become distrustful.

5 Related work

To date trust has not been investigated in the context of ontology mapping. Ongoing research has mainly been focusing on how trust can be modelled in the Semantic Web context [6] where the trust of user's belief in statements supplied by any other user can be represented and combined. Existing approaches for resolving belief conflict are based on either negotiation or the definition of different combination rules that consider the possibility of belief conflict. Negotiation based techniques are mainly proposed in the context of agent communication. For conflicting ontology alignment an argumentation based framework has been proposed [7], which can be applied for agent communication and web services where the agents are committed to a ontology and they try to negotiate with other agent over the meaning of their concepts. Considering multi-agent systems on the Web existing trust management approaches have successfully used fuzzy logic to represent trust between the agents from both individual[8] and community[9] perspective. However the main objective of these solutions is to create a reputation of an agent, which can be considered in future interactions. Considering the different variants [10] [11] of combination rules that considers conflicting belief a number of alternatives have been proposed. These methods are based on well founded theoretical base but they all modify the combination rule itself and such these solutions do not consider the process in which these combinations take place. We believe that the conflict needs to be treated before the combination occurs. Further our approach does not assume that any agent is committed to a particular ontology but our agents are considered as "experts" in assessing similarities of terms in different ontologies and they need to reach conclusion over conflicting beliefs in similarities.

6 Conclusion

In this paper we have shown how the fuzzy voting model can be used to evaluate trust, and determine which belief is contradictory with other beliefs before combining them into a more coherent state. We have proposed new levels of trust in the context of ontology mapping, which is a prerequisite for any systems that makes use of information available on the Semantic Web. Our system is

flexible because the membership functions for the voters can be changed dynamically in order to influence the outputs according to the different similarity measures that can be used in the mapping system. We have described initial experimental results with the benchmarks of the Ontology Alignment Initiative, which demonstrates the effectiveness of our approach through the improved recall and precision rates. There are many areas of ongoing work, with our primary focus being additional experimentation to investigate different kind of membership functions for the different voters and to consider the effect of the changing number of voters and the impact on precision and recall.

References

1. Nagy, M., Vargas-Vera, M., Motta, E.: Dssim - managing uncertainty on the semantic web. In: Proceedings of the 2nd International Workshop on Ontology Matching. (2007)
2. Nagy, M., Vargas-Vera, M., Motta, E.: Multi-agent ontology mapping with uncertainty on the semantic web. In: Proceedings of the 3rd IEEE International Conference on Intelligent Computer Communication and Processing. (2007)
3. Euzenat, J., Shvaiko, P.: Ontology matching. Springer-Verlag, Heidelberg (DE) (2007)
4. Baldwin, J.F. In: Mass assignment Fundamentals for computing with words. Volume 1566 of Selected and Invited Papers from the Workshop on Fuzzy Logic in Artificial Intelligence ,Lecture Notes In Computer Science. Springer-Verlag (1999) 22–44
5. Lawry, J.: A voting mechanism for fuzzy logic. *International Journal of Approximate Reasoning* **19** (1998) 315–333
6. Richardson, M., Agrawal, R., Domingos, P.: Trust management for the semantic web. In: Proceedings of the 2nd International Semantic Web Conference. (2003) 351–368
7. Laera, L., Blaoe, I., Tamma, V., Payne, T., Euzenat, J., Bench-Capon, T.: Argumentation over ontology correspondences in mas. In: AAMAS '07: Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems, New York, NY, USA, ACM (2007) 1–8
8. Griffiths, N.: A fuzzy approach to reasoning with trust, distrust and insufficient trust. In: Proceedings of the 10th International Workshop on Cooperative Information Agents. (2006) 360–374
9. Rehak, M., Pechoucek, M., Benda, P., Foltyn, L.: Trust in coalition environment: Fuzzy number approach. In: Proceedings of The 4th International Joint Conference on Autonomous Agents and Multi Agent Systems - Workshop Trust in Agent Societies. (2005) 119–131
10. Yamada, K.: A new combination of evidence based on compromise. *Fuzzy Sets Syst.* **159**(13) (2008) 1689–1708
11. Josang, A.: The consensus operator for combining beliefs. *Artificial Intelligence* **141**(1) (2002) 157–170