

# Chapter 26

## Ontology, Semantics and Reputation

Andrew Koster and Jeff Z. Pan

**Abstract** In this chapter we discuss the problem of communicating about trust and how semantic technologies can help. We briefly introduce these semantic technologies and then discuss two well-known ontologies of trust:  $\mathcal{L}_{Rep}$  and FORe. However, defining a shared language for trust ignores the personal and subjective aspect of trust, which are an important part of how it is used. We therefore discuss a number of filtering and alignment methods that allow for the processing of communicated trust evaluations without compromising the subjective aspect of trust.

### 26.1 Introduction

This section presents an overview of ontologies for reputation. Ontology is a term borrowed from philosophy that refers to the science of describing the kinds of entities in the world and how they are related. In computer science, ontology is, in general, a model of (some parts of) the world, which not only identifies important vocabulary (including classes and properties) but also specifies their meaning with a formal logic. An ontology of reputation is thus a description of the types and causes of reputation, as well as a description of all entities involved with reputation.

Ontologies are widely used to represent the *shared* understanding of a domain and, in the case of reputation, thus represent a shared meaning of reputation between individuals. This allows these individuals to freely exchange evaluations of other agents in the system, thereby propagating trust, warning against irreputable, and recommending reputable agents. In Section 26.4 we describe some ontologies for

---

Andrew Koster  
IIIA - CSIC, Campus UAB, 08193 Bellaterra, Catalonia, Spain, e-mail: andrew@iia.csic.es

Jeff Z. Pan  
University of Aberdeen, King's College, Aberdeen AB24 3UE, UK, e-mail: jeff.z.pan@abdn.ac.uk

reputation. While they enable the exchange of reputative evaluations, ontologies have difficulty providing a shared notion of trust. Trust is more of a subjective notion than reputation and, therefore, agents are less likely to agree on the meaning and causes of a trust evaluation. In Section 26.5 we discuss several methods that allow agents to communicate subjective notions of trust in a meaningful manner.

We start this section with a more general discussion of ontologies. We give a brief overview of OWL, the most widely used language for defining ontologies, in Section 26.2. Moreover, if different agents have different ontologies, then it is necessary to perform some form of agreement management among these ontologies. In this case, argumentation (cf. Section 26.3) may help work out a shared understanding and, at the same time, identify the disagreements.

## 26.2 Ontology and OWL

The most well known ontology language is the Web ontology language OWL, standardised by the World Wide Web Consortium. The more updated version of OWL is OWL 2.<sup>1</sup> The formal underpinning for OWL 2 is Description Logics [3]. OWL is considered as one of the key foundations of the Semantic Web [22].

OWL2 provides the constructors for building complex class and property descriptions from atomic ones. For example, ‘elephants with their ages greater than 20’ can be described by the following OWL class description:<sup>2</sup>

$$Elephant \sqcap \exists age. >_{20},$$

where *Elephant* is an atomic class, *age* is an atomic datatype property,  $>_{20}$  is a data range, and  $\sqcap, \exists$  are class constructors. Class and property descriptions can be used in axioms in an OWL ontology. For example, we can define the class *AdultElephant* with the following OWL axiom:

$$AdultElephant \equiv Elephant \sqcap \exists age. >_{20};$$

we can represent the constraint ‘Elephant are a kind of *Animal*’:

$$Elephant \sqsubseteq Animal;$$

we can also assert that the individual elephant *Ganesh* is an instance of the class description ‘Elephants who are older than 25 years old’:

$$Ganesh : (Elephant \sqcap \exists age. >_{25}).$$

Reasoning plays an important role in ontologies, as it could make implicit connections explicit [14] and help detect inconsistency and incoherency [12, 19, 10]. For

<sup>1</sup> <http://www.w3.org/TR/owl-overview/>

<sup>2</sup> To save space, we use DL syntax [3] rather than RDF/XML syntax.

example, in the above elephant ontology, we could infer that Ganesh is an AdultElephant.

OWL 2 provides two levels of ontology languages: the expressive and decidable language OWL 2 DL [16], and three tractable sub-languages OWL 2 EL [2], OWL 2 QL [5] and OWL 2 RL. Accordingly, OWL 2 provides three level of reasoning services:

- Sound and complete reasoning for OWL 2 DL: this allows modellers to have more expressive power for their ontologies but there is no guarantee for efficient reasoning services, due to the high computational complexity for OWL 2 DL. Available reasoners include, e.g., HermiT, Pellet, FaCT++ and RacerPro.
- Sound and complete reasoning for tractable languages (EL, QL and RL): this allows modellers to enjoy the efficient reasoning services but the available expressive power is limited. Available reasoners include, e.g., CEL, QuOnto, and TrOWL.
- Approximate reasoning services for OWL 2 DL (based on the tractable sub-languages): this allows the modellers to have more expressive power for their ontologies and enjoy the efficient reasoning services; however, theoretically the reasoning could be incomplete. A typical reasoner of this kind is TrOWL<sup>3</sup>, which implements, e.g., a faithful approximate reasoning approach [27] that has been shown to be complete for the classification service on e.g. the evaluation ontologies in the HermiT Benchmark<sup>4</sup>.

### 26.3 Ontology and Argumentation

Ontology reasoning services can be used to help manage agreements and disagreements among different ontologies from different domain experts. Before reaching agreements, argumentation support [4] is needed.

1. To detect disagreements between two expert ontologies, one could merge the two ontologies and check if the merge ontology is inconsistent or incoherent [12]. If so, disagreements exist.
2. To identify the agreed subsets, one could compute the maximally consistent (coherent) sub-ontologies of the inconsistent (incoherent) one [19].
3. To resolve the disagreement, one could debug the ontology [10] and remove the problematic parts from the inconsistent (incoherent) ontology.

In the above Steps 2 and 3, argumentation support is needed. An argument is a pair (S, c), where c is a claim and S is the support of the claim. In the case of ontology argumentation, c is an axiom in or an entailment of the given ontology, S is a justification [15] of c in the given ontology. For detailed discussions on ontol-

---

<sup>3</sup> <http://trowl.eu/>

<sup>4</sup> [http://hermit-reasoner.com/2009/JAIR\\_benchmarks/](http://hermit-reasoner.com/2009/JAIR_benchmarks/)

ogy argumentation, we refer the reader to the part in this book on ontologies and semantics.

## 26.4 Ontologies for Reputation

Insofar as we know two ontologies have been proposed for the communication of reputation. The first is FORe, described by Casare and Sichman [6] and the second is  $\mathcal{L}_{Rep}$ , introduced by Pinyol and Sabater [23]. The second ontology was briefly discussed in Section 25, but we discuss it in more detail below.

### 26.4.1 FORe

The Functional Ontology of Reputation (FORe) considers a dual definition of reputation. Reputation as a *social product* is the agreement of opinion about some target and reputation as a *social process* is the transmission of opinions through a social network to form such an agreement. FORe defines the concepts used to specify reputation in both its forms. To do this, they base their ontology on Valente's Functional Ontology of Law [30], because they claim that "the concepts of the legal world can be used to model the social world, through the extension of the concept of legal rule to social norm".

The reason the ontology is a *functional* ontology is because it models the products of reputation in terms of the function they have in the social process of reputation. The main classes of the ontology are *Reputative Knowledge*, *Responsibility Knowledge*, *Normative Knowledge* and *World Knowledge*.

*Reputative Knowledge* represents reputation in its understanding as a product, or evaluation. An instance of reputative knowledge models the specifics of an evaluation using a number of properties, the most important of which are the role of the agents involved (whether they are targets, first-hand evaluators, propagators of information or recipients of information) and the type of the reputative information (for instance, direct reputation or propagated reputation). For a full overview of the properties, we refer to [6].

*World Knowledge* represents the knowledge about the environment.

*Normative Knowledge* represents the social norms in this environment.

*Responsibility Knowledge* represents the knowledge an agent has regarding the responsibility the various agents have with regards to behaviour and norms.

Altogether the ontology allows for the modeling of the processing of reputation: the *World Knowledge* allows for the modeling of behaviour of agents and the *Normative Knowledge* contains the information of whether such behaviour is acceptable or not. Using *Responsibility Knowledge* an evaluator can decide that an agent is responsible for its behaviour and thus the evaluator's *Reputative Knowledge* regarding that agent

is affected. This reputation can be propagated, creating new instances of Reputative Knowledge for other agents in the environment.

Casare and Sichman give a short example of how this works by considering how a trust evaluation is formed, step by step, from world knowledge: the agent observes someone smoking in a closed space. This world knowledge is combined with the normative knowledge that it is forbidden to smoke in closed spaces to identify a norm violation. Further responsibility knowledge is needed to know whether the smoker is responsible for this norm violation, and the agent ascribes responsibility to the smoker for his actions. This is evaluated into a reputative evaluation and combined with other sources of knowledge about the same person. This results in an evaluation that can be communicated to other agents in the system.

FORe thus provides an ontological description of how reputation models work: it allows for the communication of the input (using the world knowledge and normative knowledge) of a reputation model and the output (an instance of reputative knowledge). However, it does not detail what happens if two agents, using the same input, obtain different output. In this case, communication might be problematic. An attempt to deal with this is given by Nardin et al. [20], who present an ontology alignment service to promote interoperability between reputation models. As a proof-of-concept they show how FORe can be used as a shared ontology and how the concepts from two different trust models can be translated in and out of FORe. However, they encountered concepts in both trust models that could not be properly translated into FORe. Furthermore, they do not present a method for automatically mapping a trust model into FORe and an agent designer must provide such a mapping manually. This limits the applicability of FORe for representing and communicating reputation.

### 26.4.2 $\mathcal{L}_{Rep}$

An entirely different approach is taken by Pinyol et al., who propose the  $\mathcal{L}_{Rep}$  language for communicating about reputation [24]. Section 25.4 describes how this language is used in argumentation about trust, but it could be used as a shared language for describing trust without argumentation as well. This language is based on a comprehensive ontology for discussing concepts of trust and reputation. The ontology defines a *social evaluation* with three compulsory elements: a *target*, a *context* and a *value*. The *context* is specified using a second language  $\mathcal{L}_{Context}$ , which is a first-order dynamic language [13] for describing the domain. The target is the agent under evaluation and the value describes the quantification of the social evaluation. We will not go into details of this quantification, but the original description of the  $\mathcal{L}_{Rep}$  language gives different alternatives for the representation of this quantification, encompassing most, if not all, representations used in modern computational trust and reputation models [24].

The taxonomy of social evaluations is given in Figure 26.1. Here we see how social evaluations are split into the different types of evaluations related to trust and

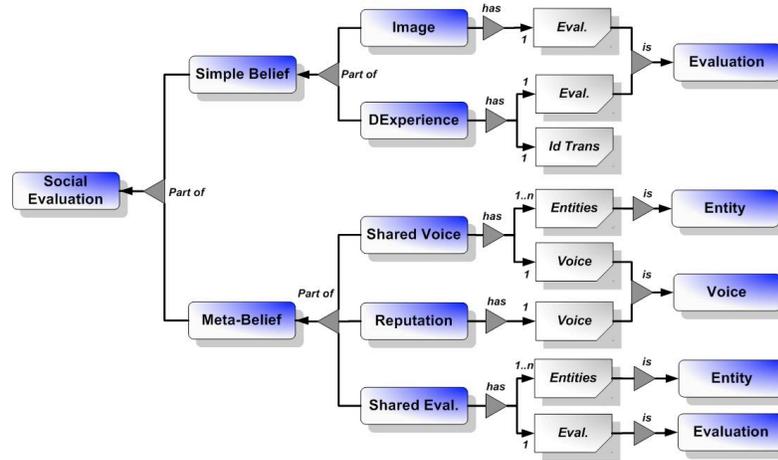


Fig. 26.1: Taxonomy of social evaluations in the  $\mathcal{L}_{Rep}$  ontology for talking about trust. Copied from [25].

reputation. This taxonomy is based on a sociological model of trust and reputation [7], which splits trust into a direct component, *image*, and a generalized concept of what the society thinks, *reputation*. These, in turn, are aggregations of direct experiences, shared voices and shared evaluations. In this way the ontology allows for the discussion of not just the final trust evaluation, but also the intermediate evaluations that are used in its calculation. The  $\mathcal{L}_{Rep}$  language is a first-order language with the vocabulary from the ontology described above and operators  $\wedge$ ,  $\neg$  and  $\rightarrow$ . A special subset of sentences in  $\mathcal{L}_{Rep}$  are ground atoms with either the predicate symbol *DExperience* or *Comm*. These are the basic elements in the ontology that are evaluations of experiences the agent has had, or communications it has received.

Pinyol et al. thus do not aim to model reputation as a process, but only the reputation (and trust) evaluations. These evaluations are modeled in detail and the language allows agents to communicate them. They acknowledge, however, that there may be subjective differences between the way different agents compute and interpret the evaluations and, for effective communication Pinyol proposes an argumentation framework for obtaining all the information about an evaluation in  $\mathcal{L}_{Rep}$  and allowing the agent to decide whether or not a communicated evaluation is acceptable [25].

## 26.5 Subjectivity of reputation

A large problem with using an ontology for trust and reputation is that trust is an inherently subjective concept. Despite the fact that reputation seems to escape this problem, because it is, per definition, “what is said about an individual by a social entity”, the question then remains of how the individual subjective trust evaluations of the members of such a social entity are combined to form this consensual reputation. Despite sharing a language, for discussing, comparing, or aggregating individuals’ subjective trust evaluations, the agents need to discuss their underlying preferences, or at the very least, an extensive description of the context in which a trust evaluation was formed.

Staab and Engel discuss this problem [28]. Rather than presenting a new trust model, they review the various problems that computational trust models encounter. One of these is the exchange of witness-information, or communicating trust evaluations directly. Their approach divides a computational trust model into three different stages and they discuss the communication at each of these. In a trust model, the first step is to monitor a target’s behaviour. This results in observations, which can be analyzed and, as the second step in their approach, interpreted as positive or negative experiences. The third step is to perform some computation, which is effectively an aggregation of the experiences with a target, resulting in a trust evaluation. At each of these steps communication can take place to augment the next process. Agents can communicate about observations, about experiences or about trust evaluations, although at each subsequent stage the level of subjectivity increases.

### 26.5.1 *Communication at the level of observations*

Communicating at the level of observations has the advantage of being a direct communication about the agent’s observations of the environment and thus being the nearest to an objective description of an interaction. Şensoy et al. use this approach to communicate witness information and incorporate it into the trust model directly [8]. They provide a trust model for evaluating web services and to communicate, they propose a dual ontology for communicating agents’ interactions with service providers. The first ontology is the base ontology, which captures the fundamental concepts that all web services have in common and a second domain ontology for describing the particular service the agents are discussing. When communicating, an agent can communicate all its observations of an interaction using these ontologies and the receiving agent can then evaluate these as if it had the interaction itself. However, there are two problems with this approach. The first is the assumption that an agent’s observations of an interaction can be objectively described. The ontology that Şensoy et al. propose includes properties like “quality” and “isAsDescribed” in the domain ontology, which they assume can be assessed in an objective manner by each agent. This seems to contradict their own assertion that any rating is always subjective, however it also begs the question of whether any domain exists in which

an interaction, that can serve to support a trust evaluation, can be satisfactorily described using only objective facts. The second problem with such an approach is more straightforward: for a receiving agent to be able to interpret an interaction as if it has observed it itself, the entire interaction must be faithfully recorded and communicated. An agent, in order to be a useful information source, must thus record details of interactions that it itself might consider trivial and never use. Furthermore, the shared ontology must include all properties all agents consider relevant for evaluating and, on top of that, it ignores possible privacy issues. While agents may be willing to communicate whether an interaction succeeded or failed, they may not be willing to communicate the exact details of that interaction. Especially if that contains sensitive information, such as, for instance, financial data.

A solution could be to allow agents to communicate partial observations of an interaction, however this leads to problems for the receiver. It has to use partial observations to evaluate the trust of a target. If the receiver needs to perform some kind of processing of the information in any case, it might be possible to use some of the more subjective information from the higher stages of trust models.

### *26.5.2 Communicating experiences and evaluations*

The reason a shared ontology for trust is infeasible is because of what Euzenat calls pragmatic heterogeneity [11]. He discusses three different levels at which ontological heterogeneity might appear. The first level is that of syntactic heterogeneity, which is quite straightforward: two agents use a different syntax to describe the same domain. An example in the real world would be a Chinese and an English speaking agent trying to communicate. At the second, or semantic, level the problem of syntax is presumed solved, but two agents use a different semantic for the same words. For instance, two agents who are discussing a minivan. One categorizes a minivan as a small bus, while the other categorizes it as a large car, so the meaning they assign to the word minivan is slightly different. This is the level at which most research into ontology alignment falls. The last level is that of pragmatics. At this level two agents agree on the syntax and the conceptual meaning of the word, however there is heterogeneity in how the word is used: this is almost always the problem when two agents try to talk about subjective concepts, such as “taste”, but also trust.

However, it is only recently that heterogeneity of trust has been considered as a problem of pragmatic heterogeneity. One of the contributing factors has been an attempt to solve the problem of heterogeneous trust concepts through techniques of ontology alignment [20]. As briefly mentioned in Section 26.4.1, Nardin et al. recognized the problem of heterogeneous trust evaluations and that ontologies did not properly capture different models' concepts. A particular problem for attempts to use ontology alignment for trust is that most trust models focus on how outcomes, that Staab and Engel call experiences, are aggregated, possibly together with other information, in order to compute the trust evaluations. The trust models therefore

already incorporate a large part of the subjectivity of trust at their lowest level. It is unclear how these outcomes can be mapped into a shared ontology.

An entirely different approach is that taken by Teacy et al. in their TRAVOS model [29] and Şensoy et al. with POYRAZ [9]. These models deal principally with liars, but their method may work equally well with trust evaluations from witnesses with too different a viewpoint. Both models learn how to filter out communicated evaluations from witnesses they mark as liars: the learning algorithm learns to distinguish liars by analyzing past experiences and consistently finding a difference between the witness' communicated evaluation and the actual evaluation after interacting. In other words, the algorithms use past experiences, together with witnesses' recommendations, to classify witnesses as either liars, or truthful agents. The main difference between TRAVOS and POYRAZ is that POYRAZ takes the context into account, in the form of an ontology for describing interactions. In fact, they use the same ontology as in [8], but rather than just communicating about the interaction, they include trust evaluations based upon the interaction. The reason this method works for detecting more than liars, is because, given a specific context, the method calculates the difference between a received evaluation and an evaluation based on personal experience. The latter is thus based on the agent's own subjective criteria, while the former is based on the witness' criteria. If these are too dissimilar too often for a single witness, this witness is considered a liar. The advantage is that this allows an agent to filter out information from agents that are too dissimilar to itself. One of the disadvantages is that it marks such agents as liars. This is problematic, because there is often a negative action attached to discovering a lying agent, such as the notification of the community that the agent is a liar. In the case of miscommunication based on subjectivity, this may lead to many agents incorrectly being marked as liars, with all its repercussions. Even if this is not the case, the filtering methods have another disadvantage: if there are many different possible criteria for calculating a trust evaluation, algorithms that learn to filter out evaluations, may very well filter out too much information for them to be viable.

An alternative is what Koster et al. [17] call trust alignment. This provides a translation, similar to the one proposed by Nardin et al., but taking the domain level information into account and, rather than attempting to translate to a central, shared ontology, attempt to learn an individualized translation between two agents' trust models.

### 26.5.3 Trust Alignment

There are a number of methods that can be considered trust alignment mechanisms. The first is described by Abdul-Rahman and Hailes' trust model [1]. This work describes a trust model that evaluates a trustee with an integer between 1 and 4, where 1 stands for *very untrustworthy* and 4 for *very trustworthy*. The alignment process uses the recommendations from another agent about *known* trustees to calculate four separate biases: one for each possible trust value. First the alignment method

calculates the own trust evaluations of the corresponding trustee for each incoming recommendation. The *semantic distance* between the own and other's trust is simply the numerical difference between the values of the trust evaluations. The semantic distances are then grouped by the value of the corresponding received trust value, resulting in four separate groups. Finally the bias for each group is calculated by taking the mode of the semantic distances in the group, resulting in four integers between -3 and 3, which can be used when the agent receives recommendations about unknown trustees. Simply subtract the corresponding bias from the incoming trust evaluation to translate the message.

This is a very simple approach to translating another agent's trust evaluation: it simply learns a vector of numerical biases and uses this, but, as shown in [17], this actually works remarkably well. However, methods that take the context into account work better. One of these is BLADE [26]. This model uses a conjunction of propositions to represent the context and a Bayesian network to learn the relation between the own trust evaluation and the other's trust evaluation given a certain context. While this works well, their representation of the context is very limited. For instance, the ontology of [8] requires a more expressive language, as does any other OWL, or even OWL Lite ontology [21]. In order to learn a context-based translation in a more expressive language, Koster et al. [18] propose to use Inductive Logic Programming. This approach learns a conjunction of Horn clauses that generalizes from a set of examples, with each example constituting the own trust evaluation, the other's trust evaluation and a description of the interaction in a first-order language. The algorithm performs regression, or if the trust evaluations are not numeric, then classification can be used, to find a translation. An approach like this, taking the context into account, is able to obtain more accurate estimates of a target's trustworthiness, as is shown in [17]. Furthermore, these methods are able to deal with a limited amount of lying, by substituting inconsistent trust evaluations with descriptions of the context and thus learn the context in which the witness' trust evaluations are inconsistent, or, if the context is specific enough, even learn a translation of a message regardless of whether the trust evaluation is a lie or not.

All these alignment methods attempt to deal with the problem of the pragmatic heterogeneity of trust by learning a specific alignment between two agents' based not on a conceptual representation of trust, but based on how an agent calculates and uses the trust evaluations. The latter two approaches do this by, additionally, taking the context into account and recognize that a trust evaluation of a target may change significantly in different contexts and thus any translation must do this too. This resolves two of the issues we discussed earlier: by using a machine learning approach, they do not require manual mappings of agents' models into a shared ontology and they do not filter out information, but instead translate it. This translation comes with an additional reliability measure, so, if the reliability is low, an agent may still make the choice to filter it out. Furthermore, it resolves some of the issues with communicating only at the lowest level that Staab and Engel [28] identify, because by using the subjective evaluations and learning a translation, the agent does not need to know all the specifics of the underlying interaction.

These advantages come at a cost. Because both BLADE [26] and Koster et al.'s approach [18] use quite complex machine learning algorithms, they need a large set of training data to learn an accurate translation. The training data consists of interactions that are shared between the requesting agent and the witness supplying information. The alignment is thus quite intensive, both from a communication and computation perspective. It also requires a domain in which it is likely that two agents can have a large number of similar interactions. However, if the conditions are met, these algorithms solve many of the issues with communicating trust and their application seems promising in a number of domains, including P2P routing, eCommerce and grid computing, although they have so far not been tested in such, realistic, application domains.

## 26.6 Conclusions

Communication with other agents is an important source of information for finding trustworthy interaction partners, however this communication is not straightforward. In this section we discuss a number of ways in which such communication can be established. The first is through a shared ontology. If the application provides an ontology for communicating trust, such as the  $\mathcal{L}_{Rep}$  or FORe ontologies that we discuss in Section 26.4, then the communication should not be problematic. The problem, however, is that such a fixed definition of trust does not allow agents to use trust as a personal and subjective evaluation: their use of trust is fixed by the shared ontology. Trust Alignment provides a solution for this, by allowing agents to learn a translation based on some shared set of evidence for each agent's trust evaluation. This allows each agent to communicate its own personal trust evaluations, which are translated by the receiving agent. A disadvantage of such methods is that a large number of shared interactions are required to learn this alignment. Another approach to communicating about trust can be found in Chapter 25 of this same book.

## References

1. Abdul-Rahman, A., Hailes, S.: Supporting trust in virtual communities. Proceedings of the 33rd Hawaii International Conference on System Sciences **6**, 4–7 (2000)
2. Baader, F., Brandt, S., Lutz, C.: Pushing the  $\mathcal{EL}$  Envelope. In: Proc. of IJCAI-05, pp. 364–369 (2005)
3. Baader, F., Lutz, C., Sturm, H., Wolter, F.: Fusions of Description Logics and Abstract Description Systems. Journal of Artificial Intelligence Research (JAIR) **16**, 1–58 (2002)
4. Black, E., Hunter, A., Pan, J.Z.: An Argument-based Approach to Using Multiple Ontologies. In: Proc. of the 3rd International Conference on Scalable Uncertainty Management (SUM 2009) (2009)
5. Calvanese, D., Giacomo, G.D., Lembo, D., Lenzerini, M., Rosati, R.: DL-Lite: Tractable description logics for ontologies. In: Proc. of AAAI 2005 (2005)
6. Casare, S., Sichman, J.: Towards a functional ontology of reputation. In: AAMAS '05: Proc. of the fourth international joint conference on Autonomous Agents and Multiagent Systems, pp.

- 505–511. ACM, Utrecht, The Netherlands (2005). DOI <http://doi.acm.org/10.1145/1082473.1082550>
7. Conte, R., Paolucci, M.: Reputation in Artificial Societies: Social beliefs for social order. Kluwer Academic Publishers (2002)
  8. Şensoy, M., Yolum, P.: Ontology-based service representation and selection. *IEEE Transactions on Knowledge and Data Engineering* **19**(8), 1102–1115 (2007)
  9. Şensoy, M., Zhang, J., Yolum, P., Cohen, R.: Context-aware service selection under deception. *Computational Intelligence* **25**(4), 335–366 (2009)
  10. Du, J., Qi, G., J. Z. Pan, Y.S.: A Decomposition-Based Approach to OWL DL Ontology Diagnosis. In: Proc. of the 23rd IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2011) (2011)
  11. Euzenat, J., Shvaiko, P.: Ontology matching. Springer-Verlag, Heidelberg (DE) (2007)
  12. Flouris, G., Huang, Z., Pan, J.Z., Plexousakis, D., Wache, H.: Inconsistencies, Negations and Changes in Ontologies. In: Proc. of AAAI2006, pp. 1295–1300 (2006)
  13. Harel, D.: First-order dynamic logic. Lecture notes in computer science. Springer (1979). URL <http://books.google.es/books?id=mp4pAQAAIAAJ>
  14. Hogan, A., Pan, J.Z., Polleres, A., Decker, S.: SAOR: Template Rule Optimisations for Distributed Reasoning over 1 Billion Linked Data Triples. In: Proc. of the 9th International Semantic Web Conference (ISWC2010) (2010)
  15. Horridge, M., Parsia, B., Sattler, U.: Justification oriented proofs in owl. In: P.F. Patel-Schneider, Y. Pan, P. Hitzler, P. Mika, L.Z. 0007, J.Z. Pan, I. Horrocks, B. Glimm (eds.) International Semantic Web Conference (1), Lecture Notes in Computer Science, pp. 354–369. Springer (2010)
  16. Horrocks, I., Kutz, O., Sattler, U.: The Even More Irresistible SROIQ. In: Proc. of KR 2006 (2006)
  17. Koster, A., Sabater-Mir, J., Schorlemmer, M.: Trust alignment: a sine qua non of open multi-agent systems. In: R. Meersman, T. Dillon, P. Herrero (eds.) Proceedings of On the Move to Meaningful Internet Systems (OTM 2011, Part I), LNCS, vol. 7044, pp. 182–191. Springer, Hersonissos, Greece (2011)
  18. Koster, A., Schorlemmer, M., Sabater-Mir, J.: Engineering trust alignment: Theory, method and experimentation. *Journal of Human-Computer Studies* **70**(6), 450–473 (2012). DOI 10.1016/j.jhcs.2012.02.007
  19. Meyer, T., K. Lee, R.B., Pan, J.Z.: Maximally Satisfiable Terminologies for the Description Logic ALC. In: Proc. of AAAI2006, pp. 269–274 (2006)
  20. Nardin, L.G., Brandão, A.A.F., Muller, G., Sichman, J.S.: A service-oriented architecture to support agent reputation models interoperability. In: Proceedings of 3rd Workshop on ontologies and their applications (WONTOiçæ2008). Salvador Bahia, Brazil (2008)
  21. Owl web ontology language overview. <http://www.w3.org/TR/owl-features>, November 12, 2009
  22. Pan, J.Z.: Description logics: Reasoning support for the semantic web. Ph.D. thesis, School of Computer Science, The University of Manchester (2004)
  23. Pinyol, I., Sabater-Mir, J.: Arguing about reputation. the lrep language. In: A. Artikis, G. O’Hare, K. Stathis, G. Vouros (eds.) Engineering Societies in the Agents World VIII: 8th International Workshop, ESAW 2007, LNAI, vol. 4995, pp. 284–299. Springer Verlag (2007)
  24. Pinyol, I., Sabater-Mir, J., Cuni, G.: How to talk about reputation using a common ontology: From definition to implementation. In: Proc. of Tenth Workshop "Trust in Agent Societies" at AAMAS ’07, pp. 90–102. Honolulu, Hawaii, USA (2007)
  25. Pinyol Catadau, I.: Milking the reputation cow: Argumentation, reasoning and cognitive agents. Ph.D. thesis, Universitat Politècnica de Catalunya (2010)
  26. Regan, K., Poupart, P., Cohen, R.: Bayesian reputation modeling in e-marketplaces sensitive to subjectivity, deception and change. In: Proceedings of the 21st National Conference on Artificial Intelligence (AAAI), pp. 1206–1212. AAAI Press, Boston, MA, USA (2006)
  27. Ren, Y., Pan, J.Z., Zhao, Y.: Soundness Preserving Approximation for TBox Reasoning. In: Proc. of the 25th AAAI Conference Conference (AAAI-2010) (2010)

28. Staab, E., Engel, T.: Combining cognitive with computational trust reasoning. In: R. Falcone, K. Barber, J. Sabater-Mir, M.P. Singh (eds.) TRUST 2008, *LNAI*, vol. 5396, pp. 99–111. Springer (2008)
29. Teacy, W.T.L., Patel, J., Jennings, N.R., Luck, M.: Travos: Trust and reputation in the context of inaccurate information sources. *Journal of Autonomous Agents and Multi-Agent Systems* **12**(2), 183–198 (2006)
30. Valente, A.: *Legal Knowledge Engineering – A Modeling Approach*. IOS Press, Amsterdam (1995)