Why does trust need aligning?

Andrew Koster

IIIA, Artificial Intelligence Research Institute
CSIC, Spanish National Research Council
Bellaterra, Spain
andrew@iiia.csic.es

Abstract. In this position paper we explain why the alignment of trust for computational agents is a problem which requires closer consideration than it has previously been given. We give a review of related work from various fields of research and propose a general framework in which a solution for the alignment of trust should be found.

1 Introduction

de gustibus non est disputandum
— Latin maxim

One of the main problems in open multi-agent systems is how the heterogeneous agents can interact with one another. Usually an ontology is given by the system designers for the agents to communicate. If there is no single ontology, but each agent uses its own, there is a large amount of work done to enable the alignment of these ontologies [1]. An ontology can be used to describe the environment and negotiate in objective terms. However, if the objects to be discussed are subject to semiotic heterogeneity [2], ontological solutions do not suffice. Semiotic, or pragmatic, heterogeneity is the problem encountered when an ontology can be interpreted in different manners. The agents agree on the syntax of the ontology and the content the ontology describes, yet still do not coincide on the *meaning* of the concepts. This is the problem we encounter when talking about taste... and also when talking about trust. We do not even have a non-ambiguous definition of what trust itself means. There are various philosophical, sociological, cognitive and economic theories of what trust means to humans [3–7] and different computational models based on these [8, 9]. Agents within an open MAS using such diverse models of trust and reputation will therefore run into communication problems if they wish to exchange information about these concepts. In [10] a strong case is made for why agents need to communicate their trust information and the manner of reputation formation is solely based on communication. As an example we can look at eBay's system [11]: their website does not include any tools for modeling trust, just an interface for communicating evaluations. The other users use this information to form their own evaluations. This simple system is one of the pillars for the website's tremendous success [12]. One of the reasons it works well is because it allows users to write a comment: not only do they give a score to the trade partner, but a short explanation *why* they give that score.

When we state opinions, we almost always give the reason for having this opinion. The same holds true for opinions about trust. Consider the phrase "I wouldn't trust Alice". If uttered sincerely, this is uttered in accordance with Grice's maxim of quality [13]: the speaker believes it to be true and the speaker can justify it based on some evidence. Yet in reality we are not so easily satisfied and we have the urge to ask for a justification from the speaker. The recognition that a justification is required, is already present in toddlers [14] and when given, such as in an utterance "I wouldn't trust Alice, because the car she sold me was a lemon", it is far more likely to be accepted by a listener. Of course, we are assuming that a "lemon" is the same to both the speaker and the listener. If not, this is an equally unsatisfactory utterance and a further explanation of what the speaker considers to be a lemon is required. We therefore continue asking for justifications of opinions until we are satisfied we understand what the speaker *means* (in [14] this is referred to as the justifications being grounded) and there is no further semiotic heterogeneity. Whether or not the listener agrees with the speaker is another issue entirely. It may very well be the case, that the listener is perfectly happy to trust Alice in her role as a car saleswoman, despite the speaker's utterance. It might even be the case that the listener disagrees with the reasoning of the speaker: Alice selling him a lemon is not a good reason to mistrust her. Perhaps even to the contrary: the listener rather dislikes the speaker and Alice selling him a lemon is all the more reason to esteem Alice. However, whatever complex reasoning lies behind the conversation partners' opinions, for them to communicate effectively they need to justify their opinions. There is evidence that this necessity of justifying a statement appears when we first begin to realize that the communication partner has mental states which may differ from our own [15]. This same difference in mental states should be assumed in autonomous computational agents [16] acting in an open MAS. All agents participate in the MAS to fulfill their own goals, based on their own beliefs, therefore we should expect communication of opinions, including opinions of trust, to be accompanied by some justification.

In this paper we discuss the problem further as well as giving a brief description of a proposed method for aligning trust. In the next section we discuss methods related to alignment of trust, such as argumentation, ontology alignment and dealing with uncertainty. In Section 3 we further expand on the problem and propose an area in which the solution could be found. We recap the main points and discuss them at the end of the paper.

2 Related work and similar problems

As mentioned in the previous section, trust alignment falls within the general area of semiotic heterogeneity, which is recognized as a sub-area of general semantic heterogeneity. While most work done on semantic heterogeneity focuses on the more traditional forms of ontology alignment, there is some work done on the problem. The field of semiotics [17] focuses on how humans interpret signs and how meanings are formed. Some work has been done on applying semiotics to AI, most notably semiotic dynamics [18] studies how semiotic systems come into being, both in human societies and in artificial societies. However, so far its application to alignment problems has been slim. Most work concentrates on evolving a semiotic system together, thus reaching a consensus of meaning. However, this is not possible with all concepts. While allowing agents to evolve a language to talk about external objects is possible, when they talk about their own mental states we specifically do not want a consensus. Each agent has their own way of interpreting the world and it is this interpretation they must try to communicate.

Communicating trust evaluations is exactly such an issue and some work has been done in trust modeling for computational systems to address it. Most trust and reputation models take both the agent's experiences and other agents' communications into account, however Abdul-Rahman & Hailes' model [19] recognizes differences in mental states to a certain extent: the interpretation of communicated trust is based on previous interactions with the same sender. This allows their model to use a heuristic to "bias" received messages depending on how far apart received trust evaluations from the same sender in the past have been from the agent's own trust evaluations. We agree that this setup is a good approach to the problem. By learning quantification of the dissimilarity between the agent's own trust evaluations and the communicated trust evaluations received from the other agent, a form of an alignment is made. However, their model misses out on some important points:

- Their alignment system only works with other agents using the same trust model. However, in an open MAS this cannot be assumed. This might be solvable using regular ontology alignment before the "biasing" of the message. Various trust ontologies have been proposed [20, 21] as well as ontology mapping service for them [22].
- More serious is the fact that they do not take the context of the trust evaluation into account. The reason trust is subjective is because each agent has its own goals and observations of the environment in which it evaluates trust. This may cause the bias to vary between trust evaluations. Their model simply averages the bias into one general numerical bias. This is a simplification we feel cannot be made.

Dealing with the context in which a trust evaluation is made complicates matters, because the bias has to be made conditional upon this context, which needs to be discovered in its own right. The approach we propose in Section 3 is inspired both by this model and the insight that context plays an important role.

Insofar as we know, this is currently the only model attempting any form of alignment of incoming trust information. There are models based on cognitive principles [23, 24] which offer the capability to do similar things, however the main focus of these models is on the cognitive modeling of trust and they leave the alignment of incoming messages as an open problem.

Another approach is to not use other agents' trust evaluations at all and instead communicate just the information about interactions. This is done in [25]. While this avoids any subjective terms in the communication, it does not work well if the information is asymmetric between the agents or there are privacy issues in communicating. For communication of just domain information to be effective, the agents have to communicate everything about an interaction, allowing the other agent to compute its own trust evaluation. A separate alignment process allows agents more freedom in communication: in this process agents only talk about interactions they have both observed. The agents both have different information available to them, which allows them to compute their own trust evaluations. The agents can then choose which properties of the interactions to communicate, in accordance with any restrictions they have. If similar situations consistently lead to similar communicated properties, then even if the other agent does not have knowledge of the interaction, it can approximate its trust evaluation through the similarity of previous situations.

2.1 Is argumentation about trust the same as aligning?

Some work has been done focusing on ways to build argumentations for trust [26, 27]. The reasons for argumentation are the same as the ones given above: agents need to explain their trust evaluations to each other [26] and to the user [27]. We agree that argumentation is an excellent domain to find a solution to the problem at hand: it gives a formal framework for building explanations. However, there are two important issues which are not answered by the work in argumentation so far. [26] describes an argumentation language in which agents can form justifications for their trust and communicate these to each other. Such a justification consists of a trust evaluation to be justified and a phrase in their justification language. However, their justification language consists of further predicates about trust, as well as agents' evaluations of interactions. This allows agents to build justifications for their ungrounded terms on further ungrounded terms. Somehow these terms need to be justified in grounded terms. As we saw in the introduction: the concept of a car being a "lemon" may be equally subjective as the trust based on that. Agents will need to justify *why* they evaluated the car as a lemon. This process should be repeated until the terms of the conversation are only terms in a shared, objective language describing the domain. The other problem is that there is no clear description of what agents should *do* with these justifications when they receive them. [26] says these can be incorporated into the trust model, however there is no description of how this should happen. Such a method of incorporation into the trust model, together with communication in objective justifications about trust is an alignment. [27] offers a different view: the justifications are used specifically to communicate to the user why a trust evaluation is given. The justifications are therefore output of the trust model to the user and the language can be grounded in the user's own terms. Furthermore they use an "opponent modeler", which learns to distinguish different situations in which recommendations are to be trusted or not. This opponent modeler could very well be seen as a method for learning an alignment, however it only takes the agent's own past experiences into account. This means it will need a large set of interactions with the same "opponent" agent to learn an accurate model, as their experiments seem to corroborate. Due to such interactions being prior to their modeling, or alignment, the agent runs the risk of such interactions being harmful, either due to malintent or simple miscommunication. Rather than only using the interactions the two agents have had with each other, the agents could learn an alignment, based on all interactions both agents have information of, thus reducing the risk from many interactions with an unknown opponent. This alignment could be formed in a separate communication process. Another facet of the approach in [27], is that they seem to focus on detecting when an opponent is dishonest. While an important facet, it is not the only situation in which an alignment method would be useful, as we will describe in the next section.

2.2 Is taking uncertainty into account sufficient?

There has been quite a lot of work done on discovering dishonesty in communicating trust evaluations. The main point of such research is to find a way of detecting *when* a communicated trust evaluation is inaccurate. In [28] these methods are divided into endogenous and exogenous methods. The endogenous methods discover unfair evaluations through statistical analysis of all ratings. This presumes that the meaning of trust is the same to all agents. Statistical methods can detect which agents diverge from the norm, but in an open MAS this doesn't automatically imply these agents are frauds. They may have different trust models. Additionally, if enough different trust models are in use, the significance of these methods drops considerably. These methods are designed to work in environments where few agents "lie", which is taken to mean their opinion deviates from the average. Thus if many different models are used, this assumption will not hold. Another assumption underlying these models is that the communication acts about trust are either public, or passed through a central unit where such statistical measures can be computed. These methods are therefore not very well adapted to use in an open MAS. The exogenous methods are more diverse and are defined by their use of additional information to determine unfair evaluations. TRAVOS [29] and BRS [30] for instance predict the reliability of a trade partner by calculating the expected value, given a probability distribution which is tailored to fit past experiences with that partner. TRAVOS in specific takes the context of the past experiences used in this calculation into account, thus discerning between similar and dissimilar situations to assess the reliability more accurately. POYRAZ [31] was developed as a combination of endogenous and exogenous methods and expands on TRAVOS' method, by taking not only the own experiences, but combining this with publicly available information, such as reputation. [31] shows experiments in environments with liars, in which POYRAZ and TRAVOS show a significant improvement over similar models which do not take contextual information into account. This confirms our earlier assertion that it is important to distinguish the context in which a trust evaluation is communicated.

However, there is an important issue which is not considered in the models discussed above: the question of *why* the information is unreliable. In the theory and in the experiments all these models make the assumption that the reason a communicated evaluation is unreliable is because it is either incompatible with the own model, or the agent is lying. We argue that these are two very different cases. In the first situation the agent's evaluation is different because it is based on a different trust model. In the second it is because the agent has malicious intentions. Models for dealing with unreliable information, however, deal with both situations in the same manner: the information is discarded. This should not be necessary in the first case, if only the agents can align their notions, such communications can be translated and used as reliable information. Furthermore, because the models don't distinguish between the two situations this may have repercussions for the truthful, but badly aligned agents, if it is assumed they are lying: when this information is propagated it may influence their reputation. Thus the statistical analysis is very necessary to discover *when* it might be useful to align, but it doesn't replace alignment. That would discard useful information, as well as negatively impacting the information-giver's reputation.

3 How to align trust?

One recurring theme, both in the theoretical approaches and in the related work we have discussed is a clear division between the subjective trust evaluations and the objective context information on which they are based. In [27] the argumentation is based on how the opponent is modeled, using the experiences the agent has had with that opponent. Similar experiences are used in TRAVOS and POYRAZ to discover unreliability in the trust evaluations. This is unsurprising, because these experiences play a central role in trust. A trust model evaluates agents based on such experiences. It can take only its own experiences into account, or also experiences the agent has observed. Furthermore such observations may be communicated. To be able to communicate about trust evaluations, it is therefore essential to also allow communication of these experiences. [31] gives an example ontology allowing for this, however each domain may have its own unique ontology to describe the interactions. In general a MAS will provide such an ontology to agents participating. An alternative is that agents have their own personal domain ontologies. These can be aligned using general ontology alignment methods [1]. Due to these being grounded ontologies about the environment we do not run into the problem of aligning subjective opinions. Once a shared domain ontology has been established the agents can exchange trust evaluations and information about the interactions such evaluations are based on. The receiving agent can use this information to form the alignment. We will describe the requirements for these parts in more detail, but first give a brief overview of the mathematical framework, giving more rigour to this idea.

3.1 Theoretical Foundations

Channel Theory [32] has been proposed as a general framework for semantic alignment [33]. This theory is a qualitative theory modeling the flow of information in distributed systems. [34] shows how dynamic situated ontology alignment can be considered in this framework. While this is a very different problem from the one we are considering, the article shows how a channel theoretic framework can aid, not only from a theoretic point of view, but also in considering how an alignment is formed. We can describe a channel in which information about trust can be transferred from one agent to another. This framework is described in detail in [35] and we give a short summary here. The intuition is that both agents can relate each others' subjective trust evaluations to the objective descriptions of interactions. By doing so they are able to find the underlying meaning of the trust evaluations.

Interaction-based alignment As we have argued, agents' interactions in the environment form the basic building blocks for trust. Such interactions are observed by different agents and each agent has an internal representation of this interaction. We make no further assumptions about such representations. As argued in [31], each agent may focus on different aspects of the interaction. Additionally agents may not receive the same information about an interaction. We further suppose that each agent may have its own way of representing such information.

These observations then lead to trust evaluations of the various agents involved. Any trust model can therefore be described as a binary relation between an agent's observations and its trust evaluations. These trust evaluations can be represented in some language $\mathcal{L}_{Trust}$, which we assume can be represented by all agents in the system. The meaning of phrases in $\mathcal{L}_{Trust}$ to the different agents is what the alignment process should uncover.

We consider trust alignment as a case-by-case problem. There is no need to align with agents there is no communication with. Furthermore, we could use a statistical analysis such as the ones used in POYRAZ or TRAVOS to filter the cases in which alignment is useful. Only in those cases should agents align. To do this, the agents need to discuss the interactions they both have observed and we assume there is a shared language to discuss these interactions. We call this language $\mathcal{L}_{Domain}$ and emphasize that it is a shared language: both the syntax and the semantics are known by all agents in the system, as opposed to the semantics of $\mathcal{L}_{Trust}$, which is interpreted differently by the agents. Each agent can relate its own internal representation of interactions to phrases in $\mathcal{L}_{Domain}$. Because this is a language about the objective, grounded, properties of the interaction, not all observations of an interaction can be communicated, however it allows us to define exactly what it means for two agents to share an interaction. A set of interactions I is shared by agents A and B if there is some $\varphi \in \mathcal{L}_{Domain}$ such that φ is in both A and B's sets of observations of the set I, or, in other words, φ is the information shared between the agents about I. The information in φ could range from a detailed description of the interaction to only the very basic fact that the interaction took place. This depends on what both agents know about the interaction and what they are willing to share as well as what can be represented in $\mathcal{L}_{Domain}$.

For aligning, the agents should *only* use the trust evaluations based on interactions *both* agents observed. For this shared *core* of interactions it is known that while each agent may have different observations, these come from the same interactions. By communicating only about the trust evaluations these interactions support, the agents guarantee that they are sharing the "building blocks" of those trust evaluations and the *channel* [32] of information flow is established.

Forming the alignment Note that such trust interactions are not necessarily the same trust evaluations that either agent actually uses: when functioning normally, most trust models use *all* information, and thus interactions, available to them. This results in "believed" trust evaluations. For alignment purposes, however, there is no reason to limit the agents to just these "believed" trust evaluations. To expedite the alignment process it is useful to consider all trust evaluations that *could be* supported by the shared interactions, rather than only those which, in fact, are.

The basis of a trust alignment is a set of messages sent from one agent to another in the form $\langle \beta, \psi \rangle$, with: β a trust evaluation of a specific target in $\mathcal{L}_{Trust}$ and ψ pinpointing the specific shared interactions this evaluation is based on in $\mathcal{L}_{Domain}$. We now see the framework for the alignment process arise, because if one agent B sends such a message, the receiving agent A can compute its own trust evaluation α, based on observations of the same interactions supporting ψ. A Specific Rule for Alignment is thus made and is the tuple $\langle \alpha, \beta, \psi \rangle$.

These SRAs must now be generalized to a predictive model, such that, for example, agent A can know what trust evaluation α' it should associate with a certain $\beta' \in \mathcal{L}_{Trust}$, given a description ψ' of some interactions, which it has not observed. It must be able to associate a trust evaluation with only the communicated information β' and ψ'.

3.2 Finding predictive rules

Channel theory gives us a theoretically sound manner to define the building blocks of an alignment, but the actual process of finding a useful set of rules, which will allow future communications to be translated is not captured by channel theory. In fact, we propose there are various methods of alignment possible and choosing which one to use is dependent on both the domain and the agents aligning. We will discuss some of the options here.

An inductive approach We start with an approach using inductive learning to find a generalization. There are various different inductive algorithms, but they all use θ-subsumption at some point or another [36]. If we consider our SRAs in a slightly different format: $\alpha \leftarrow \beta, \psi$ and restrict the predicates used to first order logic, it is easy to see how these rules constitute a logic program. By using ILP [36] we can find generalizations of these, as described in [37]. The induced rules form a different logic program, which translates the other agent's trust evaluations, given a context ψ. A similar approach to modeling other agents is taken in [27], where a fuzzy rule induction algorithm is used in their "opponent modeler". While the focus is different, because the resulting rules are used for argumentation, the basic principle is the same. However, there is another difference between the approaches: the former uses induced rules to directly translate trust evaluations whereas the latter uses such rules as an approximation of the opponent's trust model. It would require a further step to translate trust evaluations from the induced model into similar ones for the own model. This approach of modeling and then translating was taken in [38].

An ILP-based approach has one major downside, which is that ILP algorithms are very dependent on the language bias. If we do not have enough information to structure the search space quite rigorously the algorithm will either not find the correct generalization or take prohibitively long to search for it. By giving enough background information the search space can be made far more accessible, however this task needs to be done manually. Depending on the domain, the associated ontology and the agents aligning this may or may not be viable. A method of automatically generating the language bias would be through analysis of the communicated messages only and not taking the entire ontology into account. This requires the agents to only communicate the relevant properties of the interactions in $\mathcal{L}_{Domain}$ and a method for dynamically constructing the language bias for the ILP algorithm. Insofar as we know, no work has been done towards such an approach.

A context-discovery approach Instead of considering the SRAs as a logic program, we could see them as ψ being an example of the context in which α and β are aligned. Therefore finding generally predictive alignment rules equates to the problem of finding a good classification of the contexts. Some work has been done in automatic context recognition for trust [39, 40] and as mentioned in Section 2.2, there are various models to detect unreliable information based on the context. An intuitive continuation of such approaches would be to apply context recognition not only to trust evaluations but to the alignment of communication about the same. One issue with clustering- or classification-based approaches is to find an appropriate distance measure. In approaches so far, propositional or attribute-value logics have been used to describe the contexts. If the language used is more complicated, the distance measure becomes harder to define, or calculate. An example of a distance measure which could work for first order logics is [41], however in general this is an open issue.

An argumentation-based approach Another way of considering the SRA is as an argument. ψ is a justification for the trust evaluations α and β in both agents respectively. While this justification is given in a domain level language it is easy to see how this could be extended. The argumentation language proposed in [26] allows agents to communicate their justifications for trusting an agent, however it misses the link to such a grounded language. By extending this language the agents could each present their justifications for their trust evaluations. Such argumentations could either be used to negotiate about trust, or as cases in a Case-Based-Reasoning [42] algorithm. This could be used as a predictive method, by retrieving a comparable justification in the past and the corresponding trust evaluations.

4 Discussion

We have argued that trust alignment is a real problem if multiple trust models are to be effectively used in an open MAS. So far this issue has superficially been addressed by various techniques to deal with argumentation of trust or discovering unreliable communications about trust. While giving an indication of the area in which a solution must be found, they do not solve the issue themselves. We feel the general framework of alignment presented in Section 3.1 gives a good basis to work on the problem and present possible approaches to solve the alignment problem. Our own work so far has focused on inductively generating logical rulesets as an alignment, however as mentioned, this approach is very dependent on the available background information. Whether this approach is suitable or not therefore depends to a large extent on the type of trust models used and the available $\mathcal{L}_{Domain}$. In general we expect that the best approach will be largely context dependent. Some factors which may play a role in deciding which approach to use are:

– Number of agents. If there are many agents and a large number of interactions the learning-based approaches may be more suitable, while for domains with a small number of agents, argumentation may be better.
– Expressivity of $\mathcal{L}_{Domain}$. Because ILP-based solutions require a rich language, if this is not available, context-discovery may be a more suitable approach. Similarly argumentation requires an argumentation language.
– Complexity of trust models. If trust models used are very complex, machine-learning approaches may not be able to handle the task. Argumentation-based approaches may give better results.

We also do not claim to give a complete overview of methods. There may very well be other approaches we have not thought of, or combinations of the ones we have mentioned. We feel there is a lot of work to be done in this domain and we have only touched the tip of the iceberg.

Methods for discovering dishonesty should play an important role throughout the process. Not only do we feel the methods proposed could be used to find when agents could benefit from aligning, but there is a possibility for dishonesty during the alignment process. One advantage of aligning is that, because the alignment process relies on objective information, being dishonest is harder. Furthermore the results are less predictable, because the agents' trust evaluations are translated into each others', rather than being incorporated as "truthful" information. Lying during the alignment process is therefore harder, however we do not rule out the possibility and statistical methods to discover liars remain important.

References

1. Euzenat, J., Shvaiko, P.: Ontology matching. Springer-Verlag, Heidelberg (DE) (2007)
2. Bouquet, P., Ehrig, M., Euzenat, J., Franconi, E., Hitzler, P., Krötzsch, M., Serafini, L., Giorgos, S., Sure, Y., Tessaris, S.: D2.2.1 specification of a common framework for characterizing alignment. Technical report, Knowledge Web project EU-IST-2004-507482, realizing the semantic web (2004)
3. Plato: The Republic. (360BC)
4. Gambetta, D.: Can we trust trust. In Gambetta, D., ed.: Trust: Making and Breaking Cooperative Relations, Basil Blackwell (1988) 213–237

5. Dasgupta, P.: Trust as a commodity. In Gambetta, D., ed.: Trust: Making and Breaking Cooperative Relations, Blackwell (1988) 49–72
6. Bromley, D.B.: Reputation, Image and Impression Management. John Wiley & Sons (1993)
7. Celentani, M., Fudenberg, D., Levine, D.K., Psendorfer, W.: Maintaining a reputation against a long-lived opponent. Econometrica **64**(3) (1966) 691–704
8. Sabater, J., Sierra, C.: Review on computational trust and reputation models. Artificial Intelligence Review **24**(1) (2005) 33–60
9. Ramchurn, S.D., Huynh, D., Jennings, N.R.: Trust in multi-agent systems. The Knowledge Engineering Review **19**(1) (2004) 1–25
10. Conte, R., Paolucci, M.: Reputation in Artificial Societies: Social beliefs for social order. Kluwer Academic Publishers (2002)
11. Omidyar, P.: Ebay. http://www.ebay.com, retrieved September 26, 2008 (1995)
12. Reiley, D., Bryan, D., Prasad, N., Reeves, D.: Pennies from ebay: The determinants of price in online auctions. Journal of Industrial Economics **55**(2) (2007) 223–233
13. Grice, P.: Logic and conversation. In Cole, P., Morgan, J.L., eds.: Syntax and Semantics, Vol. 3: Speech Acts. Academic Press, New York (1975) 41–58
14. Orsolini, M.: "dwarfs do not shoot": An analysis of children's justifications. Cognition and Instruction **11**(3) (1993) 281–297
15. Veneziano, E., Sinclair, H.: Functional changes in early child language: the appearance of references to the past and of explanations. Journal of Child Language **22**(3) (1995) 557–581
16. Hexmoor, H., Castelfranchi, C., Falcone, R.: Agent Autonomy. Kluwer Academic Publishers (2003)
17. Chandler, D.: Semiotics: The Basics. Routledge (2002)
18. Steels, L.: Semiotic dynamics for embodied agents. IEEE Intelligent Systems **21**(3) (May 2006) 32–38
19. Abdul-Rahman, A., Hailes, S.: Supporting trust in virtual communities. Proceedings of the 33rd Hawaii International Conference on System Sciences **6** (2000) 4–7
20. 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, Honolulu, Hawaii, USA (2007) 90–102
21. 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, Utrecht, The Netherlands, ACM (2005) 505–511
22. Nardin, L.G., Brandão, A.A.F., Muller, G., Sichman, J.S.: Effects of expressiveness and heterogeneity of reputation models in the art-testbed: Some preliminar experiments using the soari architecture. In: Proc. of the Twelfth Workshop "Trust in Agent Societies" at AAMAS '09, Budapest, Hungary (2009)
23. Sabater-Mir, J., Paolucci, M., Conte, R.: Repage: REPutation and imAGE among limited autonomous partners. JASSS - Journal of Artificial Societies and Social Simulation **9**(2) (2006)
24. Hübner, J.F., Lorini, E., Herzig, A., Vercouter, L.: From cognitive trust theories to computational trust. In: Proc. of the Twelfth Workshop "Trust in Agent Societies" at AAMAS '09, Budapest, Hungary (2009) 55–67
25. Şensoy, M., Yolum, P.: A context-aware approach for service selection using ontologies. In: Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'06), Hakkodate, Hokkaido, Japan, IFAAMAS (2006) 931–938

26. Pinyol, I., Sabater-Mir, J.: Arguing about reputation. the lrep language. In Artikis, A., O'Hare, G., Stathis, K., Vouros, G., eds.: Engineering Societies in the Agents World VIII: 8th International Workshop, ESAW 2007. Volume 4995 of LNAI., Springer Verlag (2007) 284–299
27. Stranders, R., de Weerdt, M., Witteveen, C.: Fuzzy argumentation for trust. In Sadri, F., Satoh, K., eds.: Computational Logic in Multi-Agent Systems: 8th International Workshop, CLIMA VIII. Volume 5056 of LNCS., Springer Verlag (2008) 214–230
28. Jøsang, A., Ismail, R., Boyd, C.: A survey of trust and reputation systems for online service provision. Decision Support Systems **43**(2) (2007) 618–644
29. Teacy, W.T.L., Patel, J., Jennings, N.R., Luck, M.: Travos: Trust and reputation in the context of inaccurate information sources. Autonomous Agents and Multi-Agent Systems **12**(2) (2006) 183–198
30. Jøsang, A., Ismail, R.: The beta reputation system. In: Proceedings of the Fifteenth Bled Electronic Commerce Conference e-Reality: Constructing the e-Economy, Bled, Slovenia (2002)
31. Şensoy, M., Zhang, J., Yolum, P., Cohen, R.: Context-aware service selection under deception. Computational Intelligence **25**(4) (2009) 335–366
32. Barwise, J., Seligman, J.: Information Flow: The Logic of Distributed Systems. Cambridge University Press (1997)
33. Schorlemmer, M., Kalfoglou, Y., Atencia, M.: A formal foundation for ontology-alignment interaction models. International Journal on Semantic Web and Information Systems **3**(2) (2007) 50–68
34. Atencia, M., Schorlemmer, M.: A formal model for situated semantic alignment. In: Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2007), Honolulu, Hawaii, USA, IFAAMAS (2007) 1270–1277
35. Koster, A., Sabater-Mir, J., Schorlemmer, M.: A formalization of trust alignment. In Sandri, S., Sànchez-Marré, M., Cortes, U., eds.: AI Research and Development. Proc. of the Twelfth International Congress of the Catalan Association of Artificial Intelligence (CCIA 2009). Volume 202 of Frontiers in Aritficial Intelligence and Applications., Cardona, Spain, IOS Press (2009) 169–178
36. De Raedt, L.: Logical and Relational Learning. Springer Verlag (2008)
37. Koster, A., Sabater-Mir, J., Schorlemmer, M.: Inductively generated trust alignments based on shared interactions (extended abstract). In: Ninth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2010), Toronto, Canada, IFAAMAS (In Press)
38. Koster, A., Sabater-Mir, J., Schorlemmer, M.: An interaction-oriented model of trust alignment. In: Proc. of the 13th Conference of the Spanish Association for Artificial Intelligence, CAEPIA 2009, Sevilla, Spain (2009) 655–664
39. Urbano, J., Rocha, A.P., Oliveira, E.: A trust aggregation engine that uses contextual information. In: Proc. of EUMAS 2009. (2009)
40. Hermoso, R., Billhardt, H., Ossowski, S.: Dynamic evolution of role taxonomies through multidimensional clustering in multiagent organizations. In: Proc. of EUMAS 2009. (2009)
41. Ramon, J., Bruynooghe, M.: A polynomial time computable metric between point sets. Acta Informatica **37** (2001) 765–780
42. Aamodt, A., Plaza, E.: Case-based reasoning: Foundational issues, methodological variations, and system approaches. AI Communications **7**(1) (1994) 39–59