Some open questions for the integration of trust with negotiation, argumentation and semantics

Piero Bonatti
Universitá di Napoli Federico II, Italy
bonatti@na.infn.it

Eugenio Oliveira
University of Oporto, Portugal
eco@fe.up.pt

Jordi Sabater-Mir, Carles Sierra
Artificial Intelligence Research Institute, Barcelona, Spain
jsabater@iiia.csic.es, sierra@iiia.csic.es

Francesca Toni
Imperial College London, UK
ft@imperial.ac.uk

September 24, 2010

Abstract

Agreement Technologies are needed for autonomous agents to come to mutually acceptable agreements, typically on behalf of humans. These technologies include trust computing, negotiation, argumentation and semantic alignment. In this paper we identify a number of open questions regarding the integration of computational models and tools for trust computing with negotiation, argumentation and semantic alignment. We consider these questions in general and in the context of applications in open, distributed settings such as the grid and cloud computing.

1 Introduction

Agreement Technologies refer to computer systems in which autonomous agents negotiate with one another, typically on behalf of humans, in order to come to mutually acceptable agreements [37, 38]. Agreement Technologies include trust computing, negotiation, argumentation and semantic alignment.

Argumentation is a powerful technique aiming at the evaluation of possible conclusions/claims by considering reasons for and against them.

These reasons (arguments and counter-arguments) provide support for and against the conclusions, through logical reasoning. Although originally proposed within the realms of logic, philosophy, and law, in the last decade argumentation has attracted wide interest in computing to understand and meet the challenges of a number of applications characterised by the lack of certain, consistent and complete information, and when numerical (e.g. statistical) information is not available or is only partially available ([3, 46].

Negotiation allows agents to try and reach satisfactory agreements, fulfilling the needs and requirements of all agents involved. Argument-based negotiation differs from other approaches to negotiation in that it makes use of offers, counter-offers and refusals supported by arguments that explain the reasons behind them. In contrast with game-theoretic or heuristic negotiation [45], agents involved in argument-based negotiation may change their preferences, utilities, even issues and goals on-the-fly.

Semantic alignment of ontologies aims at resolving the issues arising from distributed and independent knowledge creation and management. In an open framework where different agents may use different ontologies to drive their behavior and interact with each other, it is crucial to set up methods and mechanisms to find semantic relationships among syntactically unrelated concepts, in order to reach a shared mutual understanding of requests, offers, contraints, goals, etc. [36, 12].

Trust measures are responsible for guaranteeing security on execution and help agents to determine with whom to interact and what terms and conditions to accept as a basis for the interaction [53]. Trust is usually understood as expectation on behaviour post-commitment [14]. Usually agents establish agreements that are later on executed, that execution might not be as expected due to many reasons: execution or communication errors, defection or misunderstandings, just to put a few. Reputation is usually defined as a group's opinion on an aspect of a thing [52]. For instance, the reputation of a person as a good seller. Trust and reputation are closely related as trust, built along time, is the basis to express an individual opinion later aggregated into a reputation. And a reputation value may be used to build trust.

In this paper, we consider several open questions concerning the integration of trust (and reputation) computing techniques and other Agreement Technologies (negotiation, argumentation and semantic alignment), in the context of existing work. The open questions identify trends and challenges for future work in this area. In particular, we focus on the integration of trust and negotiation models (section 2), the integration of trust and argumentation models (section 3), the use of semantic alignment techniques for trust (section 4), as well as current practical bounds (section 5) and challenges (section 6) for these agreement technologies in practical applications and in particular in open, distributed settings such as the grid and cloud computing. Finally, in section 7 we conclude.

2 Trust and negotiation

Trust/reputation and negotiation models can be integrated to benefit one another in (at least) three different ways:

- the outcome of negotiation can be a set of digital credentials, supporting the trustworthiness of potential partners in interaction and contributing to their reputation;
- trust/reputation models can help determine the right partners prior to or as a result of negotiation;
- trust/reputation models can help improve the outcome of negotiation.

We consider these three aspects of the integration in turn.

2.1 Trust Negotiation

The integration of trust and negotiation models has already been profitably worked out in the area of computer security and privacy, in the form of *trust* (or *credential*) *negotiation* [61, 55]. Trust negotiation is based on the automated negotiation of digital credentials and other forms of evidence in order to implement attribute-based access control without placing the whole burden of information exchange on the shoulders of users. The following is a classical example of trust negotiation: Alice visits for the first time an on-line bookshop that sells an interesting book at a very cheap price. By interacting with the bookshop's server Alice's personal agent finds out that she has to provide either her credit card number or a pair $(userId, password)$ for a previously created account. Alice does not want to create a new account on the fly, so releasing her credit card is the only option. However she is willing to give such information only to trusted on-line shops (e.g. belonging to the BBB[1]), therefore her personal agent asks the bookshop to provide such information. The bookshop belongs indeed to the BBB and is willing to disclose such credential to anyone. This satisfies Alice's policy, so her personal agent provides her credit card number and the transaction is successfully completed.

Some of the main advantages of trust negotiation are:

- privacy enhancement, because it is not required to disclose the full identity of the user each and every time; access control can be based on selected properties, such as age, nationality, etc.;
- digital credentials constitute a more reliable kind of information, as opposed to the unverified profiles associated to internet accounts;
- improved user experience, as the operations related to security and privacy become more transparent, and server properties (e.g., certifications) can be automatically checked, thereby making navigation safer.

Trust negotiation is a well-established area: many interesting approaches and results already exist (e.g. see [5]), so the real question is what still

[1] `http://www.bbb.org/`

needs to be done. In the remainder of this section we consider few possible directions for future work.

One important issue concerns the design of privacy-enhancing negotiation strategies that effectively reduce the amount of disclosed information. In this respect, strong guarantees are hard to achieve given the lack of trusted third parties on the web, and given the lack of mutual knowledge among the negotiating peers. Market design mechanisms seem a promising direction, as a way to establish a system of incentives and disincentives that tend to minimize the amount of user information requested by servers for access control. Another interesting open issue is how policies themselves can be negotiated, possibly relaxing some requirements in a controlled way, in order to increase the number of successful negotiations and get compensation for unsatisfied security/privacy preferences.

2.2 Trust/reputation models to select partners in negotiation

One of the standard uses of trust and reputation models is to allow an agent to choose the right partner for an interaction. In this section we focus on situations where this interaction is a negotiation, e.g. for (one or more) services in grid computing or in cloud computing settings.

One of the main aspects that a negotiator, being it human or software, has to assess in a negotiation process is the true capabilities of any potential partners in joint actions [50]. The capabilities are of two fundamental types: functional and non-functional. When these actions amount to the execution of software services, functional capabilities refer to the input/output relationship of software execution, that is *what* the software can actually compute. Non-functional capabilities refer to *how* the computing is actually executed, e.g. how long does it take, what percentage of errors are produced, how many resources are required, etc.

If we understand trust as expected behaviour after commitment [14], it is clear that trust measures can be used to assess both types of capabilities. Functional, because the input/output relationship is a commitment (specification) on behaviour, and non-functional as well, because negotiations involve quality of service agreements that are commitments that can later on be observed. For instance, if a service commits to compute a $n \times n$ matrix inversion in $n^2 \times 200ms$, the client can check whether that is the case and learn a model of service behaviour that will help the client in assessing the trust on future service commitments. This assessment will determine whether to negotiate with the service in the future and what service-level agreement to agree upon.

Using reputation models to evaluate a partner as a negotiator is not very different from using reputation models to evaluate any other characteristic of that partner. In this sense, current reputation models (for example models like [47, 35, 23]) are suitable for this task. The integration with the negotiation model implies, among other things, identifying the set of features that typify a negotiator and establishing a classification of types of negotiators. Both things would allow the exchange of reputation information regarding negotiation. Several issues are still un-

resolved/ongoing in this setting. We identify a number of these below.

2.2.1 Commitments and trust

For trust to be assessed in the context of commitments, a suitable representation and ways for determining fulfilment of commitments are needed. These may need to take into account "mitigating circumstances" that may have determined or contributed to the non-fulfilment of commitments and should not contribute negatively towards the assessment of trustworthiness of an agent. Also, events that have positively impacted on fulfilment of commitments but are outside the control of the target agent should not contribute positively towards the trustworthiness of agents. We will discuss, in section 3.1, how argumentation can support some of this reasoning.

2.2.2 Context and trust

Trustworthiness may need to refer to specific aspects of the behaviour of agents (or of the services they provide). For example, an agent may trust a service provider as far as the functional properties of the services it provides are concerned, but not about some non-functional property (e.g. timeliness of the delivery). Then, the selection of an agent to interact with becomes a multi-attribute decision-making problem. Information about the trustworthiness, or lack thereof, of an agent concerning specific aspects may determine the negotiation strategy of the other agents (see also section 2.3). For example, an agent may insist on a strongly binding service-level agreement for an aspect it does not consider the interlocutor agent trustworthy about.

Also contextual attributes that characterize the deal under negotiation should play an important role on both how to build up trust and reputation profiles and how to use them for negotiation. A real useful computational trust and reputation model, more than providing one or two indicators on whether to trust a potential partner, has also to provide the contexts for which those indications are given. A potential partner may be good at supplying some good G in general. However, if this partner is hurried to deliver G in a very short time, it may fail to deliver G appropriately. Therefore, classification and clustering algorithms can be used to attach trust and reputation measures to agents according to the desired contractual conditions so that they can be accurately used during specific negotiation processes. By using known context-based (attribute-value) features frequency metrics, and adopting inter and intra clusters similarity measures [39], clusters of agents that behave similar in closer situations can be formed and the respective stereotypes characterizing those extracted clusters [58]. Also Information Gain based methods are being successfully tried out to improve accuracy of agents behavior tendency extraction, contributing to a more reliable online, situation-aware trust model [58].

2.2.3 Bundles and trust

If a combination of items, e.g. services, is negotiated with several agents, then there is a key issue that needs to be addressed and where trust plays a role: *incompatibility between services*. Incompatibility may derive from semantic mismatches between the specification of the services or due to differences in non-functional requirements (e.g. time, quality, optimality). On the latter aspect, trust may determine the overall level of compatibility of the final solution, by computing a level of expectation over the non-functional behaviour of the solution by using previous experience with each service and by looking into the concrete orchestration and choreography of the services.

2.2.4 Groups and trust

There are two aspects that are relevant from a trust perspective when agents team up to solve a particular problem. First, the *composition* of a team: agents that may be trustworthy in isolation may turn out to be untrustworthy when teaming up with other agents, or, conversely, agents that are not to be trusted in isolation may be fine when teamed up with others. Thus, the composition of a team turns up to be a complex combinatorial problem. Second, the trust *aggregation*: as a consequence of the first issue, the trust on a group as a whole is not functional, that is, it is not a simple aggregation of the trust on the individuals. Information about the compatibility between members is required to perform this calculation. Again, a combinatorial problem appears.

2.2.5 Negotiation skills and trust

If an interaction involves a negotiation process, choosing the right partner should imply not only the evaluation of the partner from the point of view of the service it can offer but also the expected characteristics of the negotiation that will take place before agreeing upon the service. Indeed, it is worthless for an agent to select a partner that hypothetically can offer the best service if it is known a-priori that the interaction will never happen because the negotiation process will fail. This failure in the negotiation stage can be for different reasons and some of them can be foreseen by using trust and reputation models. One possible reason for failure is the lack of suitable negotiation skills. For example, it may be the case that an agent's negotiation mechanism is too simple to obtain a good deal when the agent interacts with partners with greater negotiation skills. Knowing the reputation of the possible partners as negotiators can help an agent to balance (i) the quality of the service they potentially can offer with (ii) the actual possibilities for the agent to obtain that quality after a negotiation process.

2.2.6 Organisations and trust

Whenever there is a need for entities to negotiate possible joint activities, seeking information about potential partners may be accompanied

by noisy information and possibly by cheating, depending on the degree of competitiveness involved in that particular negotiation.

When negotiation between different entities implies some kind of competition, available information about different negotiation players needs to be clearly backed by the system/institution through which the negotiation process takes place. A clear separation has to be enforced between situation-based (*context sensitive*) trust and socially generated information (reputation). The former is typically local, sporadic, associated to an agent-specific functionality, activity or even role; also, it is mostly dynamic: it decays over time unless new evidence arrives [34], it is asymmetric, in the sense that it grows slowly and decays fast, and it embeds the maturity property [59]. The latter is more general, broader in scope, sometimes fuzzy, and can only be acceptable if and only if it obeys some explicit, institutional, normative rules governing trust information propagation in a network in such a way as to prevent unfair generalization and abuse of the transitivity property[57].

Agents interoperability, negotiation, cooperation and joint work only become credible if regulated by institutional norms. Norms, prescribing obligations, sanctions and, possibly, rewards, are enforced to entities (agents) by some kind of accepted institutional power (electronic institution). Usually the outcome of a negotiation contract is a (formal or informal) contract, whose clauses may trigger, in the normative environment, specific norms. According to a model inspired by contract law theory, prescribing the use of norms and default rules for different contexts, trust and reputation measures can be beneficial to select appropriate norms that have proven to be more adequate to situations where specific types of Agents (more or less reputable) are present. This will lead the normative environment to adapt to the current situation, trying to be fair according to what has been perceived before, about the contracting agents. Also, incorrect broadcasting, through computational trust and reputation mechanisms, of one agents perception on other agents' trust that could lead to incorrect bad reputation propagation, will be monitored and sanctioned by the normative environment. These procedures will entangle the computational trust and reputation mechanism with both the negotiation process and the monitoring capabilities of a normative environment leading to more fairness in norms application.

2.3 Trust models to improve the outcome of negotiation

In the previous section we have seen how trust models can help choosing a partner according to the negotiation skills. Once the negotiation process has started, the trust and reputation model can play a new role, the role of conducting or at least influencing how the negotiation process takes place. This use of the trust and reputation model not as a mere partner selection mechanism but as a piece of the decision making module implies a tighter integration of the model with the rest of elements of the agent architecture, something that in current models has been somehow neglected. Not only that, the integration has to be also made at the orga-

nizational level, involving the trust and reputation model in the decisions about communication and social related actions.

2.3.1 Trust/reputation to direct negotiation

Once the negotiation process starts, all the knowledge agents have about the negotiation strategies used by their partners can help to improve the outcomes of the negotiation. Again, the reputation of a partner as negotiator (not only how good it is at negotiating but also the kind of negotiation mechanism it uses) can be useful to direct the negotiation process.

For example, if our partner has the reputation of trying to force always a quick negotiation because it has a very good algorithm for a specific kind of negotiations, we can decide to slow down the rythm of the negotiation hoping for possible mistakes in our favour. Or if we cannot force the negotiation conditions, we can simply use the reputation information to choose the right negotiation algorithm (in this case a fast and probably simple algorithm) that allow us to be competitive. Of course, the context plays a very important role in this use of reputation for negotiation. The selected negotiation strategy usually will not only depend on the reputation of the partner as a negotiator but also on the context surrounding that particular negotiation. In our previous example we can probably slow down the rythm of the negotiation only if we have a stronger position than our partner.

2.3.2 Context-based contracts negotiation

Normative environments [9] may impose, to consortia resulting from negotiated agreements, for *contracts* to be of different types according to partners' degree of mutual trustworthiness. Partners with lower credibility may get tougher electronic contracts including more strict clauses and harder sanctions in case of commitment failure.

When trust is used for partner selection (see section 2.2) trust values can be seen as hard constraints (for example a pre-requisite for B2B partner selection processes) prior to negotiation. However, simultaneously, trust values can be seen as a kind of soft constraint influencing (driving) a (possibly argumentation-based) negotiation process whose outcome is the final decision on whether to work jointly or not. This is often the case in real B2B negotiations. In this setting, a progressively more detailed information exchange, starting with asking information about past partners, may end by asking private, financial information together with specific guarantees about future behavior. This means that "*proposing/counter-proposing*" goes along with "*asking for/answering to*" certain kinds of guarantees closely related with the importance of what is being negotiated and the involved calculated risk.

Negotiated contracts can be used through conceptual *clustering* (non-supervised learning) algorithms to create "trustworthiness" stereotypes, namely sets of the most discriminate features characterising the elements of each cluster [60]. Associating meaningful stereotypes to clusters of previously known entities (for example, negotiation partners) may be of great help when dealing with *newcomers* trying to temporarily include

them in known clusters, and deducing some kind of expectation on their future behaviour without being either too unfair or too naive. When more interactions take place and more information becomes available, these more recent known entities can either confirm or deny their classification.

Computational trust and reputation models have already been included in Electronic Institution platforms, mainly for regulating the all Virtual Enterprises life cycle [9, 59]. However, we envisage, for the future, that these modules will be at the heart of the agents decision making whenever they decide to communicate in the framework of a social network [28]. This will permit to evaluate which kind of relationships have been established in the network, regarding stability, openness, specific nodes criticality and centrality, sub-networks aggregation or disaggregation forces, etc. Playing with different trust and reputation models may lead to very different (social) networks, with different characteristics. Conclusions will thus be drawn about what features should be included in the trust and reputation modules that may lead to either more social aware or risk tolerant social networks.

3 Trust and argumentation

Argumentation can be used to *justify* an observed behaviour, and therefore modify the impact on the model of expected behaviour that is at the base of trust models [52]. For example, a partner in an interaction could provide an evaluator with a justification of the form "I could not deliver the package yesterday because the driver had a flu". The recipient of the information may decide to believe that and thus not alter (decrease) the trust on the interaction partner (the courier service company in our example).

There are (at least) three forms of integration, beneficial to either or both models:

- a trust model can be used to determine the reliability of an argument
- an argumentation process can be used to improve the performance of a trust model
- argumentation can be used in support of trust negotiation (see section 2.1)

We consider these three aspects of the integration in turn.

3.1 Trust for argumentation

3.1.1 Provenance and argumentation

During an argumentation process there is an exchange of arguments that are built using different pieces of knowledge. Whether an argument is finally accepted or not depends, amongst other things, on the structure of the argument and the "truth" behind the knowledge used to build the argument. More than often, the "truth" of the knowledge that is used to build the arguments may depend on the source of that knowledge. The stream of information received from an informant together with its

reputation are the only elements an agent (or human user) can use to decide whether an argument, built from that information, can be accepted or not.

The use of trust models for this purpose is straightforward. Currently, one of the main uses of trust models in multi-agent systems is the evaluation of a piece of information regarding the source (who is the origin) of that information. Extending this kind of evaluation to each one of the elements of a (structurally correct) argument, we could evaluate the truthfulness of that argument. It would be interesting to explore definitions of acceptability of arguments using a combination of standard dialectical acceptability and trustworthiness/reputation of the source of the argument.

3.1.2 Social relations and argumentation

The importance of social relations in the acceptance of arguments has been long studied by psychologists [25]. Arguments produced by someone holding authority over the hearer usually have a large impact in terms of bigger changes in the beliefs although the impact is shorter along time and hearer to their previous beliefs faster than in the case of a peer generating the argument. Exploring the social dimension seems an interesting line of work that has not been studied in depth yet.

3.2 Argumentation for trust

A different but complementary approach is to use argumentation to improve the performance of the trust (or reputation) mechanisms. This can happen when trust is computed by an agent in isolation (relying on past interactions with a target) or when agents can exchange and share information about the trustworthiness of possible targets (or one another).

Computing trust is a problem of reasoning under uncertainty, requiring the prediction and anticipation by an agent (the evaluator) of the future behaviour of another agent (the target). Despite the acknowledged ability of argumentation to support reasoning under uncertainty (e.g. see [27]), only Prade [44], Dondio & Barret [15] and Parsons et al [40] have considered the use of arguments for computing trust in a local trust rating setting. Dondio & Barret [15] propose a set of trust schemes, in the spirit of Walton's argument schemes, and assume a dialectical process between the evaluator and the target whereby the evaluator poses critical questions against arguments by the target concerning its trustworthiness. Prade [44] proposes an argumentation-based approach for trust evaluation that is bipolar (separating arguments for trust and for distrust) and qualitative (as arguments can support various degrees of trust/distrust). Parsons et al [40] define an argumentation logic where arguments support measures of trust, e.g. qualitative measures such as "very reliable" or "somewhat unreliable".

There are several non-argumentation based methods to model the trust of the evaluator in the target. Sabater and Sierra [48] classify approaches to trust as either "cognitive", based on underlying beliefs, or "game-theoretical", where trust values correspond to subjective probabilities and

can be modelled by uncertainty values, Bayesian probabilities, fuzzy sets, or Dempster-Shafer belief functions. The latter approach is predominant nowadays for trust computing. However, Castelfranchi and Falcone [11] argue against a purely game-theoretic approach to trust and in favour of a cognitive approach based upon a mental model of the evaluator, including goals and beliefs. Moreover, some works (e.g. [54]) advocate the need for and benefits of hybrid trust models, combining both the cognitive and game-theoretical approach.

Recent work [32] proposes a hybrid approach for constructing Dempster-Shafer belief functions modeling the trust of a given agent (the evaluator) in another (the target) by combining statistical information concerning the past behaviour of the target and arguments concerning the target's expected behaviour. These arguments are built from current and past contracts between evaluator and target, and are integrated with statistical information proportionally to their validity. The method extends a standard method for trust [63] that relies upon statistical information only. The two methods have identical predictive performance when the evaluator is highly "cautious", but the hybrid method gives a significant increase when the evaluator is not or is only moderately "cautious". Moreover, with the hybrid method, target agents are more motivated to honour contracts than when trust is computed on a purely statistical basis.

3.2.1 Argumentation profiles and trust

When agents engage in argumentation dialogues they show different profiles. For instance, an agent may be more persuasive than others in the sense that the final opinions of the others vary more significantly, or it may predetermine the result of the argumentation process as the final joint opinion is closer to its initial position than to the others'. Given that trust is usually related to expected behaviour, the composition of a team or a committee where the members have particular profiles is key in order to assess the trustworthiness of the team in reaching decisions. This aspect of argumentation has not been studied in detail. Some initial ideas in the area of reputation can be found in [52] and in the reputation model of the LiquidPub[2] project.

3.2.2 Trust information exchange

Agents can exchange trust/reputation information in the form of arguments, e.g. as in [30]. One of the key aspects in argumentation [51] is *when* and *whom* to disclose *what* type of information. The behaviour of agents, services and humans is just one type of such information that can be disclosed. The information that "My butcher gave me shocking chops yesterday!" may persuade others to buy somewhere else. This information may be incorporated in trust models of others depending on their interpretation of the information and their trust in turn on the agent providing the information [52].

[2] `www.liquidpub.org`

3.2.3 Communicating justifications

Instead of just communicating an isolated trust/reputation value, an agent can communicate the value plus a justification of that value. After that, an argumentation-based dialog can help to decide if that trust/reputation value is worth to be added to the agent's mind as new knowledge [41]. Given that, how the reputation model calculates the reputation value and therefore how this justification is built becomes very relevant. If we want that agents and humans share common virtual spaces and therefore interact among them, using a cognitive approach in the design of reputation models is the best approach. In these kind of environments we need agents that can build justifications like humans so other humans can easily understand those justifications. The agents have to be able to understand also arguments given by humans. To achieve that, from a computational point of view we need several elements:

- A cognitive theory on reputation. This is the starting point and for that, we can rely on social sciences that have studied reputation from a cognitive perspective for many years. A good example of cognitive reputation theory is that from Conte and Paolucci [13].
- A computational model of that theory. At the end, we want the reputation model to be into an agent's mind so the initial reputation theory has to be formalized and implemented.
- A language that allows to express the concepts related to reputation as logical predicates. The agents have to be able to analyze the justification and reason about it.
- A framework that integrates everything in a cognitive architecture. Although simple reputation models can be considered isolated and therefore the justification can be built only from internal elements of the model, once you move to more complex reputation models the connection with other modules of the agent needs to be taken into account. Trying to justify properly a reputation value without taking into account the planner or the goals of the agent is not feasible. A proper justification will imply to follow the justification thread outside the reputation model. Because we want a cognitive justification, having a cognitive agent architecture that allows to follow a justification thread outside the reputation model is the best approach.
- A reputation ontology and a language based on that ontology to talk about reputation. At this point, the agent is able to build a justification taking into account not only the elements strictly attached to the reputation model but also to other relevant modules of the architecture. Now it has to be able to communicate that justification and therefore a language and a common ontology are required.
- An argumentation framework and a dialog protocol to allow agents to seek information. We have the justification, we have a language that is understood by the agent partners and can be used to communicate the justification. The last thing needed is an argumentation framework and a dialog protocol that allow the agents to argue about that reputation value.

The final result of all this process is that the agents should be able to exchange reputation values, justify why they think those values are right and discuss with others about them changing their minds if the arguments are convincing enough. Although complex, all this process can considerably improve the usefulness of reputation values, giving them an extra level of reliability.

3.2.4 Ignorance and Trust

A negotiation system, and an electronic institution "supervising" the negotiation process, should be clear about the contexts in which it is possible for entities to provide faithful information on trust. Moreover, an electronic institution should be able to prevent trust measures from being formed and shared based on misleading grounds. This implies that there are situations in which trusting/not trusting are not the only options. Indeed, some kind of measure of "*ignorance*" should be computed (as it is the case with the use by Dempster-Shafer models of the plausibility concept). An argumentation-based process can then be followed for the sake of decreasing the current ignorance if this is higher than a threshold, in order to reach a situation in which agents believe that either their trust in an entity (agent) is reliable or not. Normative environments, by making clear rules available as well as norms that may sanction undesirable behaviors [9], can be designed to guarantee prevention of incorrect, unfair or unjustified agents reputation measures broadcasting, mostly in the same manner as anti-spam systems are needed to block undesirable messages to propagate along a communication network.

3.3 Argumentation for trust negotiation

Argumentation has not (yet) been applied in the context of trust negotiation. In most trust negotiation systems access control and credential disclosure decisions are actually based on proofs that constitute arguments in favor of permitting the incoming request. Negative decisions usually result implicitly from the failure of all policy rules (i.e., there are no counter-arguments).

In the broader literature on computer security, one can find security models supporting positive and negative authorization inheritance with exceptions, that more closely resemble an argumentation system: the proof of a specific (positive or negative) authorization (e.g. "staff can access confidential files") may defeat a recommendation for a more general case (such as "users cannot access confidential files"; see also the survey in [4]). Conflict handling is specified via explicit rules, see for example [24].

3.3.1 Extending trust megotiation with negative authorizations and conflict resolution

In principle, in trust negotiation there is space for such nonmonotonic mechanisms because (i) authorization overriding provides an appealing way of formulating policies incrementally, through iterative refinements,

and (ii) different contexts require different forms of conflict resolution, therefore a flexible, rule-based approach is attractive.

Such extensions, however, raise at least three open questions. First, it is not yet known how negation and conflict resolution would affect the computational complexity of automated trust negotiation.

Second, a similar question applies to negotiation *strategies*. It is not only a matter of complexity; in the literature strategies are usually required to satisfy an *interoperability* property (strategies should not cause unnecessary negotiation failures) that currently has been proved only in the absence of overriding and conflict handling [64, 65, 2]; how would the latter interfere with strategy interoperability?

Third, the impact of the new features on usability – e.g. due to the subtleties of nonmonotonic reasoning – should be systematically investigated, as we know that lay users find it difficult to use even extremely simple and streamlined policy languages, see for example the user studies reported in [49]. Intelligent authoring tools (like the policy learning approach of [49]) may help to address such problems, however we are only at the beginning of this line of research, and again the more powerful the policy language, the more difficult the development of effective authoring tools.

3.3.2 The potential of argumentation: unreliable knowledge and policy negotiation

So far we have only touched integrations of argumentation and trust negotiation that – to a certain extent – are limited to extending trust negotiation with mechanisms that are already supported in logic-based policy frameworks without negotiation. However, such rudimentary, specificity-based forms of argumentation do not address a number of interesting issues that naturally arise in modern web scenarios, where general argumentation techniques may bring significant benefits and show their full potential.

First consider that the quality of knowledge may be highly dishomogeneous across different sources. Some works report systematic knowledge representation errors that trigger inferences that are logically correct (w.r.t. the encoded knowledge) but do not reflect the true state of affairs [22]. So a natural question is: can argumentation notions such as attacks and counterattacks be enriched with provenance, reliability and authoritativeness meta-information in such a way that argumentation can be used to address the problem of policy enforcement in the presence of unreliable and noisy knowledge? By comparing argument reliability and authoritativeness, it may be possible to improve the quality of (possibly distributed) policy-related decisions. Furthermore, similar ideas can be exploited to improve efficiency by restricting reasoning to a reliable subset of the relevant sources available on the web, as in [22].

Another topic – of great interest for companies – concerns extending automated trust negotiations with a bargain-like mechanism that may increase the number of successful transactions by progressively relaxing agent policies, for example by trading service improvements for more information about users. In this perspective, policy rules are treated like soft

constraints that may be relaxed if this brings sufficient benefits. Again, traditional trust negotiation mechanisms are too rigid for this kind of *policy negotiation*, and the role of argumentation techniques in evaluating the pros and cons of relaxing a policy is still to be explored.

4 Trust and semantics

Trust information is "linguistic" —thus semantics is essential. It is linguistic in the sense that behaviour is 'described' by words. For instance, if we need to describe the input/output relationship of a service, the input and output descriptions may not be equally understood by everybody. Semantic alignment permits to solve potential meaning mismatches. Also, when extracting information from past negotiation experiences, semantic similarity may help in extracting the most of it. For instance, if the promised lamb chops appeared to be excellent then the kid leg I'm offered might be similarly good. Semantics and ontologies are fundamental in human trust modelling and should be equally important in automated trust models.

4.1 Semantics for trust

Semantics are fundamental for trust (and reputation) models as they are deployed in open environments where agents designed and implemented by independent developers have to interact and use trust and reputation mechanisms to increase the robustness of the overall system. In this setting we cannot guarantee that all agents will use the same trust and reputation mechanism. In fact, the current situation is that there is a big diversity of models, each approaching the problem from a different perspective. These models frequently even differ on basic concepts such as what trust is and what reputation is. On the other side, trust (and especially reputation) models often require that there is communication among agents and exchange and sharing of information (as we discussed in section 3.2.2). Given these considerations, it is clear that some kind of semantic alignment is necessary to allow the agents to exchange information about trust and reputation.

A straightforward approach to the problem is the use of a common ontology that can work as a bridge between the different models. Although some attempts to define such an ontology can be found in the literature [42, 10], it seems quite difficult to finally obtain an ontology that can be accepted by all the scientific community and that can give support to the different levels of complexity required by the different models. For example, there are models that use labels to represent reputation values while others are using fuzzy numbers or probability distributions. A possible alternative [26], currently under consideration, is to use dynamic semantic alignment. In that case, the alignment is performed for each specific interaction. This would make possible to tailor an alignment that exploits much better the characteristics of the models being aligned. Whatever be the approach, what is clear is that some kind of alignment is necessary to allow heterogeneous models in the same environment.

Trust arises from the analysis of a wide range of evidence of various strength, incrementally gathered during a sequence of interactions (e.g. transactions) in direct as well as indirect ways (like reputation systems do, for example). Semantic techniques address relevant problems such as recognizing problematic agents, combining heterogeneous information. Remarkably, in some trust negotiation systems like Protune [7, 6] agents exchange their requirements by disclosing selected parts of their policies in the form of logic programming rules. Such *shared knowledge* is a semantic technique in itself: it allows agents to interoperate through a knowledge-based formulation of their policies grounded on concrete data items such as X.501 digital credentials. All that agents need to understand a priori in Protune is the meaning of three predicates, because all the other application-dependent and domain-dependent abstractions adopted in Protune policies can be defined in terms of those three predicates, in a machine understandable way. While this is an interesting proof of concept, showing that an ontology infrastructure is not necessarily a prerequisite for interoperability in a trust negotiation system, we remark that in the domains where such a semantic infrastructure is available, semantic policies can benefit from the use of the available ontologies. Hybrid formalisms (integrating rules and description logics) are going to play an essential role in this area; researchers on policy languages should evaluate the many existing approaches w.r.t. complexity and expressivenenss, taking in due account the specific requirements of policy languages.

Finally, when negotiation is integrated with information exchange to ascertain trustworthiness (see section 2.3), a shared ontology of concepts and properties in the domain is required. When this is not available, semantic alignemnt techniques are needed.

4.2 Trust for semantics

A much less explored area of research is the use of trust techniques for semantic purposes. In particular there are two research venues that require further attention: development of *semantic services*, e.g. for making alignments, and trust on *semantic resources*, e.g. trust on an ontology.

The development of semantic services seems unavoidable in an open world where the ontologies are being developed in large quantities. When orchestrating services to form complex solutions, semantic services will be required to make the glue among them (e.g. align ontologies of different services) as well as to facility the choreography between services that communicate using different messaging languages.

When querying ontologies, the answers' quality will depend among other issues on the provenance of the ontology. For instance, answers to queries on astronomy will probably be more trustworthy if coming from an ontology developed by NASA. Thus, the concept of *trust on an ontology*, or on combinations of ontologies, takes shape as a new concept that requires deeper analysis.

Of course, knowledge quality is not the only relevant parameter in this context; the methods and the algorithms that implement a semantic-based service are equally important, as well as the related issue of how much a

particular service or agent can be trusted. A natural approach consists in evaluating a trust measure by observing the behavior of a service or agent across a sufficently long sequence of transactions. However, for this purpose, it should be possible to tell whether two or more transactions involve essentially the same agent; this may be difficult when agents can clone or masquerade themselves. This problem has been partially addressed by the security community in several ways. For example, mobile code is usually signed by the producer and privileges are assigned to the code according to its signature (see for example the latest Java security model [19]). In this case, trust is actually placed in the producer; however, the signature can be used to identify the code, as well.

Another possible approach to recognizing an agent across multiple transactions is *certificate fingerprinting* [29]. Roughly speaking, the idea is that even if an agent (or a user) is not authenticated and transactions are guaranteed only by means of anonymous certificates, it is frequently possible to tell that a set of transactions has been performed by the same agent because some specific certificates have been repeatedly used in those transactions. Certificate fingerprinting actually identifies the agent's *principal* (i.e., the entity on whose behalf the agent is acting, that is the real owner of the certificates). In order to focus on agent recognition, certificates should be assigned directly to agents.

Whether these methodologies can be effectively applied for agent trust evaluation, and whether semantic techniques can improve their precision, are two interesting open questions that deserve further investigations.

5 Towards real world scenarios

If agreement technologies, including negotiation, argumentation, trust and semantics techniques, are to be used in real world scenarios, practical bounds to these techniques need to be removed. To do that there are two needs that have to be satisfied: testbeds to compare different approaches and efficient models that can scale.

5.1 Testbeds for trust

There are not many testbeds developed for trust in combination with negotiation, semantics or argumentation. We mention two of them that offer interesting opportunities.

5.1.1 ART Testbed

The ART testbed[3] [18] is a testbed for computational reputation models designed to serve in two roles: (1) as a competition forum in which researchers can compare their technologies regarding reputation models against objective metrics, and (2) as a suite of tools with flexible parameters, allowing researchers to perform customizable, easily-repeatable experiments. The testbed presents an scenario with simulated clients that

[3] http://megatron.iiia.csic.es/art-testbed/index.html

request appraisals for paintings from different artistic eras. The participating agents have to provide these appraisals to the clients, however they are not experts in all the eras so if an appraiser does not have the expertise to complete the appraisal, it can request opinions (paying accordingly for them) from other appraiser agents. The dilemma the agent has to confront is either to help other agents appraising as best as possible the requested paintings (and therefore wining a good reputation among appraisers as informer) or cheat them so the clients who ordered the appraisal to the requester agent be unhappy. The reputation model is the tool that the appraiser uses to take decisions at this level. Unfortunately the ART testbed, though still available, is no longer maintained.

5.1.2 DipGame Testbed

The DipGame testbed[4] is based on the classical Diplomacy board game. The game is played by seven players incarnating political powers at the beginning of the XXth Century with the goal of conquering Europe. It has the interesting property that there are no chance moves and that the main ability to win is to be a good negotiator. Not only that, negotiations are private but decisions are public, therefore it is easy to verify whether secret pacts are kept or not, and therefore a model of trust on the other players can be built along a game, as there are two negotiation rounds per year (the game progresses along time) and games usually last for quite a number of years. The language used in the negotiation interchanges can be set to be as complex as desired (i.e. up to natural language) and therefore can potentially be used to also test argumentation models.

Powers can be incarnated by either humans or by softbots making the testbed ideal to compare negotiation strategies among themselves and with humans. The testbed offers facilities to build your own softbot in Java. See [17] for details.

5.2 Scalability

The most urgent practical problem to address in trust and its relation with negotiation or semantics is scalability (e.g. how to reason over different ontologies and how to negotiate with different agents). In networks of thousands or millions of interoperating services/agents, a solution to scale up agreements is needed. Social techniques might be required to focus the search, around known and trusted neighbours, and semantics might be required to scale up the search for similar interests. Also, organisations might help by following a divide and conquer approach, clustering agents according to criteria like similarity in the preferences or in semantics.

By contrast, scalability is not a major problem with access control policies. Indeed, currently, real world policies are not very large, partly because trust negotiation systems are not yet widely deployed therefore the application of articulated policies would place an excessive burden on users. Experimental result such as those reported in [6] show that there are no immediate issues related to performance when using these policies.

[4] `www.dipgame.org`

There are, however, serious difficulties related to usability. Recent studies show that users write very poor policies, even if the policy language is extremely simple [49]. An improper formulation of the user policy undermines the effectiveness of trust negotiation, and may cause privacy violations.

Some technological issues form serious practical bounds for agreement technology. For example, the credential infrastructure (e.g. Digital Ids) have not been picked up yet by industry — standard login/password are used instead.

6 Challenges in large-scale open distributed systems

6.1 Integration with other techniques

Models of negotiation, argumentation and trust, as well as semantics techniques, typically need to be deployed in large-scale open distributed systems such as the grid and cloud computing platforms. Moreover, they may need to be integrated with current technological applications such as crowd-sourcing, that are also deployed in large-scale open distributed systems. This integration presents in itself several technical and conceptual challenges, e.g. concerning the automated assessment of interactions before this assessment can be fed into the tust and reputation model used.

6.2 Usability

Usability is one of the major challenges to be addressed to foster widespread adoption in large scale open systems. More generally, the research community working on negotiation, argumentation, trust and semantics will have to convincingly show that their techniques are compatible with a wide range of application requirements, and provide concrete, measurable evidence of the benefits associated to such models and methodologies. Indeed, the community would benefit from a coherent set of benchmarks, pilot applications, and demonstrators providing such concrete evidence.

6.3 Deployment

Despite the promise of argumentation-based (cognitive or hybrid) approaches, these approaches have not been deployed in fully-fledged systems and applications. Future work is needed to tranfer these predominatly theoretical works into practice.

6.4 Quantifying benefits

Another important challenge amounts to quantifying the benefits brought about by agreement technologies in order to engage industry. The computations involved in argumentation techniques are more complex than other ad-hoc solutions (as, for example, the explicit conflict handling rules adopted in policy languages). In order to foster the adoption of

argumentation systems in real-world applications it is essential to address concrete use cases of industrial interest where the benefits of argumentation solutions cleary emerge (esp. from the perspective of adopters and end users) and can be directly compared with standard implementations. Similar exercises provide also useful feedback to research, e.g. by highlighting pragmatic issues (such as usability) that reduce the effectiveness of theoretically appealing techniques or make them hard to include in a particular application context or work environment.

6.5 Online Dispute Resolution

An important issue when attempting to deploy geographically distributed platforms to support applications is the fragmented nature or even absence of cross-border legislation [56]. This can result in legal disputes that are hard to mediate and may undermine the confidence of users in the fruitfulness of deploying these platforms.

Thus, one of the biggest challenges for electronic commerce is how to deal with disputes (especially cross-border ones).[5] Distance, different laws and jurisdiction are all potential obstacles to online business. Online Dispute Resolution (ODR), which is currently going through a process of regulation, is a clear mechanism for building trust in online transactions in situations where relationships are new or where there is a lack of efficient institutions that cover the transactions. The initial assumption that online markets would just need low prices has been shown to be wrong, as low risk transactions are equally relevant. ODR is also important for simple interpersonal dispute resolution as shown in [66] or as used in eBay[6]. See [43] for recent advances in the field. A preliminary solution to support online dispute resolution for contract-regulated interactions, using argumentation, has been proposed in the ARGUGRID[7] project [16].

6.6 Engaging industry

Trust plays a similar role in cloud computing as mentioned in section 2. Indeed, trust permits to assess the reliability of resources (e.g. PaaS, platform as a service, SaaS, software as a service or Iaas, Infrastructure as a service). Clouds require the signing of service level agreements (SLAs) in order for them to operate. One of the challenges is how to allow the interoperation among emerging commercial clouds (e.g. Amazon, Google) and non-commercial ones. There is room for semantics to help with this interoperability problem. Negotiation and argumentation techniques are going to be fundamental for the settlement of cloud computing agreements.

Some doubts have been shed on the willingness of industry to adopt the cloud computing paradigm:[8]

1. It is not secure.

[5] United Nations, E-Commerce and Development Report 2003, `http://www.unctad.org/ecommerce/2003`.

[6] `http://pages.ebay.com/services/buyandsell/disputeres.html`

[7] `www.argugrid.eu`

[8] `http://gigaom.com/2008/07/01/10-reasons-enterprises-arent-ready-to-trust-the-cloud/`

2. It cannot be logged
3. It is not platform agnostic.
4. Reliability (trust) is still an issue.
5. Portability is not seamless. Formats are not standard.
6. It is not environmentally sustainable. Outsource garbage (i.e. non-efficient computing) to third world countries.
7. Cloud computing still has to exist on physical servers. Patriotic act law applying to servers in the USA.
8. The need for speed still reigns at some firms.

Agreement technologies can help meet these challenges. Indeed, semantics techniques can help on issues 3 and 5 by supporting interoperability amongst platforms and by helping in the translation amongst formats. Trust can address security issues (1), on the selection of sites (7), and on provenance (6). Negotiation and argumentation can determine the caracteristic of SLAs addressing concerns like (8).

6.7 Benchmarking

There are several computational trust and reputation models that can be put into action and compared along several dimensions when facing realistic situations in open, distributed settings. Specific B2B operations (for example in the automotive, textile or oil industries [20], disaster rescue [8]) as well as particular social networks (for example for health care) could provide interesting scenaria for the sake of benchmarking of these models providing some kind of validation through industry-based scenaria.

7 Conclusions

Agreement technologies show promise in setting where autonomous agents can suitably support humans to identify mutually acceptable and profitable agreements. In this paper, we have considered four specific agreement technologies, namely trust, argumentation, negotiation and semantic alignment, and in particular the integration of the former with the latter three. The focus has been on open questions and challenges, building up on current achievements. We have looked at open challenges in general as well as in a specific class of applications, arising in large-scale, open distributed systems.

We have ignored other agreement technologies, notably the use of *norms* to regulate interactions and the possibility of agents to team up in *organisations*. Naturally, there are some interesting open challenges also concerning the integration of trust with these additional technologies. For example, trust may need to be applied in the context of norms [31], and trust is important in determining the composition of organisations during their formation and evolution [62, 33]. Moreover, we have focused on the integration of *trust* with other agreement technologies. However, it would be interesting to consider the integration of any agreement technologies with any other agreement technologies, e.g. argumentation and norms [21] and norms and organisations [1].

Acknowledgements

This wark was partially supported by the Agreement Technology COST action (IC0801). The authors would like to thank for helpful discussions and comments all participants in the panel on "Trust, Argumentation and Semantics" on 16 December 2009, Agia Napa, Cyprus.

References

[1] E. Argente, N. Criado, V. Julián, and V. J. Botti. Designing norms in virtual organizations. In T. Alsinet, J. Puyol-Gruart, and C. Torras, editors, *Artificial Intelligence Research and Development, Proceedings of the 11th International Conference of the Catalan Association for Artificial Intelligence, CCIA 2008, October 22-24, 2008, Sant Martí d'Empúries, Spain*, volume 184 of *Frontiers in Artificial Intelligence and Applications*, pages 16–23. IOS Press, 2008.

[2] S. Baselice, P. A. Bonatti, and M. Faella. On interoperable trust negotiation strategies. In *POLICY*, pages 39–50. IEEE Computer Society, 2007.

[3] P. Besnard and A. Hunter. *Elements of Argumentation*. MIT Press, 2008.

[4] P. Bonatti and P. Samarati. Logics for authorizations and security. In *Logics for Emerging Applications of Databases*, pages 277–323. Springer, 2003.

[5] P. A. Bonatti, J. L. D. Coi, D. Olmedilla, and L. Sauro. Rule-based policy representations and reasoning. In F. Bry and J. Maluszynski, editors, *REWERSE*, volume 5500 of *Lecture Notes in Computer Science*, pages 201–232. Springer, 2009.

[6] P. A. Bonatti, J. L. De Coi, D. Olmedilla, and L. Sauro. A rule-based trust negotiation system. *IEEE Trans. on Knowledge and Data Engineering*, 22(11):1507–1520, 2010.

[7] P. A. Bonatti and D. Olmedilla. Driving and monitoring provisional trust negotiation with metapolicies. In *POLICY*, pages 14–23. IEEE Computer Society, 2005.

[8] L. Camarinha-Matos, I. S. I., H. Afsarmanesh, and A. Oliveira. Towards a framework for creation of dynamic virtual organizations. In L. Camarinha-Matos, H. Afsarmanesh, and A. Ortiz, editors, *Collaborative Networks and their Breeding Environments*, pages 69–80. Springer Berlin, 2005.

[9] H. L. Cardoso and E. Oliveira. Virtual enterprise normative framework within electronic institutions. In A. O. M.P. Gleizes and F. Zambonelli, editors, *Engineering Societies in the Agents World V*, pages 14–32, 2005.

[10] S. J. Casare and J. S. Sichman. Towards a functional ontology of reputation. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi Agent Systems*, pages 505–511, 2005.

[11] C. Castelfranchi and R. Falcone. Trust is much more than subjective probability: Mental components and sources of trust. In *HICSS*, 2000.

[12] N. Choi, I.-Y. Song, and H. Han. A survey on ontology mapping. *SIGMOD Rec.*, 35(3):34–41, 2006.

[13] R. Conte and M. Paolucci. *Reputation in artificial societies: Social beliefs for social order*. Kluwer Academic Publishers, 2002.

[14] J. Debenham and C. Sierra. Trust and honour in information-based agency. pages 1225–1232, 2006.

[15] P. Dondio and S. Barrett. Presumptive selection of trust evidences. In *AAMAS*, 2007.

[16] P. M. Dung and P. M. Thang. Modular argumentation for modelling legal doctrines in common law of contract. *Artif. Intell. Law*, 17(3):167–182, 2009.

[17] A. Fabregues and C. Sierra. A testbed for multiagent systems. Technical report, IIIA-CSIC, Bellaterra, Barcelona, 10/2009 2009.

[18] K. K. Fullam, T. B. Klos, G. Muller, J. Sabater, A. Schlosser, Z. Topol, K. S. Barber, J. S. Rosenschein, L. Vercouter, and M. Voss. A specification of the agent reputation and trust (art) testbed: Experimentation and competition for trust in agent societies. In *4th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2005)*, pages 512–518, 2005.

[19] L. Gong and G. Ellison. *Inside Java(TM) 2 Platform Security: Architecture, API Design, and Implementation, 2nd edition*. Pearson Education, 2003.

[20] M. Grabowski and K. Roberts. Risk mitigation in virtual organizations. *Organization Science*, 10(6):704–721, 1999.

[21] S. Heras, N. Criado, E. Argente, and V. Julian. Norm emergency through argumentation. *Journal of Physical Agents*, 3:31–38, 2009.

[22] A. Hogan, A. Harth, and A. Polleres. Scalable authoritative owl reasoning for the web. *Int. J. Semantic Web Inf. Syst.*, 5(2):49–90, 2009.

[23] T. Huynh, N. Jennings, and N. Shadbolt. An integrated trust and reputation model for open multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 13(2):119–154, 2006.

[24] S. Jajodia, P. Samarati, M. L. Sapino, and V. S. Subrahmanian. Flexible support for multiple access control policies. *ACM Trans. Database Syst.*, 26(2):214–260, 2001.

[25] M. Karlins and H. I. Abelson. *Persuasion*. Crosby Lockwood & Son, London, UK, 1970.

[26] A. Koster, J. Sabater-Mir, and M. Schorlemmer. Inductively generated trust alignments based on shared interactions. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi Agent Systems*, pages in–press, 2010.

[27] P. Krause, S. Ambler, M. Elvang-Gøransson, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Comp. Int.*, 11:113–131, 1995.

[28] L. Lacomme, Y. Demazeau, and V. Camps. Personalization of a trust network. In *Proceedings of the 1st International Conference on Agents and Artificial Intelligence (ICAART'09)*, pages 408–415, 2009.

[29] A. J. Lee and M. Winslett. Virtual fingerprinting as a foundation for reputation in open systems. In K. Stølen, W. H. Winsborough, F. Martinelli, and F. Massacci, editors, *iTrust*, volume 3986 of *Lecture Notes in Computer Science*, pages 236–251. Springer, 2006.

[30] G. Lenzini, N. Sahli, and H. Eertink. Agents selecting trustworthy recommendations in mobile virtual communities. In *AAMAS-TRUST*, volume 5396 of *LNCS*. Springer, 2008.

[31] M. Luck, S. Munroe, R. Ashri, and F. L. y López. Trust and norms for interaction. In *Proceedings of the IEEE International Conference on Systems, Man & Cybernetics: The Hague, Netherlands, 10-13 October 2004*, pages 1944–1949. IEEE, 2004.

[32] P.-A. Matt, M. Morge, and F. Toni. Combining statistics and arguments to compute trust. In W. van der Hoek and G. A. Kaminka, editors, *Proceedings of the 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, 2010.

[33] J. McGinnis, K. Stathis, and F. Toni. A formal framework of virtual organisations as agent societies. *EPTCS*, 16:1, 2010.

[34] D. Melaye and Y. Demazeau. Bayesian dynamic trust model. In P. P. M. PechouceK and L. Varga, editors, *Lecture Notes in Artificial Intelligence*, pages 480–489. Springer Berlin, 2005.

[35] L. Mui, M. Mohtashemi, and A. Halberstadt. A computational model for trust and reputation. In *Proceedings of the 35th Hawaii International Conference on System Sciences*, 2002.

[36] N. F. Noy. Semantic integration: a survey of ontology-based approaches. *SIGMOD Rec.*, 33(4):65–70, 2004.

[37] S. Ossowski. Coordination and agreement in multi-agent systems. In M. Klusch, M. Pechoucek, and A. Polleres, editors, *Cooperative Information Agents XII, 12th International Workshop, CIA 2008, Prague, Czech Republic, September 10-12, 2008. Proceedings*, volume 5180 of *Lecture Notes in Computer Science*, pages 16–23. Springer, 2008.

[38] S. Ossowski. Coordination in multi-agent systems: Towards a technology of agreement. In R. Bergmann, G. Lindemann, S. Kirn, and M. Pechoucek, editors, *Multiagent System Technologies, 6th German Conference, MATES 2008, Kaiserslautern, Germany, September 23-26, 2008. Proceedings*, volume 5244 of *Lecture Notes in Computer Science*, pages 2–12. Springer, 2008.

[39] C. Paliouras, V. Karkaletsis, C. Papatheodorou, and C. D. Pyropoulos. Exploiting learning techniques for the acquisition of user stereotypes and communities. In *Proceedings of the Seventh International Conference on User Modeling*, 1999.

[40] S. Parsons, P. McBurney, and E. Sklar. Reasoning about trust using argumentation: A position paper. In *Proceedings of the Seventh International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2010), affiliated to AAMAS 2010*, 2010.

[41] I. Pinyol and J. Sabater-Mir. Towards the definition of an argumentation framework using reputation information. In *The 12th Workshop on Trust in Agent Societies (TRUST@AAMAS'09)*, 2009.

[42] I. Pinyol, J. Sabater-Mir, and G. Cuni. How to talk about reputation using a common ontology: From definition to implementation. In *Ninth Workshop on Trust in Agent Societies*, pages 90–102, May 2007.

[43] M. Poblet, editor. *Proceedings of the 5th International Workshop on Online Dispute Resolution, in conjunction with the 21st International Conference on Legal Knowledge and Information Systems (JURIX 2008), Firenze, Italy, December 13, 2008*, volume 430 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2008.

[44] H. Prade. A qualitative bipolar argumentative view of trust. In *SUM*, 2007.

[45] I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. *The Knowledge Engineering Review*, 18(04):343–375, 2003.

[46] I. Rahwan and G. Simari, editors. *Argumentation in AI: The Book*. Springer, 2009.

[47] K. Regan and R. Cohen. Indirect reputation assessment for adaptive buying agents in electronic markets. *Business Agents and the Semantic Web workshop*, 1, 2005.

[48] J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artif. Intell. Rev*, 24(1):33–60, 2005.

[49] N. M. Sadeh, J. I. Hong, L. F. Cranor, I. Fette, P. G. Kelley, M. K. Prabaker, and J. Rao. Understanding and capturing people's privacy policies in a mobile social networking application. *Personal and Ubiquitous Computing*, 13(6):401–412, 2009.

[50] C. Sierra and J. Debenham. The logic negotiation model. In *Sixth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1026–1033, 2007.

[51] C. Sierra and J. Debenham. Information-based argumentation. pages 155–170, Sydney, 2008.

[52] C. Sierra and J. Debenham. Information-based reputation. In *First International Conference on Reputation: Theory and Technology*, 2009.

[53] C. Sierra and J. Sabater-Mir. Review on computionational trust and reputation models. *Artificial Intelligence Review*, 24(1):33 – 60, 2005.

[54] E. Staab and T. Engel. Combining cognitive with computational trust reasoning. In *AAMAS-TRUST*, volume 5396 of *LNCS*. Springer, 2008.

[55] S. Staab, B. K. Bhargava, L. Lilien, A. Rosenthal, M. Winslett, M. Sloman, T. S. Dillon, E. Chang, F. K. Hussain, W. Nejdl, D. Olmedilla, and V. Kashyap. The pudding of trust. *IEEE Intelligent Systems*, 19(5):74–88, 2004.

[56] K. Stanoevska, D. Parrilli, and G. Thanos. BEinGRID: Development of business models for the grid industry. volume 5206 of *LNCS*. Springer, 2008.

[57] M. Tavakolifard, P. Herrmann, and P. Ozturk. Analogical trust reasoning. In *Trust Management III*, pages 149–163. Springer Boston, 2009.

[58] J. Urbano, A. Rocha, and E. Oliveira. Trustworthiness tendency incremental extraction using information gain. Technical report, University of Oporto, Department of Computer Science, 2010.

[59] J. Urbano, A. P. Rocha, and E. Oliveira. Computing confidence values: Does trust dynamics matter? In L. S. Lopes, P. M. N. Lau, and L. M. Rocha, editors, *Lecture Notes in Artificial Intelligence*, pages 520–531. Springer Berlin, 2009.

[60] J. Urbano, A. P. Rocha, and E. Oliveira. Trust estimation using contextual fitness. In *Proceedings of 4th International KES Symposium on Agents and Multi-agent Systems, Technologies and Applications*, 2010.

[61] M. Winslett, T. Yu, K. E. Seamons, A. Hess, J. Jacobson, R. Jarvis, B. Smith, and L. Yu. Negotiating trust on the web. *IEEE Internet Computing*, 6(6):30–37, 2002.

[62] J. Young. Trust in virtual organisations: a synthesis of the literature. *IJNVO*, 5(3/4):244–258, 2008.

[63] B. Yu and M. P. Singh. Distributed reputation management for electronic commerce. *Computational Intelligence*, 18(4):535–549, 2002.

[64] T. Yu, M. Winslett, and K. Seamons. Interoperable strategies in automated trust negotiation. In *CCS '01: Proceedings of the 8th ACM conference on Computer and Communications Security*, pages 146–155. ACM Press, 2001.

[65] T. Yu, M. Winslett, and K. Seamons. Supporting Structured Credentials and Sensitive Policies through Interoperable Strategies in Automated Trust Negotiation. *ACM Transactions on Information and System Security*, 6(1), Feb. 2003.

[66] J. Zeleznikow. Beyond interest based bargaining - incorporating interests and fairness in the development of negotiation support systems. In Poblet [43].