

Some thoughts on using argumentation to handle trust

Simon Parsons^{1,2}, Yuqing Tang², Kai Cai², Elizabeth Sklar^{1,2}, and Peter McBurney³

¹ Department of Computer & Information Science, Brooklyn College,
City University of New York, 2900 Bedford Avenue, Brooklyn, NY 11210 USA
{sklar, parsons}@sci.brooklyn.cuny.edu

² Department of Computer Science, Graduate Center
City University of New York, 365 Fifth Avenue, New York, NY 10016, USA
{ytang, kcai}@gc.cuny.edu

³ Department of Informatics, Kings College London,
Strand Building, peter.mcburney@kcl.ac.uk

Abstract. This paper describes some of our recent work on using argumentation to handle information about trust. We first discuss the importance of trust in computer science in general and in multi-agent systems in particular. We then describe the setting of our work, situating it within the broad area of work on trust. Next we provide an overview of two lines of work we are currently pursuing — using argumentation to reason about which individuals to trust, and using argumentation to relate sources of information to conclusions drawn from information provided by those sources. Finally, we outline our current initiatives and briefly highlight other work that is closely related to ours.

1 Why trust is important

Trust is a mechanism for managing the uncertainty about autonomous entities and the information they deal with. As a result, trust can play an important role in any decentralized system. As computer systems have become increasingly distributed, and control in those systems has become more decentralized, trust has become an increasingly more important concept in computer science [3, 12].

Much of the work on trust in computer science has concentrated on dealing with specific scenarios in which trust has to be established or handled in some fashion. Thus, we see work on trust in peer-to-peer networks, including the EigenTrust algorithm [19] — a variant of PageRank [27] where downloads from a source play the same role as outgoing hyperlinks and which is effective in excluding peers who want to disrupt the network. [1], also in the area of peer-to-peer networks, develops a mechanism that prevents peers manipulating their trust values to get preferential downloads. [44] is concerned with slightly different issues in mobile ad-hoc networks, looking to prevent nodes from getting others to transmit their messages while refusing to transmit the messages of others.

The internet, as the largest distributed system of all, is naturally a target of much of the research on trust. There have, for example, been studies on the development of trust in ecommerce through the use of reputation systems [34] and studies on how such systems perform [33, 39] and how such systems can be manipulated [21]. Another area

of concern has to do with the reliability of sources of information on the web. [43], for example, investigates mechanisms to determine which sources to trust when faced with multiple conflicting sources, while [4] looks at the related question of how to resolve conflicting information, and [2] extends this idea to rate the individuals who provide information by looking at the history of the information they have provided. Issues related to trust in the social web have also attracted much attention [11, 22, 39, 42].

Trust is an especially important issue from the perspective of autonomous agents and multiagent systems. The premise behind the multiagent systems field is that of developing software agents that will work in the interests of their owners, carrying out their owners' wishes while interacting with other entities. In such interactions, agents will have to reason about the amount that they should trust those other entities, whether they are trusting those entities to carry out some task, or whether they are trusting those entities to not misuse crucial information. As a result we find much work on trust in agent-based systems [35], including the influential model proposed by [6].

We are studying the use of formal systems of argumentation for handling trust. In this paper we present a brief, and largely informal, overview of our work to date, with pointers to the detailed treatment of all the topics we introduce.

2 The setting for our work

Our work considers an agent Ag which is part of a society of agents Ags . Each agent $Ag_i \in Ags$ has a knowledge base Σ_i that contains the information that the agent has about the world. The agents communicate, and so each Ag_i can make use of not only the information that is in its Σ_i , but also the information that comes from other agents. We model this situation by considering that each Ag_i has a *commitment store* CS_i which contains information that that agent has communicated. Since at the moment we are only concerned with the perspective of the one agent Ag , we only consider the commitment stores of agents other than Ag to hold information that the respective agents have communicated to Ag . That is we ignore issues around how the information was communicated, who else knows the same information, and so on. We just consider all information that some agent Ag_i has communicated to Ag to be contained in CS_i .

We assume that all agents that have communicated information to Ag are members of Ag 's social network, and so it is possible to construct a graph which relates Ag to all these agents. We further assume that it is possible to attach a numerical measure to each link in this social network to quantify the extent to which an agent trusts those to which it is linked in the social network. In common with the literature on trust in social networks, we call such a structure a *trust network*. Figure 1 is a fragment of the trust network for an agent $john$ identifying $john$'s relationships with the agents $dave$, $mary$, $alice$ and $jane$. (This social network is taken from an example in [20], which itself is drawn from the FilmTrust network [7, 11], with the trust values normalized between 0 and 1.)

Our work has concentrated on establishing mechanisms by which our featured agent Ag , $john$ in the case of Figure 1, can use the information it obtains from its acquaintances in a way that reflects its trust in them. Thus, there is an assumption that if some




Fig. 1. A simple trust network

proposition $p \in CS_j$, then Ag will accord that proposition the belief $bel(p)$ where:

$$bel(p) = ttb(tr(Ag, Ag_j)) \quad (1)$$

That is, $bel(p)$ is some function of the trust that Ag has in Ag_j regarding its belief in p . Thus, in the example of Figure 1, we consider that if $john$ is told p by $jane$, we assume that $john$'s belief in this is a function of his trust in $jane$.

Thinking a little about $john$'s relationship to $jane$ raises some issues, one directly, and others that follow from that first issue. The first issue is that according to Figure 1, $john$ doesn't know $jane$. $john$ only knows $mary$, and so only directly states his trust in $mary$. However, it is common in the literature of trust in social networks to assume that trust propagates (or, to describe it in another way, is transitive) so that since $john$ trusts $mary$ and $mary$ trusts $jane$, then $john$ can trust $jane$. The assumption of transitivity (and it is a big assumption, as discussed in [8]) then raises additional issues.

The first of these is the context in which $john$ trusts $jane$. As many authors have pointed out, not least [8], trust is highly context dependent. $john$ may trust $jane$ in some domains, but not in others. To use a common example, $john$ may trust $jane$, a car mechanic, for information on cars and how to fix them. However, just because $john$ trusts $jane$ about cars does not mean that he will trust her to recommend a restaurant or someone to babysit his daughter. Indeed, as [17] points out, the semantics of the links between $john$ and $mary$ and between $john$ and $jane$ are rather different. Consider that $john$ is, as in the example from which Figure 1 is taken in [20], looking for recommendations about which film to watch. If $john$ solicits information from his friend $mary$, he is trusting her to recommend films he would like to watch. However, if he accepts an indirect recommendation from $jane$, whom he trusts because $mary$ trusts her, he is relying on $jane$ being a good recommender of films but is relying on $mary$ being a good recommender of film recommenders. In other words, $john$ is not trusting $mary$'s film knowledge in this second case, but her knowledge of people who can recommend films. [17] distinguishes between *functional* trust, the trust in an individual to carry out some task (such as recommending a film), and *referral* trust, the trust in an individual's recommendation of another individual. In our work, we ignore these dis-

tinctions, assuming that the trust networks capture a context in which transitivity holds. (We can easily imagine an agent having different trust networks for different contexts, each composed of functional and referral trust links.)

3 How we use argumentation in handling trust

In this section we describe how our work makes use of argumentation to handle trust.

3.1 How argumentation can help

Some models of trust, for example the influential model from [6], start from the position that trust in an individual is only ever important when the trusting party needs that individual to perform an action. While it is clear that in such cases trust is important — if we are going to construct a plan where a critical action is to be performed by individual I then we certainly need to trust I to do what we need them to — we, like [23] for example, believe that trust in individuals also has a role to play when those individuals provide us with information.

In particular, we believe that knowing the sources of the information and the way in which it is used, that is the *provenance* of the information, is important when the sources of the information may not be completely trustworthy. As [26] points out,

By knowing the provenance of data, users can often better understand, trust, reproduce, and validate it.

The relationship between trust and provenance has been explored by a number of authors, for example [10, 32, 40, 41]. As we argued in [29], since argumentation records the data from which conclusions are drawn, it provides a natural mechanism to capture provenance. This is especially the case when the argument records not only the data from which conclusions are drawn, but also the full derivation of those conclusions (as, for example, in [9, 31]).

3.2 Our contribution so far

We are working on two related lines of research. In both, we use argumentation, as described above, to make explicit the connection between sources of data and the conclusions derived from that data. In one line of research, we use argumentation to carry out reasoning about the trustworthiness of sources — that is to compute trust propagation — with the aim of permitting reasoning about the validity of different forms of propagation. In the other line of research, we use a graphical representation of arguments as a means of communicating provenance information, as will be explained below.

In using argumentation to compute trust propagation [30], we assume that each agent Ag_i has a knowledge base, $\Delta_i^{tr} \subseteq \Sigma_i$, containing information about who trusts whom. Table 1 contains Δ_{john}^{tr} , the knowledge base for *john*, constructed from the example in Figure 1. Each element of Δ_{john}^{tr} has the form:

$$(\langle index \rangle : \langle data \rangle : \langle value \rangle)$$

The first is a means of referring to the element, the second is a formula, and the third is the degree of trust between the individuals.

Table 1. Knowledge base containing *john*'s beliefs

$$\begin{aligned} \Delta_{john}^{tr} & (t1 : trusts(john, mary) : 0.9) \\ & (t2 : trusts(mary, jane) : 0.7) \\ & (t3 : trusts(mary, dave) : 0.8) \\ & (t4 : trusts(alice, jane) : 0.6) \\ & (t5 : trusts(alice, paul) : 0.4) \end{aligned}$$

Arguments can then be constructed from Δ_{john}^{tr} using the rules in Figure 2⁴. For example, using the first two rules, from Figure 2, Ax^{tr} and dp , it is possible to construct the argument:

$$\Delta_{john}^{tr} \vdash_{tr} (trusts(john, jane) : \{t1, t2\} : \{Ax^{tr}, Ax^{tr}, dp\} : \tilde{t})$$

where all arguments in our approach take the form:

$$(\langle conclusion \rangle : \langle grounds \rangle : \langle rules \rangle : \langle value \rangle)$$

The $\langle conclusion \rangle$ is inferred from the $\langle grounds \rangle$ using the rules of inference $\langle rules \rangle$ and with degree $\langle value \rangle$. In this case the argument says *john* trusts *jane* with degree \tilde{t} (which is $0.9 \otimes^{tr} 0.7$), through two applications of the rule Ax^{tr} and one application of the rule dp to the two facts indexed by $t1$ and $t2$ ⁵. Using just Ax^{tr} and dp captures the transitivity of trust, and the crucial inference step is what [13] calls “direct propagation” (hence the name of the rule dp). Figure 3a shows the result of establishing all the indirect links possible using direct propagation where \otimes^{tr} is taken to be \min , as we do in [30].

The reason for constructing arguments about trust is primarily so that it is possible to tell on what basis the conclusion to trust a particular source has been drawn. We do this because there are a variety of reasons that one might have for trusting a source, and it may be necessary to identify which reasons(s) have been used in a particular case in order to be able to dispute or defend them. Making the reasons explicit, as our approach does, unleashes this possibility. As an example of an alternative to transitivity as a form of trust propagation, consider the rule cc from Figure 2. This captures a form of reasoning that [13] calls *co-citation*. [13] describes this as:

For example, suppose i_1 trusts j_1 and j_2 , and i_2 trusts j_2 . Under co-citation, we would conclude that i_2 should also trust j_1 .

⁴ Note that the consequence relation in Figure 2 is not intended to be comprehensive. There are many other ways to construct arguments about trust — for some examples see [28] — which could be included in the definition of \vdash_{tr} .

⁵ As mentioned above, there are good reasons for using the formulae themselves in the grounds and factoring the whole proof into the set of rules (as we do in [29]) to obtain structured arguments like those in [9, 31]. However, for simplicity, here we use the relevant indices.

$$Ax^{tr} \frac{(n : trusts(x, y) : \tilde{d}) \in \Delta_i^{tr}}{\Delta_i^{tr} \vdash_{tr} (trusts(x, y) : \{n\} : \{Ax^{tr}\} : \tilde{d})}$$

$$dp \frac{\Delta_i^{tr} \vdash_{tr} (trusts(x, y) : G : R : \tilde{d}) \text{ and } \Delta_i^{tr} \vdash_{tr} (trusts(y, z) : H : S : \tilde{e})}{\Delta_i^{tr} \vdash_{tr} (trusts(x, z) : G \cup H : R \cup S \cup \{dp\} : \tilde{d} \otimes^{tr} \tilde{e})}$$

$$cc \frac{\Delta_i^{tr} \vdash_{tr} (trusts(x, y) : G : R : \tilde{d}) \text{ and } \Delta_i^{tr} \vdash_{tr} (trusts(x, z) : H : S : \tilde{e}) \text{ and } \Delta_i^{tr} \vdash_{tr} (trusts(w, z) : K : T : \tilde{f})}{\Delta_i^{tr} \vdash_{tr} (trusts(w, y) : G \cup H \cup K : R \cup S \cup T \cup \{cc\} : \tilde{d} \otimes^{tr} \tilde{e} \otimes^{tr} \tilde{f})}$$

Fig. 2. Some rules for propagating trust

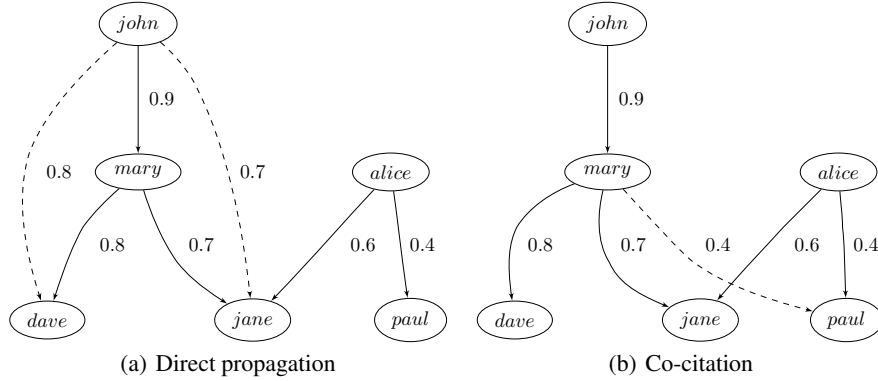


Fig. 3. Two types of trust propagation — in each figure the dashed lines show the indirect trust relationships that are inferred.

In our example (see Figures 1 and 3b), therefore, co-citation suggests that since *alice* trusts *jane* and *paul*, and *mary* trusts *jane*, then *mary* should trust *paul*. The idea is that since *alice* and *mary* agree on the trustworthiness of *jane*, *mary* should trust *alice*'s opinion about *paul*. [13] also tells us how trust values should be combined in this case — *mary*'s trust in *paul* is just the combination of trust values along the path from *mary* to *jane* to *alice* to *paul*. Figure 3b shows the direct trust link that co-citation implies in our example where the values are again combined using min, with *mary* trusting *paul* to degree 0.4.

Combining the application of *cc* with the use of *dp*, as above, allows the construction of the argument:

$$\Delta_{john}^{tr} \vdash_{tr} (trusts(john, paul) : \{t1, t2, t4, t5\} : rules_1 : \tilde{r})$$

indicating that *john* trusts *paul* (based on *john*'s knowledge base as in Table 1), where $rules_1$ is:

$$\{Ax^{tr}, Ax^{tr}, Ax^{tr}, Ax^{tr}, cc, dp\}$$

and \tilde{r} is $0.9 \otimes^{tr} 0.7 \otimes^{tr} 0.6 \otimes^{tr} 0.4$ (which also comes to 0.4 when \otimes^{tr} is min).

Once these trust values have been established, *john* can then apply Equation 1 to establish a degree of belief for anything the other agents in the trust network tell him, and then combine this information with what he already knows. [30] describes formally how this can be done.

The description above concerns the first of the two lines of work we are pursuing: the use of argumentation to capture trust propagation. Our second line of work employs a graphical model of arguments as the basis of communicating provenance information. The model we have developed [38, 37] is best explained through the use of an instance of a graphical argument of the kind generated by the model. Such a graph is given in Figure 4. In short, the graph contains three components. First, there is a trust graph. In this case, it is a subset of the trust graph from Figure 1 — exactly that bit of the graph from Figure 1 which contains the agents from whom *john* can infer trust using

direct propagation. Second, there are arguments. These are proof trees where the conclusions of the trees are formulae of interest, and each formula that is not an inference is linked to the agent that supplies the information. In this case, the arguments — taken from the example in [20] — concern whether or not to watch the Pedro Almodovar film “Hable con ella” (abbreviated “hce” in Figure 4). Again Equation 1 tells us how to establish belief values for information that comes from different agents in a way that depends on trust in the source agent(s). Third, there are arcs that identify conflicts between arguments. In Figure 4, these are just between the conclusions of arguments, but more general conflicts are also identified in the full model [37]. Together these three components make up a *trust-extended argumentation graph*.

The full model [37] not only includes a formal description of each of these components, but also translates standard approaches for evaluating arguments [5] into criteria that can be evaluated on the graph. In addition, we have identified algorithms for building the graph and then evaluating the arguments that it contains.

3.3 Current and future work

Our long term aim is to build tools that help users to reason with information from sources of varying trustworthiness. Our hypothesis is that the graphical argumentation model we illustrated above is a useful way to do this. Now that we have a formal model which makes it possible to automate reasoning and link sources to conclusions, we are working towards testing this hypothesis. To do this, we need an interactive tool, and implementing such a tool is one of our current foci. One important aspect of the implementation is the interface. While we believe that it is useful to show users how sources of information relate to conclusions, we do not think that users will respond well to complete graphs of the kind in Figure 4 — there is too much information for this to be easily digested by the average user. Instead we are looking at ways of allowing users to navigate sections of the graph, zooming in on areas of interest. To help us understand the best way to do this, we are running user studies where human subjects develop argument graphs for some simple problem scenarios.

4 Conclusions

We have described our work on using argumentation to handle trust. We focus on the situation in which some entity uses information that comes from sources of varying trustworthiness, and look at how argumentation can be used to capture the provenance of the information used to derive answers to queries of interest. There are two main aspects of our work to date. One is the use of argumentation as a means of capturing different mechanisms for propagating trust. The other is the generation of a graphical model that can be used to communicate provenance information to users. Our current work is concentrating on implementing the second model and experimentally determining how best to present results to users.

Finally we should note that though the particular combinations of argumentation and trust that we are studying are novel, the idea of combining trust and argumentation is not. Four lines of work on trust and argumentation that are complementary to ours

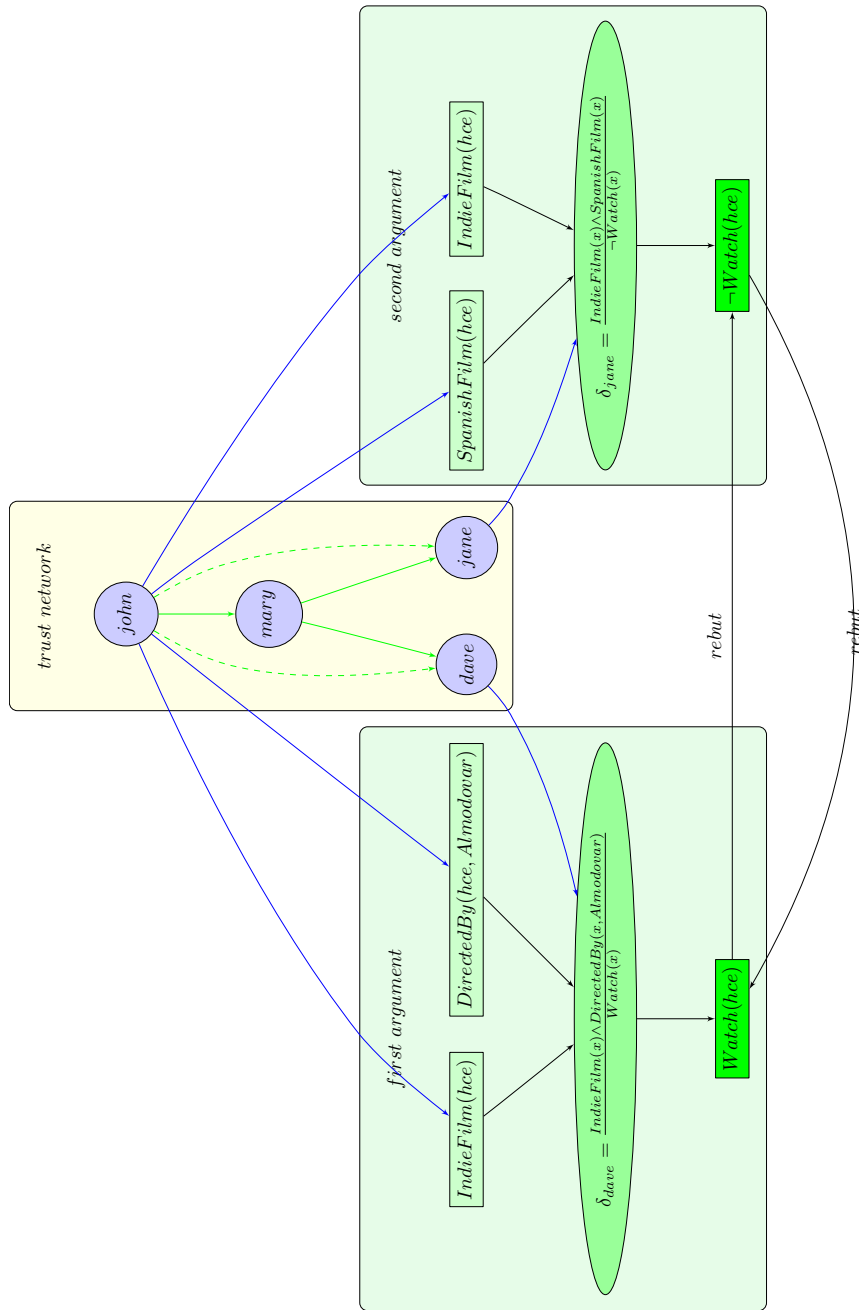


Fig. 4. An example of a trust-extended argumentation graph

are those of Harwood [14, 15], Matt *et al.*, Hunter [16], [24], and Stranders [36]. For a detailed comparison of this work and ours, see [37]. In addition, argumentation has been used in the past to reason about risk [18, 25], a subject closely related to trust; though the cited work looks at risk of carcinogenicity given chemical structure rather than risk due to untrustworthiness.

Acknowledgments

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

1. Z. Abrams, R. McGrew, and S. Plotkin. Keeping peers honest in EigenTrust. In *Proceedings of the 2nd Workshop on the Economics of Peer-to-Peer Systems*, 2004.
2. B. T. Adler and L. de Alfaro. A content-driven reputation system for the Wikipedia. In *Proceedings of the 16th International World Wide Web Conference*, Banff, Alberta, May 2007.
3. D. Artz and Y. Gil. A survey of trust in computer science and the semantic web. *Journal of Web Semantics*, 5(2):58–71, June 2007.
4. X. L. Dong, L. Berti-Equille, and D. Srivastava. Integrating conflicting data: The role of source dependence. In *Proceedings of the 35th International Conference on Very Large Databases*, Lyon, France, August 2009.
5. P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n -person games. *Artificial Intelligence*, 77:321–357, 1995.
6. R. Falcone and C. Castelfranchi. Social trust: A cognitive approach. In C. Castelfranchi and Y. Tan, editors, *Trust and Deception in Virtual Societies*, pages 55–99. Kluwer Academic Publishers, 2001.
7. <http://trust.mindswap.org/FilmTrust/>.
8. R. Francone and C. Castelfranchi. Transitivity in trust: A discussed property. In *Proceedings of the Undicesimo Workshop Nazionale "Dagli Oggetti agli Agenti"*, Rimini, September 2010.
9. A. J. García and G. Simari. Defeasible logic programming: an argumentative approach. *Theory and Practice of Logic Programming*, 4(1):95–138, 2004.
10. J. Golbeck. Combining provenance with trust in social networks for semantic web content filtering. In *Proceedings of the International Provenance and Annotation Workshop*, Chicago, Illinois, May 2006.
11. J. Golbeck and J. Hendler. Filmtrust: Movie recommendations using trust in web-based social networks. In *Proceedings of the IEEE Consumer Communications and Networking Conference*, 2006.
12. T. Grandison and M. Sloman. A survey of trust in internet applications. *IEEE Communications Surveys and Tutorials*, 4(4):2–16, 2000.

13. R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *Proceedings of the 13th International Conference on the World Wide Web*, 2004.
14. W. T. Harwood, J. A. Clark, and J. L. Jacob. Networks of trust and distrust: Towards logical reputation systems. In D. M. Gabbay and L. van der Torre, editors, *Logics in Security*, Copenhagen, Denmark, 2010.
15. W. T. Harwood, J. A. Clark, and J. L. Jacob. A perspective on trust, security and autonomous systems. In *Proceedings of the New Security Paradigms Workshop*, Concord, MA, 2010.
16. A. Hunter. Reasoning about the appropriateness of propositors for arguments. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, Chicago, Illinois, July 2008.
17. A. Jøsang, E. Gray, and M. Kinatader. Simplification and analysis of transitive trust networks. *Web Intelligence and Agent Systems*, 4(2):139–161, 2006.
18. P. N. Judson, J. Fox, and P. J. Krause. Using new reasoning technology in chemical information systems. *Journal of Chemical Information and Computer Sciences*, 36:621–624, 1996.
19. S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The EigenTrust algorithm for reputation management in P2P networks. In *Proceedings of the 12th World Wide Web Conference*, May 2004.
20. Y. Katz and J. Golbeck. Social network-based trust in prioritized default logic. In *Proceedings of the 21st National Conference on Artificial Intelligence*, 2006.
21. J. Lang, M. Spear, and S. F. Wu. Social manipulation of online recommender systems. In *Proceedings of the 2nd International Conference on Social Informatics*, Laxenburg, Austria, 2010.
22. K. Lerman and A. Galstyan. Analysis of social voting patterns on Digg. In *Proceedings of the 1st Workshop on Online Social Networks*, Seattle, August 2008.
23. C.-J. Liau. Belief, information acquisition, and trust in multi-agent systems — a modal logic formulation. *Artificial Intelligence*, 149:31–60, 2003.
24. P.-A. Matt, M. Morge, and F. Toni. Combining statistics and arguments to compute trust. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagents Systems*, Toronto, Canada, May 2010.
25. P. McBurney and S. Parsons. Dialectical argumentation for reasoning about chemical carcinogenicity. *Logic Journal of the IGPL*, 9(2):191–203, 2001.
26. S. Miles, P. Groth, S. Munroe, and L. Moreau. PrIME: A methodology for developing provenance-aware applications. *ACM Transactions on Software Engineering and Methodology*, (to appear), 2010.
27. L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank citation ranking: Bringing order to the Web. Technical Report 1999-66, Stanford InfoLab, 1999.
28. S. Parsons, K. Haigh, K. Levitt, J. Rowe, M. Singh, and E. Sklar. Arguments about trust. Technical report, Collaborative Technology Alliance, 2011.
29. S. Parsons, P. McBurney, and E. Sklar. Reasoning about trust using argumentation: A position paper. In *Proceedings of the Workshop on Argumentation in Multiagent Systems*, Toronto, Canada, May 2010.
30. S. Parsons, E. Sklar, and P. McBurney. Using argumentation to reason with and about trust. In *Proceedings of the 8th International Workshop on Argumentation in Multiagent Systems*, Taipei, Taiwan, 2011.
31. H. Prakken. An abstract framework for argumentation with structured arguments. *Argument and Computation*, 1:93–124, 2010.
32. S. Rajbhandari, I. Wootten, A. Ali, and O. F. Rana. Evaluating provenance-based trust for scientific workflows. In *Proceedings of the Sixth IEEE International Symposium on Cluster Computing and the Grid*, pages 365–372, Singapore, 2006. IEEE Computer Society Press.

33. P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system. In M. R. Baye, editor, *The Economics of the Internet and E-Commerce*, pages 127–157. Elsevier Science, Amsterdam, 2002.
34. P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems: Facilitating trust in internet interactions. *Communications of the ACM*, 43:45–48, 2000.
35. J. Sabater and C. Sierra. Review on computational trust and reputation models. *AI Review*, 23(1):33–60, September 2005.
36. R. Stranders, M. de Weerd, and C. Witteveen. Fuzzy argumentation for trust. In F. Sadri and K. Satoh, editors, *Proceedings of the Eighth Workshop on Computational Logic in Multi-Agent Systems*, volume 5056 of *Lecture Notes in Computer Science*, pages 214–230. Springer Verlag, 2008.
37. Y. Tang, K. Cai, P. McBurney, E. Sklar, and S. Parsons. Using argumentation to reason about trust and belief. *Journal of Logic and Computation*, (to appear), 2011.
38. Y. Tang, K. Cai, E. Sklar, P. McBurney, and S. Parsons. A system of argumentation for reasoning about trust. In *Proceedings of the 8th European Workshop on Multi-Agent Systems*, Paris, France, December 2010.
39. C-Y Teng, D. Lauterbach, and L. Adamic. I rate you. You rate me. Should we do so publicly? In *Proceedings of the 3rd Workshop on Online Social Networks*, Boston, June 2010.
40. P. Victor, C. Cornelis, M. De Cock, and P. Pinheiro da Silva. Towards a provenance-preserving trust model in agent networks. In *Proceedings of the Workshop on Models of Trust for the Web*, 2006.
41. X. Wang, K. Govindan, and P. Mohapatra. Provenance-based information trustworthiness evaluation in multihop networks. In *Proceedings of the 53rd Annual IEEE Global Communications Conference*, Miami, FL, December 2010.
42. S. Ye and S. F. Wu. Measuring message propagation and social influence on Twitter.com. In *Proceedings of the 2nd International Conference on Social Informatics*, Laxenburg, Austria, 2010.
43. X. Yin, J. Han, and P. S. Yu. Truth discovery with multiple conflicting information providers on the web. In *Proceedings of the Conference on Knowledge and Data Discovery*, 2007.
44. S. Zhong, J. Chen, and Y. R. Yang. Sprite: A simple cheat-proof, credit-based system for mobile ad-hoc networks. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies*, 2003.