# Trust Analysis Through Relationship Identification

**Ronald Ashri**[1] and **Jordi Sabater** [2] and **Sarvapali D. Ramchurn** [3] and **Michael Luck** [4] and **Nick Jennings** [5]

**Abstract.** Current mechanisms for evaluating the trustworthiness of an agent within an electronic marketplace depend either on using a history of interactions or on recommendations from other agents. In the first case, these requirements limit what an agent with no prior interaction history can do, and in the second case, they transform the problem into one of trusting the recommending agent. However, these mechanisms do not consider the relationships between agents that arise through the *interactions* between them, such as buying or selling, or through overarching *organisational* structures, which can also aid in evaluating trustworthiness. In response, this paper outlines a method that enables agents to evaluate the trustworthiness of their counterparts, based solely on an analysis of such relationships: relationships are identified using a generic technique in conjunction with an ontological model for agent-based marketplaces; they are then interpreted through a trust model that enables the inference of trust valuations based on the different types of relationships. In this way, we provide a further component for a trust evaluation model that addresses some of the limitations of existing work.

## 1 Introduction

Agents generally interact by making commitments to, or contracts with, one another to carry out particular tasks. However, in most realistic environments there is no guarantee that a contracted agent will actually enact its commitments (because it may defect to gain higher utility, or because there is uncertainty about whether the task can be achieved). In such situations, computational models of trust (here defined as the positive expectation that an interaction partner will act benignly and cooperatively in situations in which defecting would prove more profitable to itself [3]) have an important role to play. First, they can help determine the most reliable interaction partner (i.e. those in which the agent has the highest trust) and second, they can influence the interaction process itself (e.g., an agent's negotiation stance may vary according to the opponent's trust level).

However, when an agent first enters an environment, it has no history of interactions (with the other agents in that environment) that it can analyse to decide who to trust. In such circumstances, current research suggests two possible solutions. It could interact with all agents and then derive trust measures from the history of interactions [9]. Alternatively, it could request reputation information acquired from an existing social network [10], in which reputation is understood as a third party's estimate of trustworthiness. However, there are a number of problems in each of these alternatives. Firstly, if the agent interacts with each agent, it inevitably risks making losses if the counterparts it interacts with are not trustworthy. Secondly, if the agent relies on reputation information, then it cannot be sure that the agents providing the information are indeed truthfully providing information about their counterparts. In both cases, others may be unreliable because they have conflicting interests with the agent (e.g., if they compete in the same market) or because they can collude to exploit the agent (e.g. if some agents know each other and all share their gains from exploiting an agent). In consequence, if an agent could take into account relationships such as competition or situations such as collusion it could produce more robust valuations. However, existing mechanisms do not adequately consider the relationships between agents that arise through the *interactions* between them, which may lead to competitive relationships, or through overarching *organisational* structures, which may lead them to collude. Furthermore, when relationships *are* taken into account in a limited manner, such as in [10], the information used is implicitly assumed and no mechanisms are provided to enable the agent to *discover* that information dynamically, nor to react to changing relationships between counterparts. An agent also needs to be able to *identify* and *interpret* such relationships in a changing environment.

In this paper we address just this need by developing a method for identifying such relationships between agents in an electronic marketplace and then using this information to enhance trust valuations. We advance the state of the art in the following ways. First, we develop a process for agents to dynamically identify relationships between agents in an electronic marketplace. Second, we identify the general types of relationships that should be considered with regards to trust and discuss the types of reasoning such information can enable. Finally, we make use of an ontology-based framework to analyse relationships, providing a realistic application of semantic web technologies.

The following section provides an overview of our approach, while Section 3 describes the agent-based market model that provides the types of information required in order to determine what relationships exist. Section 4 describes the relationship identification process used and how that is mapped to the specific context of an electronic marketplace. Section 5 introduces the most relevant relationship types with regard to trust and Section 6 discusses how knowledge of such relationship can affect trust valuations as well as how the work here can directly be used by existing trust models. Conclusions and further work are given in Section 7.

## 2 Approach

Our overarching aim is to enable the derivation of trust information based on the relationships that one agent can infer its counterparts

[1] University of Southampton, Southampton, SO17 1BJ, United Kingdom, email: {ra}@ecs.soton.ac.uk
[2] Institute of Cognitive Science and Technology, CNR, Roma, Italy, email: {jsabater}@iiia.csic.es
[3] University of Southampton, Southampton, SO17 1BJ, United Kingdom, email: {sdr01r}@ecs.soton.ac.uk
[4] University of Southampton, Southampton, SO17 1BJ, United Kingdom, email: {mml}@ecs.soton.ac.uk
[5] University of Southampton, Southampton, SO17 1BJ, United Kingdom, email: {nrj}@ecs.soton.ac.uk

have with each other and itself. The various steps to achieve this aim are described below.

1. ***Relationship Identification*** First, we *identify* relationships between agents. In order to do so, we must have some information about agents, and a process that can lead from that information to knowledge about relationships. The types of information that we can expect agents to use to identify relationships are captured and related through an ontology that represents an *Agent-Based Market Model* (ABMM) and is based on a typical e-commerce example of an electronic marketplace. This information is then mapped to a generic relationship identification process [1] that uses a model of agent interaction with their environment to infer when agents may be related given their individual capabilities: in this case, we consider what an agent can buy or sell, and their goals.
2. ***Relationship Characterisation*** Having identified the possible relationships, within the context of the ABMM, we then distinguish the types of relationships that are most relevant with regard to trust. These types provide us with relationship *patterns* that an agent can use in trust evaluation.
3. ***Relationship Interpretation*** Using these relationship patterns, and their interpretation through additional information about the specific context in which an agent operates, which is also captured by the ABMM, we discuss how an agent can use a trust model to derive trust valuations.

## 3 Agent-Based Market Model

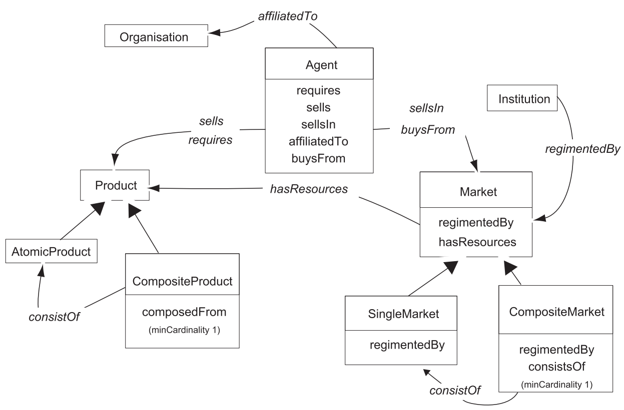


**Figure 1.** Agent-Based Market Model

The ABMM aims to capture most of the features of a typical e-commerce scenario by which sellers and buyers trade in an online market.[6] The model is defined using the OWL (Web Ontology Language) [7] standard, which is a natural choice, since it enables us to check the consistency of the model and reason about it through the underlying description logics [2], using widely available tools.[7] In addition, we benefit from a standardised representation, allowing for easier acceptance and integration into existing agent toolkits.

[6] We believe these features of the model are necessary rather than sufficient ones. The model can be easily adapted to cope with less or more features if necessary as will be discussed later on.

[7] The ontology was developed using the Protege toolkit [8] with the OWL plugin and tested with the RACER engine [6].

The information types and the relationships between them are illustrated in Figure 1, in which large arrowheads represent inheritance relationships while smaller arrowheads represent property references. In this model, an *Agent* is considered to be any entity that *requires* or *sells* a number of *Products* (e.g. memory chips or computer processors), and can be affiliated to an *Organisation*. We distinguish between an *AtomicProduct* (e.g. computer chips), for which no further division of the product into components takes place, and a *CompositeProduct* (e.g. desktop computer), which comprises several atomic products. This distinction allows us to better represent the situation in which an agent requires a composite product whose atomic components must be sourced from a number of seller agents. Now, an agent *buys from* or *sells in* a *Market*, which has a number of resources of type *Product*. A *Market* is *regimented by* an *Institution*. An *Institution* is an entity that regiments the roles and relationships of the interacting agents, and determines the rules of encounter that prescribe what an agent can do at what point in time [5].

Similarly to a *Product*, a *Market* can be a *SingleMarket* or a *CompositeMarket*. A composite market is a market in which the goods traded are inter-related (e.g. buyers and sellers of particular car parts and as well as second-hand cars may trade on a single website). The composite market may also be different from a single market in that it is regimented by more than one institution.

Once an agent has information about its counterparts as described and related by the ABMM model, it can begin the process of identifying how they are related. For example, if two agents sell the same products in a market, we can assume that they are competitors. If they sell complementary products in a composite market and belong to the same organisation, then we could assume that their opinions of each other may be biased. Furthermore, the use of OWL enables us to improve the process by performing some basic types of reasoning about the information captured in the model. For example, if one agent sells computer components and another sells memory chips, we could infer that they are competitors if we also had access to an ontological model describing computer components, which would allow us to identify that a memory chip is a type of a computer component.

## 4 Identifying relationships

The ability to identify the different types of relationships in a marketplace allows us to reason about the possible underlying motives of agents and, as a result, derive trust valuations. For example, consider a situation in which agent A sells a product to agents *B* and *C*, and *A* and *B* belong to the same organisation. Then, if we ask *B* for a rating of A's product quality we should not place much credence on the reply, since an ulterior motive, relating to the overall gain of the organisation that *A* and *B* belong to, could bias the reply.

However, in order for the identification process to be widely applicable we require a principled approach that can be made part of an agent program. In this section, we adapt an existing relationship identification process to the task of identifying relationships within the context of the ABMM. This process is particularly suitable since it is based on a model of interaction between agents and their environment that makes no assumptions about any internal agent components, since they cannot be observed. The focus, therefore, is on the *interface* between individual agents and their environment, through the capabilities of agents. The notions that underpin this model are based on the SMART framework [4], and are discussed in more detail in [1], so they are only briefly described below. We first present the underlying agent model provided by SMART and then explain how we use it to create a model of interaction with the environment, and by consequence other agents.

**Agents** For the purposes of relationship analysis, an agent is considered an entity described by a set of *attributes*. Attributes are simply describable features of the *environment*, and are the only characteristics that are manifest. Agents are able to perform *actions*, which can change the environment by adding or removing attributes. Agents also pursue *goals*, which are desirable environment states described by non-empty sets of attributes.

In the specific case of market agents, the attributes include information such as the organisation an agent belongs to, the market in which it operates, the available products, and so forth. The basic actions are the ability to buy or sell products. Finally, goals represent the *desire* to buy or sell a specific product.

**Agent Perception and Action** Agent actions are divided into those that retrieve the values of attributes, representing the agent's *sensor capabilities*, and those that attempt to *change* attribute values of the environment, representing the agent's *actuator capabilities*.

With regard to the ABMM, sensor capabilities are those that allow the agent to perceive other agents in a market and identify relevant information such as available products. The most relevant actuator capabilities are those that allow it to sell or buy a product.

**Viewable Environment and Region of Influence** Given that agents interact with the environment through actuators and sensors, and that the environment as a whole is defined through a set of attributes, we can intuitively think of actuators and sensors as defining *regions of the environment*, or subsets of the entire set of attributes that make up the environment. The attributes that an agent's actuators can *manipulate* define a *Region of Influence* (*RoI*), while the attributes that an agent's sensors can *view* define a *Viewable Environment* (*VE*).

The *VE* and the *RoI* of an agent provide us with a model that relates an agent and its individual capabilities to the environment. In order to identify relationships between agents we need to look at how their *VEs* and *RoIs* overlap. The different ways in which these overlaps occur plays a role in determining the possible relationships between them. In Figure 2, these concepts are illustrated by using an ellipse to represent the *VE* and a pentagon shape for the *RoI*. We use this notation throughout when illustrating different situations.

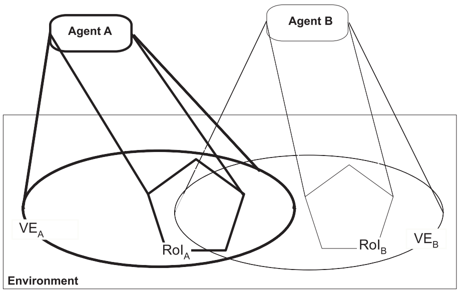


**Figure 2.** Region of influence affects viewable environment

With respect to the ABMM, we assume that the *VE* of an agent defines a region of the market that an agent is able to view. Furthermore, if two agents belong to the same organisation, we could decide, depending on the nature of the organisation, to represent the *VE* of each individual agent as the sum of the *VEs* of each member of the organisation, reflecting the assumption that agents would share information. Now, the *RoI* of an agent represents the products that an agent can sell or buy in a market. In this context, the *RoI* can be divided into an *RoI* related to buying products and an *RoI* related to selling products. In this paper, we make the assumption that the *RoI* for buying products is the entire *VE* since an agent could decide to buy anything in a market provided it has the credit to buy it. There is thus no need to represent the *RoI* for buying, since it overlaps with the *VE*, and the only *RoI* we represent graphically covers products that an agent is able to *sell* in the market. This simplification does not detract from the analytical power of the model and makes the overall analysis more straightforward. In order to clarify these notions we illustrate them through an example.

Figure 2 shows the situation in which *A*'s *RoI* overlaps with *B*'s *VE*, and both agents' *VEs* overlap. Given this information, we can infer that *A* and *B* operate in a common market where the *VEs* overlap. Furthermore, *A* sells a product in that market where its *RoI* overlaps with the common *VEs*.

Now, assume that *B* has the goal to buy a product from *A*. *B* uses its sensory capabilities to identify the price and other relevant attributes of the product. In addition, assume that *B*'s *RoI* represents the sale of a product that *B* constructs using, in part, the products bought from *A*. As a result, *B* now becomes *dependent* on *A* making that product available at an appropriate price. Thus, whenever *A* performs an action that affects that product in some way, it will eventually *influence B*'s actions, since *B* must now react to the changes when producing its own product. We discuss in Section 6.1 the implications of such a relationship with regard to trust.

**Goals** Knowledge of an agent's goals, in addition to its *RoI* and *VE*, can provide better information about its possible relationships with other agents, as we have seen in the example above. In the ABMM, the relevant goals are the desire of an agent to buy or sell something. If an agent needs to buy something, then it must find another agent whose *RoI* contains the required product while an agent can only sell something which, by definition, be part of its own *RoI*.

In the next section, the ABMM and the relationship identification process are used to identify a number of relationship types that are needed when evaluating the trustworthiness of another agent.

## 5 Interpreting relationships

By combining the process described above with the ABMM, we can now identify a variety of different types of relationships between agents. Even so, it is the *interpretation* of these relationships that ultimately determines how the trust valuation for an agent may change because of them. However, as there is a large number of possible relationships and combinations of relationships, as well as a large number of possible interpretations of those relationships, it is necessary to minimise the amount of decision making needed in order to interpret the significance of the relationship. In this section, to tackle this problem we define some *basic* types of relationships that are clearly relevant to trust valuations and can be combined to describe more complex types. In the next section we discuss how these relationships can be interpreted.

The relationship types defined here build on results from an existing trust model [10], and represent the most salient types with regards to trust. Each type is represented by a *pattern* representing a specific configuration of *VEs*, *RoIs* and goals. In addition, the final interpretation depends on other issues that are not directly captured by the pattern alone. Issues such as the abundance of a product, the number of sellers of the product, and the amount being bought, define what we term the *intensity* of the relationships. We discuss some of these issues as we present the various patterns.

Note that in order to make this analysis we need to consider a specific context (in our case the context of e-markets). The same configuration of *VEs* and *RoIs* in a different context, or based on a different ABMM, would possibly have a completely different interpretation in terms of relationships.

Below, we provide brief definitions of the four basic types we deal with, and then discuss how they can be identified through patterns.

**Trade** relationships exist when a transaction takes place.

**Competition** relationships exist when agents pursue the same goals and may need the same (usually scarce) resources.

**Dependence** relationships exist whenever one agent depends on another for it to accomplish its goal.

**Collaboration** relationships exist whenever two agents depend on each other to accomplish their respective goals.

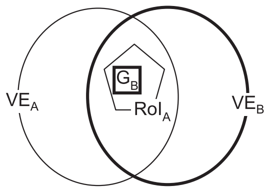


**Figure 3.** Trade-Dep: *A* sells goods to B

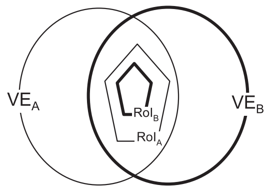


**Figure 4.** Comp-Sell: *A* and *B* are competing in *A*'s RoI

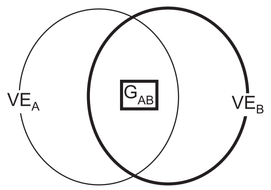


**Figure 5.** Comp-Buy: The goal of *A* is the same that the goal of *B*

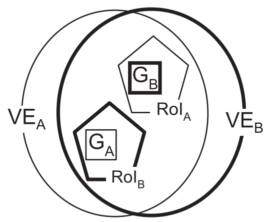


**Figure 6.** Coll: *A* sells to *B* and *B* sells to *A*

**Trade and dependency** The configuration in Figure 3 corresponds to an agent *A* selling goods in a market that *B* can view and buy from. At the same time, *B* has the goal (represented by a square) to buy the goods *A* is selling in that market. This configuration, termed a *Trade-Dep* configuration, identifies two kinds of relationships: a trade relation and a dependence relation. The *trade* relation is obvious; *A* is selling goods to *B*. The intensity of this relation is associated with the amount (or the value) of goods in the transactions. The *dependency* relation is due to the fact that *B* needs the products that *A* is selling. The intensity of this relationship depends on several factors: the number of sellers in that market that provide the same product, the abundancy of these products in general and the intensity of the *trade* relation. Notice that these factors can ultimately determine who depends on whom. For instance, in a market where there are many sellers providing the same product, and very few buyers interested in that product, the roles in a dependency relation like the one described above are interchanged. Then the seller depends on the buyer.

**Competition** In the configuration of Figure 4, termed a *Comp-Sell* configuration, the RoIs of agents *A* and *B* intersect. This implies that both agents are selling the same goods in the same market. This reflects a competition in that area of influence. The intensity of this competitive relation is determined by several factors such as market share, profit, and cost of goods.

The configuration of Figure 5 also reflects a competitive relationships, termed (*Comp-Buy*). In this case, *A* and *B* have the same goal, indicating that they want to buy the same products. The intensity of this competitive relation is based on factors similar to the dependence relation presented in Figure 3: the number of sellers in that market that provide the products required by *A* and *B*, and the abundance of these products in general.

**Collaboration** Figure 6 shows a configuration in which *A* has a goal in the *RoI* of *B* and *B* has a goal in the *RoI* of *A*. This means that *A* is selling goods to *B* and, at the same time, *B* is selling (different) goods to *A*. This configuration, called a *Coll* configuration, is a composition of two Trade-Dep configurations (see Figure 3). The relationships generated by the two Trade-Dep configurations are a trade and dependence relation between *A* and *B*, and a trade and dependence relation between *B* and *A*. If *A* depends on *B* and *B* depends on *A*, we say there is a *collaboration* between *A* and *B*.

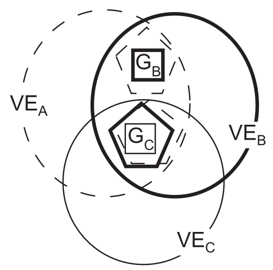


**Figure 7.** *A* and *B* competing to sell to *C* while *A* is selling to *B*

**Tripartite relationships** In this paper, we focus on relationships between just two agents. However, in this section we provide an indication of how relationships between more agents can be considered, if at least one more agent is added to the analysis.

In the majority of cases, the resulting configurations can be decomposed in terms of the configurations above. One of the exceptions to this is the configuration showed in Figure 7, in which *A* is selling a product that *B* requires, while at the same time *A* and *B* are competing in order to sell *C* a product. Thus, there are three *Trade-Dep* configurations and a *Comp-Sell* configuration. However in this configuration there is a special situation, since while is a competitive relation between *A* and *B* at the same time *B* depends on *A*. This situation gives a privileged position to *A* with respect to *B*, which has to be considered when analysing trust.

## 6 Trust Valuations

Having introduced some of the most salient relationships that can be directly inferred using the ABMM and the relationship identification process, we discuss here how we can make use of them to derive trust valuations. Given our initial definition of trust (in section 1), we argue that an agent should *distrust* its counterpart whenever the latter has an opportunity to defect, as can be inferred from the relationships with the counterparts and the counterpart's relationships with others. We deal with the trust valuations in two parts: (i) where one agent tries to infer the trustworthiness of its counterpart (bipartite relationships) and (ii) where one agent tries to infer the trustworthiness of its counterpart and both or one of them is related to other agents (multipartite).

**Bipartite Relationships** The reasoning that knowledge of bipartite relationships enables with regard to trust is described below for each type of relationship identified above.

1. *Trade-Dep* – if agent *B* is dependent on *A*, *A* may have an opportunity to exploit *B*, if *B* has no other choice than *A* as an interaction partner. In the case where the intensity of dependence is high (e.g. in terms of amount of goods traded and percentage of total costs to *B*), *B*'s trust in *A* should be the lowest possible (and conversely if the dependence is low). Thus, when trust is low, *B* should negotiate more strongly for very low prices and for more stringent contracts to reduce the possibility for *A* attempting exploit it. However, if the institution *A* and *B* interact in prevents *A* from behaving unfairly (e.g. by charging high prices), then *B* does not need to increase or decrease its trust in *A* since its behaviour is independent of its motives. [8]
2. *Comp-Sell* and *Comp-Buy* – these competitive relationships obviously do not favour trust between the agents since it is in their interest to undermine each other in all possible ways (as regimented by an institution). In such cases it becomes more important to understand how such agents are related to other agents as we will see later.
3. *Coll* – in this case, both agents gain by not defecting during their interactions since they both depend on each other to achieve their goals. Depending on how intense this relationship is, we would expect these agents to trust each other highly if they are strongly interdependent, and not place much trust in each other if they are not equally dependent on each other (e.g. if *A* depends on *B* more than *B* depends on *A*, *B* could defect on *A* and *A* will not be able to compensate such defections by defecting on *B*).

**Multipartite Relationships** We can now discuss how the intensity of relationships and their combinations (i.e. where agents can be related to more than one other agent) can lead to more informed decisions given the environment these agents interact in. Some examples are given below.

1. *A* depends on *B* in *Trade-Dep* while *B* and *C* are in *Coll*. In this case *B* may have an incentive to misrepresent the reliability of *C* to *A*. This may happen because *B* could gain from a more profitable *C*, and would therefore provide unrealistically high rating for *C*. Alternatively, *B* may wish to continue to hold *C* captive in their collaboration, and provide unrealistically low ratings so that *C* is not able to become more independent. Given the intensity of these relationships, the credibility of *B*'s reports will be decreased or increased to different degrees. However, if roles between *A* and *B* where reversed, and *B* depended on *A* in *Trade-Dep*, then the converse of the above reasoning may apply and *A* might give more value to *B*'s reports on *C*.
2. *A* is in *Comp-Sell* or *Comp-Buy* with *B* and *B* is in *Coll* with *C*. In this case, *A* will obviously distrust *B*'s reports about *C* since *B* could gain from giving false reports about *C* to *A*, as discussed above.
3. *A* is in *Coll* with *B* and *B* is in *Coll* with *C* and *A* is in *Coll* with *C*. All the agents should trust each other fully. This situation may arise if all the agents form part of the same organisation or form a cartel. The latter form of collaboration might be very profitable to the agents *A*, *B* and *C*, but might affect the performance of the system within which they operate with other agents. In such circumstances, it is up to the institution to ensure good behaviour.

One particular model that studies how such combinations of relationships influence trust valuations of agents is the ReGreT model [10]. ReGreT uses fuzzy sets to capture the intensity of relationships that might exist between agents or groups of agents in order to elicit quantitative measures of trustworthiness of agents or the ratings they may provide. For example, in ReGreT if a witness *A* is in a 'strong' *Coll* relationship with *B*, then agent *C* who is strongly dependent on *A* in *Trade-Dep* might value ratings of *A* about *B* very 'low', where 'strong' and 'low' represent fuzzy sets characterising the intensity of relationships.

Similarly, and also by means of fuzzy rules, the ReGreT system uses relationships together with direct experiences in order to assign trust values to other agents. In other words, it uses prejudice as a mechanism for evaluation. For example, if an agent *A* usually offers good quality resources and agent *B* has a 'strong' *Coll* relation with it, the system will assign to agent *B* a high trust value associated to the quality of the products (see [10] for more details).

In this sense, trust and reputation systems like ReGreT that rely on relationships to improve the computation of reputation and trust values can take advantage of the work presented in this paper.

## 7 Conclusions and Further Work

In this paper, we have presented a novel process for identifying relationships between agents in an electronic marketplace and discussed how this information can be used to reason about the trustworthiness of agents. By doing this, we address a shortcoming of existing trust models, since they typically do not consider such relationships and where they attempt to incorporate them in a trust model they provide no mechanisms for *identifying* them automatically and do not enable sophisticated reasoning about the complex range of scenarios that may arise. Furthermore, we have discussed how the work presented here can be used directly within existing trust models, such as ReGreT, to supplement its existing approach. By combining existing models with this work we further the development of robust trust valuation models in a constructive and immediate manner.

In the future, we aim to deal with more complex combinations of relationships (i.e. more than 3 agents in all relationships) and explore ways of analysing such combinations of relationships using learning or case-based reasoning tools.

## REFERENCES


[1] R. Ashri, M. Luck, and M. d'Inverno, 'On Identifying and Managing Relationships in Multi-Agent Systems', in *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, eds., George Gottlob and Toby Walsh, pp. 743–748. Morgan Kaufmann Publishers, (2003).

[2] *The Description Logic Handbook*, eds., F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, Cambridge University Press, 2003.

[3] P. Dasgupta, 'Trust as a commodity', in *Trust: Making and Breaking Cooperative Relations*, ed., D. Gambetta, 49–72, Blackwell, (1998).

[4] M. d'Inverno and M. Luck, *Understanding Agent Systems*, Springer, 2nd edition edn., 2004.

[5] M. Esteva, D. de la Cruz, and C. Sierra, 'ISLANDER: an electronic institutions editor', in *The First International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 1045–1052. ACM Press, (2002).

[6] V. Haarslev and R. Möller, 'Description of the RACER System and its Application', in *International Description Logics Workshop (DL-2001)*, eds., C. A. Goble, D. L. McGuinness, R. Möller, and P. F. Patel-Schneider, volume 49 of *CEUR Workshop Proceedings*, (2001).

[7] D. L. Mcguinness and F. van Harmelen. OWL Web Ontology Language: Overview. http://www.w3.org/TR/2003/PR-owl-features-20031215/.

[8] N. F. Noy, M. Sintek, S. Decker, M. Crubezy, R. W. Fergerson, and M. A. Musen, 'Creating semantic web contents with protege-2000', *IEEE Intelligent Systems*, **16**(2), 60–71, (2001).


[8] The effect of institutional regulation is discussed in [9].

[9] S. D. Ramchurn, C. Sierra, L. Godo, and N. R. Jennings, 'A computational trust model for multi-agent interactions based on confidence and reputation', in *Workshop on Deception, Trust, and Fraud, AAMAS*, eds., R. Falcone, S. Barber, L. Korba, and M. Singh, pp. 69–75, (2003).

[10] J. Sabater and C. Sierra, 'REGRET: a reputation model for gregarious societies', in *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multi-Agent Systems*, eds., C. Castelfranchi and L. Johnson, pp. 475–482, (2002).