

A Situation-Aware Computational Trust Model for Selecting Partners

Joana Urbano, Ana Paula Rocha, and Eugénio Oliveira

LIACC – Laboratory for Artificial Intelligence and Computer Science
Faculdade de Engenharia da Universidade do Porto – DEI
Rua Dr. Roberto Frias, 4200-465, Porto, Portugal
{joana.urbano, arocha, eco}@fe.up.pt
<http://www.fe.up.pt>

Abstract. Trust estimation is a fundamental process in several multi-agent systems domains, from social networks to electronic business scenarios. However, the majority of current computational trust systems is still too simplistic and is not situation-aware, jeopardizing the accuracy of the predicted trustworthiness values of agents. In this paper, we address the inclusion of context in the trust management process. We first overview recently proposed situation-aware trust models, all based on the predefinition of similarity measures between situations. Then, we present our computational trust model, and we focus on Contextual Fitness, a component of the model that adds a contextual dimension to existing trust aggregation engines. This is a dynamic and incremental technique that extracts tendencies of behavior from the agents in evaluation and that does not imply the predefinition of similarity measures between contexts. Finally, we evaluate our trust model and compare it with other trust approaches in an agent-based, open market trading simulation scenario. The results obtained show that our dynamic and incremental technique outperforms the other approaches in open and dynamic environments. By analyzing examples derived from the experiments, we show why our technique get better results than situation-aware trust models that are based on predefined similarity measures.

Keywords: computational trust systems, dynamics of trust, situation-aware trust, multi-agent systems.

1 Introduction

Computational trust systems are systems that collect information, such as evaluation of direct experiences, reputation, recommendations or certificates, about agents (e.g. representing individuals or organizations) and that computes the trustworthiness of these agents based on the collected information. In this document, we envision trust as the confidence that an evaluator agent has on the capabilities and the willingness of another agent (the trustee) in fulfilling some kind of task in a given context.

Trust management is in these days a fundamental research topic in agent-based and multi-agent systems, and its appliance concerns decision making processes in almost all electronic forms of commerce and social relationship. In this context, several trust and reputation models have been proposed to represent and aggregate social evaluations into trustworthiness scores (see [1] for an overview). However, the majority of these models are based on simple aggregation engines that do not differentiate between situations and therefore use all available evidences about an agent to compute its trustworthiness, independently of the context of the current assessment, potentially jeopardizing the accuracy of the estimated trust value.

In order to address this limitation, some models have been recently proposed that bring situation into the trust loop [2][3][4][5][6][7]. Some of them rely on the existence of taxonomies that allow defining a situation, and all of them use predefined domain-specific similarity metrics. Our belief is that these systems are more adequate to environments where the evidences about the agents in evaluation are abundant and where there is a relative stability concerning the context definition. However, there are real world scenarios where these assumptions are not guaranteed. For example, in e-sourcing activities performed in open and global markets, it is expected that agents risk new partnerships outside the embedded set on known acquaintances, which, consequently, reduces the availability of direct evaluations of the potential partners in the several distinct contexts that can be considered. In such harsh scenarios, we need different and more dynamic types of trust models that are able to reasoning even when the available evidences about the agents in evaluation are scarce, contextual and heterogeneous.

In this paper, we address the development of trust models that are adequate to these scenarios. First, we present our situation-aware trust model that is based in three fundamental characteristics:

- The inclusion in the aggregation procedure of important properties of the dynamics of trust, as addressed in research areas related to social sciences and psychology;
- The dynamic detection of tendencies of agents' behavior in the presence of different situations;
- The ability to infer the trustworthiness of agents even when the available trust evidences are scarce.

Then, we set up a multi-agent simulation environment where textile client agents seek for optimal deals by selecting from a range of suppliers with different behaviors. We are interested in evaluating the benefits of using our own trust proposal for the selection of partners in open and global scenarios and to compare it with other trust approaches. Also, we intend to evaluate how the different trust methods can support the exploration of new potential partners in such a way that the risk associated to trading with strangers is decreased by the method. A final achievement of the work presented in this paper is the development of a simulation environment that can be used to support important studies about parochialism and trust that are being done in the social sciences area [8] [9].

This paper is structured as follows: Section 2 briefly introduces situation-aware trust and overviews the most significant computational situation-aware trust models existing to date. Section 3 presents our own contextual trust model, and presents the scenario and notation used in this paper. Section 4 evaluates our approach and compares it with other computational trust models, in the open market scenario defined. Finally, Section 5 concludes the paper and refers future work.

2 Situation-Aware Trust

There are several definitions of trust in the literature, and most of them refer context as a dimension to take into consideration in the trust assessment process. For instance, in one of the earliest research on computational trust, Marsh (1994) considers situational trust as “the amount of trust an agent has in another in a given situation”, and gives a helpful example: “I may trust my brother to drive me to the airport, I most certainly would not trust him to fly the plane” [10]. Another well known definition of trust, from the economics area, is given by Dasgupta, where trust reflects the expectation on the activities of an entity when it reacts on a given context [11]. In the same line of thought, Dimitrakos defines trust as a measurable belief that a given entity has on the competence of the trustee entity in behaving in a dependably way, in a given period of time, within a given context and relative to a specific task [12].

Therefore, the need for some kind of computational situation-aware trust seems evident. In one hand, it is realistic to assume that not all the past evidences are equally relevant to future interactions, as it is common sense that a given entity might behave differently in different contexts, as shown by Marsh definition above. It is the same to say that the trust aggregator algorithm shall not take all positive evidences concerning the car driving ability of the author’s brother when assessing his ability to drive an airplane.

This same reasoning further extends to reputation models based on transitivity graphs, much used in recommendation systems. In fact, although existing models of reputation are generally based on the transitivity of trust, some authors consider that trust is not transitive [13]. In this context, Tavakolifard (2009) suggests a theoretically approach that incorporates a situational dimension into graph-based reputation models [14].

There are other theoretically benefits that can derive from the consideration of context into computational trust models. One of these benefits is the bootstrapping of unanticipated situations, where missing information from a given target agent can be inferred from similar situations (e.g. “A person trusting Bob as a good car mechanic will not automatically trust him also in undertaking heart surgeries (. . .) [but] he probably could be quite good in repairing motorcycles” [2]); the other is the management of newcomers, where the use of similarities between both agents and situations can generate an approximate estimate for the initial trustworthiness value of the newcomer. Finally, Neisse et al. consider that the consideration of context may allow the reduction of the complexity of management of trust relationships [7].

2.1 Overview on Situation-Aware Models

Although the need for contextual trust has been identified several years ago, only recently there started to appear practical approaches of situation-aware trust models. However, the majority of them is based on ontologies or taxonomies and implies the predefinition of domain-based similarity measures of situations or roles.

The first example we analyze is defined in the domain of message-based communications. In this work, the authors propose to use rules based on the explicit context of messages in order to determine the trustworthiness of these messages [15]. In order to capture the message content and the context-dependent trust relations, the authors extend a trust ontology proposed by Golbeck et al. [16] in order to include the notion of context-sensitivity trust in messages. However, the authors do not address important questions, such as how the context is extracted from messages and how to compute the context-sensitive factor that allows determining the overall trustworthiness of a message.

Another example of such a model is the Context Management Framework (CMF) model, where trust relations in one domain are used to infer trust relations in similar domains [2][14][17]. The model uses case-base reasoning techniques to estimate trust in unanticipated situations, by retrieving the most similar cases from a case base. In this approach, the authors use a trust-based ontology ([18]) to specify situations and they measure the similarity between situations using a relational similarity model based on the SimRank algorithm [19]. A major drawback of this approach resides in the weak assumption made by this algorithm about the similarity between different objects, as even signaled by the authors ([2]) which prevent the model to adequately scale to more complex representations of contexts.

The CMF model drinks from the collaborative filtering research area that use taxonomy-based similarity measures for recommendation purposes. There, the similarity between users is normally derived from the items or classes of items they tend to share (also known as co rating behavior). A recent extension on this topic is cross-domain recommendation. In this context, a recent work [20] proposes a model that allows users that share similar items or social connections to provide recommendations chains on items on other domains. This model uses Web taxonomies made available by service providers. In order to allow measuring the similarity of users that do not have rated the same items, the model first computes class classifications from the individual classifications of items of the class, and then computes the similarity between users taking into account the resulting similarities of users in each class. This model implies, however, that subtleties between items within classes cannot be taken into account because of the generalization process it assumes.

The next model also uses an ontology for describing situations, or concepts (e.g. business orders) [3]. In this approach, the expected behavior of an agent on a given situation is represented as a conditional probability distribution function (PDF) over the possible observations given the possible agreements. As in the previous model, a concrete experience about a commitment can be used to update the expectation of behavior over semantically close commitments, allowing for faster

bootstrapping of unanticipated situations. Although the authors state that the PDF is initialized using background knowledge on the other agents, the model suffers from the a priori information limitation of conditional probability distributions. Also, the authors do not present an evaluation of the model, which prevents a further understanding of its potential benefits.

A related approach by Hermoso and Billhardt dynamically builds a taxonomy of organizational roles from trust information maintained by the agents, using clustering techniques [4]. When evaluating the trustworthiness of a target agent, the evaluator uses the direct experiences it has with the agent and weights them according to the similarity between the role assumed by the agent in the specific experience and the role of the current situation.

Finally, another approach in the area of social browsing [21] allows the representation of context based on personal ontologies of Web collaborators and provide context mapping methods.

Although the use of taxonomies and ontologies is increasing in the Web, both in social networks and e-business activities, the main limitation associated with the above referred models is the necessity to predefine adequate similarity measures for all possible situations in assessment before such situations are even presented to the evaluator. This is a domain specific, hard tuning process that may be a challenge in dynamic environments with complex representations of contexts.

A somewhat different approach is presented in the Context Space and Reference Contexts model [5][22][23]. This approach defines the context space as a Q -dimensional metric space, with one dimension per each represented situation feature, and places $n < Q$ reference contexts (either regularly or adaptively) over the context space. There exist as many context spaces as agents in evaluation. Then, every time a new trust evidence is generated for a given agent, the trustworthiness at each one of the reference contexts of the agent's context space is updated with the outcome of the evidence, weighted by the similarity between the reference context and the context of the new evidence. This interesting model is, however, also dependent on predefined measures of similarity between contextual attributes. Also, the consideration of multiple dimensions can lead to an exponential number of reference contexts that each evaluator needs to keep for every agent, jeopardizing the scalability of the model to complex contextual scenarios. Another limitation of this model is that it can only be used complementary to traditional trust systems that aggregate evidences using weighted means approaches.

Finally, there is a related approach [6] to the previous model that differentiates from it on the definition of the context space, where each dimension of the context space is now a function of many context attributes, instead of one sole attribute. However, all limitations pinpointed above for the previous model are still present in this approach.

The situation-aware trust model we propose in this paper distinguishes from the above models by performing online evaluation of the trust evidences, as it is able to dynamically extract tendencies of the (probably changing over time) behavior

of the trustee, on a given specific situation. This means that our approach does not rely on a priori definition of similarity measures. It is our firm belief that such a model is more adequate to detect subtleties between two similar contexts and then to perform better in open and dynamic environments than the models we overviewed above. In Section 4, we present the results of the experiments we have performed that strongly support our belief.

3 The Proposed Trust System

We developed a trust method envisioning its use in future semi-automatic and open business-to-business markets, taking into special consideration the exigent requirements of such environments. Namely, we are concerned with the performance of the method when the trust evidences available about a given agent are scarce and heterogeneous, and when the activity of the agents under evaluation can span through different situations and contexts. The current implementation of our system that encompasses the proposed method is composed of two different components, as depicted in Figure 1:

- The aggregator component, which is responsible for aggregating the available trust evidences of an agent into a trustworthiness score for this agent. Several trust engines that are defined in the literature can be used (cf. [1]), although we are particularly interested on engines that model the dynamics of trust, as they appear to perform better than the traditional statistical approaches [24][25][26];
- The Contextual Fitness component, which tunes the outcome of the aggregating step by taking into consideration the specificities of the current business opportunity and the adequacy of the agent in evaluation to the specific situation under assessment.

The idea beyond this extension is that if the trust system detects that an agent has some kind of handicap related to the current necessity, the overall trustworthiness of the agent for this necessity will be zero; otherwise, the trustworthiness score is the value computed by the aggregator component. One good characteristic of this approach is that it can be used with any conventional trust aggregation engine, being it based on statistical, probabilistic or heuristic-based models. Before we describe the Contextual Fitness component in more detail, we first introduce the notation and the scenario used all over this paper.

3.1 Scenario and Notation

The scenario used in this paper situates on the automation of the procurement processes of business-to-business relations in the textile industry. Ideally, the development of adequate agreement technologies would allow opening the space of business opportunities, where agents would have the means and the confidence knowledge to search for good business opportunities outside their limited sphere of breeding relations. However, the fear of risking unknown partners is one of




Fig. 1. The current implementation of our trust system

the biggest barriers to trade in a truly open market environment. Taking our subcontracting scenario in the textile industry as an example, the client part of the relation can be deceived by the provider part in several different ways [27]:

- A delay in delivery, which affects all the supply chain;
- The quality received, as specified by affordability, safety and degree of uniqueness parameters;
- The quantity received (too much or too less);
- The violation of intellectual property rights;
- Ethical problems;
- Other problems, such as price rise and legislative changes.

Without some kind of a trust mechanism, it is reasonable to conclude that business partners would preferentially adopt parochial environments in detriment to more aggressive exploration of deals outside the space of already known partner relationships. For instance, in the fashion retail industry, clients often rely on knowledge available through textile fairs and textile agents to make the bridge between brands and the trustable and reliable textile suppliers. However, even with this form of trust guarantees, the space of available suppliers is relatively small and it is strongly supported on the expected behavior of the partner, rather than on the potential real utility of the business transaction.

In our scenario, business clients in the textile industry try to select the best suppliers of textile fabric. For this, they announce their business needs and wait for proposals from the suppliers. The need is announced through a call for proposals (CFP) and concerns the delivery of some quantity of a fabric due in some delivery time. An example of such a need is the following: (fabric = cotton, quantity = 900000 meters, delivery time = 15 days). In this context, we define $at \in AT$ as the description of the need, i.e. an instance of the space AT of all possible combinations of attribute-value pairs that describe the need (good, product or service).

Then, after issuing the CFP, the client agent selects the best proposal based on the evaluation of the trustworthiness of each proponent supplier, for the situation announced in the need, taking as input contracts the suppliers have established in the past and their respective outcome.

This way, we define $trust_{ac}(as) \in [0, 1]$ as the trustworthiness value of a trustee agent as , in the eye of the evaluator agent ac , as computed by a traditional trust aggregator engine; and *adequacy trust* $ad_{ac}(as, at) \in [0, 1]$ as a binary operator for situation-awareness purposes, where $ac \in AC$ is an agent from the set AC of client agents and $as \in AS$ is an agent from the set AS of supplier agents.

Therefore, the trustworthiness value of agent as as seen by agent ac in the specific context at is given by the following equation:

$$trust_{ac}(as, at) = trust_{ac}(as) * ad_{ac}(as, at) \quad (1)$$

This is the same as to say that, in a given moment, an agent may be qualified as trustworthy in a specific situation and as untrustworthy in a (maybe slightly) different situation.

Finally, a contractual evidence is represented by the tuple $\langle ac, as, at, t, o \rangle$, where t is the timestamp of the transaction that occurred between client agent ac and supplier agent as and $o \in \{true, false\}$ is the outcome of this transaction. Here, a *true* outcome means that agent as has succeeded in providing the good in the contractual terms and a *false* outcome means that as has violated the contract. Therefore, in our scenario, each supplier agent as has a history of its past contractual evidences.

3.2 The Contextual Fitness Component

The Contextual Fitness component is based on an online, incremental and flexible technique of extraction of behavior tendencies that we have developed. We have been testing different methods for extracting these tendencies from the historical set of agents' evidences, such as the use of the increase in frequency metric, as explained in [28]. Our current version, introduced in [29], uses the information gain metric. This is a metric used in the machine learning area (such as in the simple decision tree learning algorithm ID3 [30]) for classification purposes. It is typically used as an offline process, implying that the training and testing phases occur before the actual classification of new instances is performed.

The information gain metric is based on the entropy concept of information theory, and is defined as following:

$$Gain(S, A) \equiv Entropy(S) - \sum_{v \in Values(A)} \frac{|s_v|}{|S|} Entropy(s_v), \quad (2)$$

where $Gain(S, A)$ is the information gain of attribute A relative to a collection of samples S , $Values(A)$ is the set of all possible values for attribute A , and s_v is the subset of S for which attribute A has value v [30].

In our approach, we use this metric in a novel way, as it is used to dynamically learn a decision tree from the history of evidences of a given agent as, *every time* a client agent needs to verify the adequacy of as proposal to its current business need. In fact, we use all the evidences available about the supplier to build the decision tree, which normally consists of a dataset with some dozens of evidences, if that much. This means that no training or testing phases are performed. After

that, the failure tendencies of the agent in evaluation are extracted from the rules pointing to *false* outcomes. Figure 2 depicts a decision tree that was learned for a given supplier in a specific experiment we have run (we use the Weka API [31] in our simulations).

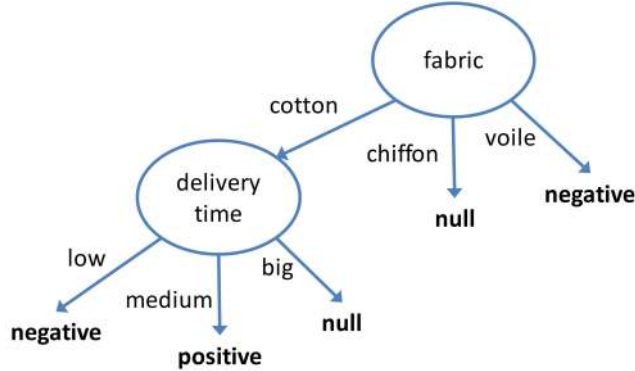


Fig. 2. Decision tree generated in our simulations

Concerning the tree above, our algorithm was able to identify that, at the time of this particular assessment, the supplier showed the failure tendencies (fabric = cotton, *, dtime = low) and (fabric = voile, *, *). Therefore, the trustworthiness value $trust_{ac}(as, at)$ of agent as , as given by Equation 1, would be zero if situation at matched any of the failure tendencies derived from the learned decision tree. Otherwise, the trustworthiness value of the target agent for the considered situation would be given by the $trust_{ac}(as)$ component of Equation 1.

For now, we are only exploring the *false* branch of the generated trees, which means that currently Contextual Fitness is a binary operator. However, we leave for future work the extraction of the *true* and *null* branches of the trees, in order to further distinguish between positive tendencies and uncertainty.

4 Experiments

We set up a multi-agent simulation scenario where business clients in the textile industry try to select the best suppliers of textile fabric, i.e. the ones that would maximize the utility of the clients. At every round of the simulation, each client issues a call for proposals (CFP) specifying a business need. This need is established randomly for each agent at setup time and is an instance $at \in AT$, as described in section 3.1. The clients keep their needs constant in all the experiment rounds. Moreover, all values of possible quantities and delivery times are fuzzified in the categorical values depicted in Table 1.

After issuing the CFP, every client selects the more adequate proposal for the current need based only on the trustworthiness value of all the proposals received, as computed by the trust model used by the client. In these experiments, we use three different trust models, including our situation-aware approach, in order to evaluate and compare the ability of each model in supporting better selection decisions.

Moreover, most of the suppliers that exist in our market have different handicaps on performing some particular aspect of a business transaction. For example, some suppliers tend to fail to deliver fabric in short delivery times, while others might fail to deliver high quantities of any fabric type. The aim of these experiments is to evaluate how effective are different trust models and correspondent aggregation techniques in detecting and acting upon these handicaps, in order to better assist the partners' selection decision. It is assumed that every supplier is able to provide any different type of fabric.

In this paper, we run the experiments with three different populations of suppliers, as described in the following sections. In the next section, we describe the generic configuration common to all the experiments.

4.1 Generic Configuration

Table 1 presents the configuration parameters that are common to all experiments.

Table 1. Configuration of the experiments

| Parameter | Value |
|-----------------------|---|
| Fabrics | Chiffon, Cotton, Voile |
| Quantities | Low, Medium, High |
| Delivery Time | Low, Medium, Big |
| # buyers | 20 |
| # of sellers | 50 |
| Seller stocks | 4 contracts per round |
| Types of sellers | Uniform distribution over the types considered in population (described later on) |
| # rounds | 60 |
| # runs per experiment | 15 |

Trust Models Evaluated. This section introduces the trust models evaluated in these experiments. The first model is the SA approach, which represents *SinAlpha* [25], a traditional, situation-less trust aggregation system that we have developed which uses properties of the dynamics of trust. In previous work [26], we thorough compared *SinAlpha* to a traditional statistical aggregation engine that uses weighted means by recency and we have shown that *SinAlpha* outperforms this model. In the present paper, we use the SA model in order to compare the behavior of both non situation and situation-aware approaches in open contextual environments.

The second model is the CS approach [5][22][23], which goes one step further traditional trust models by considering contextual aspects of the business in assessment. As we mentioned before in Section 2.1, it is a situation-aware technique that defines a context space as a multidimensional metric space with one dimension per each represented situation feature. It is able to estimate the trustworthiness values of agents in unanticipated situations using the similarity between different situations. We chose this approach to represent situation-aware trust models that use domain specific, predefined similarity metrics to predict unanticipated situations.

Concerning our implementation of this model in current experiments, we placed the reference contexts regularly over the combinations of all possible values of the contractual attributes. The distance function we used for attribute fabric is given in Equation 3. As can be observed, the distance is minimum (zero) if both contexts c_1 and c_2 have the same fabric and maximum (one) otherwise.

$$d^{fabric}(c_1, c_2) = \begin{cases} 0, & \text{if } fabric_1 = fabric_2, \\ 1, & \text{if } fabric_1 \neq fabric_2. \end{cases} \quad (3)$$

For the remaining attributes considered in the experiments (quantity and delivery time), the distance function is given in Equation 4.

$$d^{attr}(c_1, c_2) = |\ln(attr_1) - \ln(attr_2)|. \quad (4)$$

In the equation above, considering first the attribute *quantity*, $attr_i$ takes the value of 1, 3 or 5, depending on the value of the quantity is *low*, *medium* or *high*, respectively. In the same way, for attribute *delivery time*, $attr_i$ takes the value of 1, 3 or 5 for values of *low*, *medium* or *big*. The total distance between the two contexts is a weighted means of the three distances calculated above, with all dimensions equally weighted. Finally, the weight used to evaluate the relevance of a context c_1 accordingly to its similarity with context c_2 is given in Equation 5. All the remaining formulas needed to compute the trustworthiness scores of agents were implemented accordingly to one of the authors' paper [5].

$$w_i = e^{-d(c_1, c_2)}. \quad (5)$$

The last model to be evaluated is CF, our trust model that uses the Contextual Fitness technique described in section 3. In the experiments, we used SA as the aggregator component of the trust system (Figure 1). As with happens with the CS approach defined in the previous point, CF is a situation-aware trust model. It was designed to fit well to non parochial open market scenarios, where the available trust evidences for a particular partner agent might be scarce.

Populations of Suppliers. Table 1 characterizes the different types of suppliers used in the experiments.

All suppliers used in the experiments, excepting *RBad* agents, have a given handicap related to one or more contractual dimensions. For example, a supplier initialized with the *HQT* behavior (standing for 'Handicap in Quantity') has a

Table 2. Characterization of the populations of suppliers

| Supplier | Prob. Success | PopA | PopB | PopC |
|--|-------------------------------------|------|-------|------|
| HFab, HQt, HDt, HFabQt, HFabDt, HQtDt | 0.05 (handicap) 0.95 (otherwise) | 100% | 66.7% | |
| I-HFab, I-HQt, I-HDt, I-HFabQt, I-HFabDt, I-HQtDt | | | | 100% |
| RBad | 0.50 | | 33.3% | |

handicap in providing high quantities of any fabric; this way, if it is selected to a business transaction that involves the delivery of high quantities of fabric, it has a probability of 95% of failing the contract. In any other transaction that the supplier is involved, it has a probability of 5% of failing the contract. Besides HQT, we used five other types of behavior that represent five other different types of handicap: on a given fabric (*HFAB*), on low delivery times (*HDT*), on high quantities of a given fabric (*HFABQT*), on low delivery times for a given fabric (*HFABDT*), and on high quantities to be delivered in low delivery times (*HQTDT*). As happens in the example before, a supplier has a probability of 95% of violating a contract if the current CFP matches the supplier’s handicap, and 95% of probability of succeeding the contract otherwise.

Handicapped suppliers with name starting with ‘*I*’ are similar to the handicapped suppliers just described, excepting the fact that, with a probability of 66.6%, they change their handicap when they reach half of the total number of the rounds.

Finally, suppliers of type *RBad* have a probability of failing a contract of 50%, independently of the context of the contract established with the client.

Evaluation Metrics. We use three different metrics in order to evaluate how client agents tend to behave in terms of selection of partners – and how good is their decision on that – when using each one of the trust approaches in evaluation in scenarios that might involve scarcity of trust evidences. The first metric is the *average total utility* got by all clients in all rounds, given by the ratio of the number of successful contracts got by all clients in one round over the total number of contracts established in the round, averaged over all rounds of the experiments. The second metric is the *average utility per round*, which measures the same ratio but on a per round basis. Finally, the third metric is the *average number of different suppliers per round*, which measures how many distinct suppliers are selected by all the negotiating clients at every round, averaged over all rounds of the experiments.

4.2 Experiments with Population A

We started our experiments with a population of distinct handicapped suppliers, as can be observed from Table 2. These suppliers normally tend to fail in very specific situations (e.g. low delivery times), but perform well in all other

situations. Therefore, a good trust model should be able to detect the situations where these agents fail and to prevent their selection (only) in these specific situations.

Results. Figure 3 shows the average utility achieved by all clients in all rounds, with the results of population A being presented in the leftmost plot.




Fig. 3. Average total utility for populations A, B and C

As can be observed from the graph, both situation-aware approaches outperform SA in population A: SA fulfils in average 83.95% of the established contracts (with standard deviation of 2.87%), where CS succeeds in 86.01% of the contracts (sd. 4.73%) and CF gets the best results by succeeding in 90.57% of the contracts (sd. 1.34%).

The average utility per round and the average number of different suppliers selected per round are presented in Figure 4. As can be seen at the right graph, clients that use the CF approach start, since the first rounds, exploring a larger number of different suppliers than clients that use the other approaches, and they keep showing this behavior all over the rounds. The described behavior of the trust approaches seems to be related with the utility that clients achieve using these approaches, as suggested in the Figure (*left*). We observe that the most exploratory approach, CF, is the one that gets higher utility, at every round.

Another important result obtained with CF is that, after some quick learning at the first rounds, the number of succeeded contracts with this method is very close to the maximum of 19 (out of possible 20) contracts that can succeed per round, i.e. to the 95% probabilistic limit imposed in population A.

4.3 Experiments with Population B

In the second set of experiments, we wanted to verify if the promising good results of the CF model would maintain in an environment with fewer trustworthy suppliers. Therefore, we introduced population B, where one third of the available suppliers have erratic, non contextual behavior, succeeding only 50% of the contracts they enter, irrespective of the context of the contract (cf. Table 2).

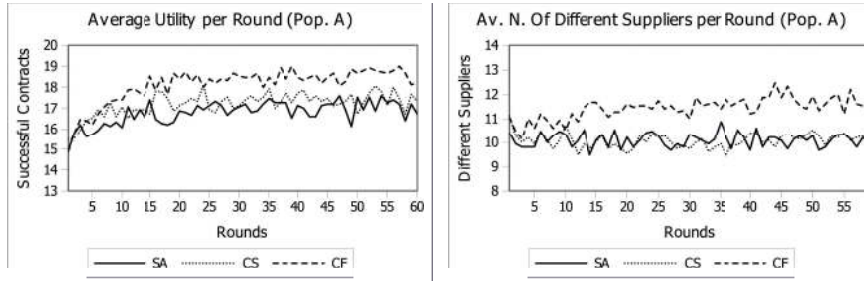


Fig. 4. Results per round for Population A

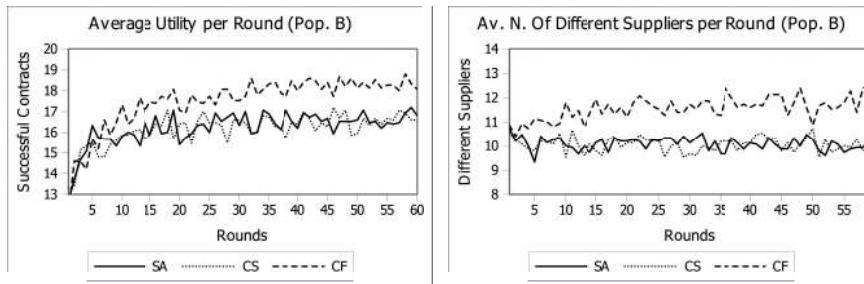


Fig. 5. Results per round for Population B

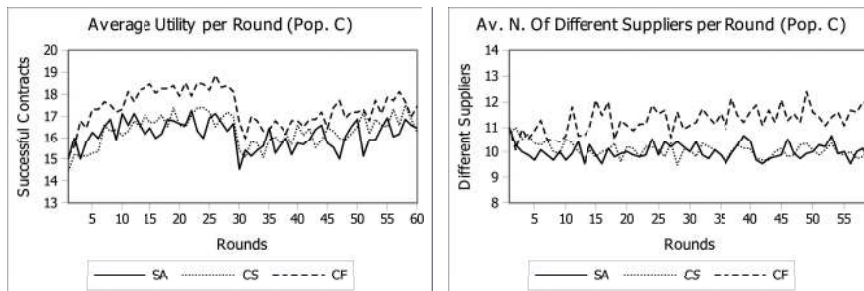


Fig. 6. Results per round for Population C

Results. Figure 3 (*central plot*) presents the results for population B concerning the average total utility achieved by the trust models in evaluation. As can be observed, there is a general decrease in utility for all approaches when compared to population A, as it would be expected, once the set of potential trustworthy suppliers is reduced in one third. In these experiments, SA clients succeeded in 81.08% of the contracts (sd. 2.03%), while CS was successful 80.73% of the times (sd. 4.38%) and CF succeeded 87.38% of the times (sd. 2.15%). In these experiments, the CS approach was significantly more penalized by the introduction of the *RBad* suppliers than the other two approaches, and its performance was as low as the performance of the situation-less approach.

Figure 5 plots the average utility per round (*left*) and the average number of different suppliers selected per round (*right*) for the three trust models in evaluation, for population B. The effect of introducing the *RBad* suppliers is visible in the very first rounds of the experiments, with all trust approaches obtaining less utility than in similar rounds of the experiments with population A. However, as happened in population A, we also observe here that the more exploratory behavior of CF allows clients using this approach to achieve higher utility associated with the choice of suppliers.

4.4 Experiments with Population C

In the last set of experiments, we wanted to evaluate how each trust model in evaluation reacts to a dynamic change of behavior in the suppliers' population. This way, we used the handicapped suppliers of population A, but they have now a 66.6% chance of changing their handicap at round 30 (half of the rounds considered in the experiments).

Results. The results for population C are shown in Figure 3 (*right*) and Figure 6. SA got 80.78% (sd. 4.23%) of average total utility and was the approach that had less suffered from the change in handicaps, as compared to the results obtained with population A. CS got 81.66% (sd. 2.63%) of utility and it was the approach more penalized by the reversal in the behavior of suppliers. Finally, CF still outperforms the other two approaches, obtaining 86.93% (sd. 2.32%) of utility, but it was slightly more penalized in relative terms by the change of the behavior of the suppliers than the traditional approach.

4.5 Discussion

We start the discussion of the results analyzing the relation between the average number of different suppliers selected per round and the average utility achieved per round, for each trust model evaluated. At a first sight, we could expect that an approach that explores more partners in the scenario described would have a smaller number of succeed contracts, at least in the first rounds of suppliers' exploration, where the partners are rather unknown. However, the results show that the CF approach does not perform worse than the remaining representative approaches at this first exploration phase and performs significantly better

than the others in the remaining steps of the experiments. This happens due to the CF capability to extracting tendencies of behavior even in the presence of a reduced number of available trust evidences, and to its capacity of doing that in an incremental and dynamic way. This way, when a client makes a bad decision concerning the exploration of a new supplier, the consideration of this new evidence (and corresponding *false* outcome) in a subsequent assessment will update the extracted tendency of failure of the supplier and approximate it to the true handicap of the supplier.

On the contrary, the SA approach tends to select the agents that have the highest values of trustworthiness until that date. As handicapped suppliers have the same probability of succeeding outside the context of their handicap, there is a strong chance that the first partners to be selected are the ones that incrementally increase their trustworthiness in the eye of the evaluator. This parochial strategy results in the undesirable behavior of keep choosing the same suppliers that occasionally fail the contracts for which they have a handicap for the current CFP, not giving a chance to explore other suppliers. In a way, this reflects what succeeds is real life subcontracting in the textile industry.

Concerning the situation-aware CS approach, we observe that it sensitively outperforms the SA approach in populations where all suppliers have some sort of handicap (A and C). This is due to the fact that the CS client reasons based not only on the global trustworthiness of the agent in assessment, but also on the behavior of the supplier in the context of current CFP. On the other hand, the introduction of one third of suppliers that fulfill or violate contracts regardless of the context of those contracts appear to reduce the effectiveness of this trust model based on reference contexts. Moreover, we can see from Figures 4, 5 and 6 that CS also behaves in a conservative way when selecting new partners, being outperformed by CF in terms of utility since the first rounds of the simulation. This is due to the fact that, although CS is a situation-aware technique, it functions in a rather different way than CF. In order to better understand the differences between CS and CF, we present next two examples taken from the experiments we have run.

In the first example we use the past contractual evidences of supplier *as* as depicted in Table 3. For the sake of readability, only the *at* and the *o* elements of the contractual evidence tuple (cf. Section 3.1) are presented in the table.

Table 3. Contractual evidences of supplier *as* (simplified)

| evd # | <i>at</i> | <i>o</i> |
|-------|-----------------------|----------|
| 1 | voile, low, medium | true |
| 2 | chiffon, low, low | true |
| 3 | chiffon, high, medium | false |
| 4 | voile, medium, medium | true |
| 5 | cotton, low, low | true |
| 6 | cotton, medium, big | true |
| 7 | voile, low, low | true |

Now, suppose the current business need of a client is the instance at_x : ($fabric = chiffon, quantity = high, dtime = medium$), and that the client needs to evaluate the proposal made by as . Using the CS approach, we verify that the context space of as is composed of 27 (3^3) reference contexts, and only seven of them directly correspond to the contexts of the past contractual evidences of the agent in evaluation. Let us also name the reference context corresponding to at_x as rc_x . Processing the first two evidences, they are located at distances $d_{1,x}$ and $d_{2,x}$ from rc_x , and then rc_x is updated one evidence at a time with the value of its outcome (*true*) weighted by a function of the distance, as given in Equation 5. This means that the trustworthiness of the supplier at rc_x is increased proportionally to the similarity between the reference context and the context of the evidence. The third evidence coincides with at_x and therefore the weight w_x of the evidence is maximum for rc_x , lowering the trustworthiness value of as at rc_x in a significant way as the outcome of this evidence is *false*. Finally, the processing of the last four evidences of the supplier, all positive, raise again the trustworthiness value at rc_x , even if attenuated by the distance of the context of each evidence to the context of rc_x . However, due to the scarcity of evidences about supplier as , its final trustworthiness score (strongly supported by the rc_x value) is still positive, and therefore bigger than the trustworthiness values of all other suppliers that have not yet been explored. This explains why, in these conditions, the approach has a tendency to select, from the set of the more fitted suppliers, the ones that have been involved in more contracts to date, acting in a rather parochial way. From our analysis, we can conclude that the interesting characteristic of bootstrapping of the CS approach can also be somewhat disappointing in open and dynamic environments where the available evidences for every agent can be scarce.

Now, let us look how our CF approach would behave in the same situation of example number one. Taking all available evidences of supplier as as depicted in Table 3, the Contextual Fitness technique would be able to extract the failure tendency ($*, quantity = high, *$). Applying Equation 1, the global trustworthiness of supplier as would be zero for current need at_x , as a failure match was detected, which means that the chances of this supplier being selected for current CFP would be low, allowing for the selection of a more adequate proposal, or even the exploration of a new partner.

At this point, it is necessary to say that a match between a supplier's handicap and the current situation in assessment does not exclude the supplier from the selection process, it just lowers its trustworthiness to zero. In the absence of better alternatives, the supplier can always be selected by the client to do business together.

We are now in conditions to analyze the second example. Table 4 illustrates an excerpt of the past contractual evidences of supplier as obtained from another set of experiments we have run using the CS model. In these experiments, supplier as is handicapped on low delivery times, which means that it has a high probability of failure in satisfying any ($*, *, dtime = low$) need.

Table 4. Contractual evidences of supplier *as* (simplified)

| evd # | <i>at</i> | <i>o</i> |
|-------|----------------------|----------|
| 1 | voile, medium, big | true |
| 2 | chiffon, high, big | true |
| 3 | cotton, low, low | false |
| 4 | cotton, medium, big | true |
| 5 | chiffon, high, big | true |
| 6 | cotton, high, medium | false |
| 7 | voile, high, low | false |
| 8 | voile, medium, low | false |
| 9 | chiffon, medium, big | true |
| 10 | voile, low, big | true |
| 11 | chiffon, high, big | true |
| 12 | voile, medium, big | true |
| 13 | chiffon, low, low | false |
| 14 | voile, medium, big | true |
| 15 | voile, high, low | false |

As can be observed from the table above, supplier *as* keeps being selected by clients that use the CS approach for business transactions for which it is not fitted due to its handicap on low delivery times. The problem here concerns the use of predefined similarity distances amongst contexts by the CS approach.

For instance, let us imagine that there is a new contractual evidence with $at_{16} : (fabric = cotton, quantity = low, dtime = medium)$ and $o_{16} = true$; also, let us focus on two specific reference contexts: rc_y , which matches $at_y : (fabric = cotton, quantity = low, dtime = low)$, and rc_z , corresponding to $at_z : (fabric = voile, quantity = low, dtime = medium)$. Using the values considered in Section 4 and equations 3 and 4, we verify that at_{16} is close enough from the reference context rc_y and consequently the trustworthiness of the supplier at rc_y would increase in a significant way with the consideration of this new evidence, regardless its handicap on low delivery times. On the other hand, the same new *true* evidence would increase the reference context rc_z in a less significant way, even though it corresponds to a context for which the supplier does not present a handicap. Figure 7 illustrates this scenario.

This example shows the limitations associated to situation-aware trust models that rely on predefined measures of similarity in defining how similar or different two contexts must be. More concretely, an agent that shows a good behavior in a context might fail in a *very similar* context and succeed in more *distant* contexts. Our study of the CS model in the proposed scenario gives us the strong belief that, even dedicating a team of experts to tune the distance functions, the use of predefined similarity functions, even those based on taxonomy-based similarity, can fail in adequately detecting the contextual subtleties exposed in this last example. Moreover, in the specific case of the CS approach, this hard tuning effort can be even compromised with the addition of more contextual dimensions.

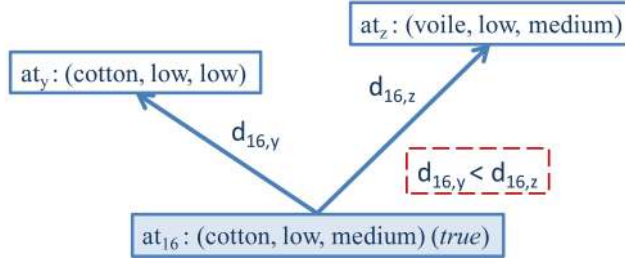


Fig. 7. Distances between a new evidence and two distinct reference contexts

For the sake of comparison, the application of our CF approach to the same scenario (cf. Table 3) after the third evidence would generate the tendency (fabric = cotton, *, *). In this case, CF would not be able to detect the true handicap of the supplier with the three available evidences and would even lower the probability of the supplier being selected to any contract involving *cotton*, which does not integrate its handicap. Using CF after the sixth evidence, our trust model would generate the tendencies (*, *, dtime = low) and (*, *, dtime = medium). In one hand, the CF client would have a high probability of wrongly missing the opportunity of transacting with this supplier in contracts stipulating *medium* delivery times. On the other hand, it would also have a high probability of not choosing the supplier in transactions that involve low delivery times (the supplier’s handicap), avoiding all the unsuccessful contracts that derived from the choices done by the CS clients.

In this last example we observed that CF can be sometimes too restrictive, by overfitting the existing evidences. For instance, with only six evidences, the trust approach detects the true handicap of the supplier but adds a new, non-existent failure tendency concerning medium quantities. We used population B in order to see if this characteristic would prevent the CF approach of doing good deals. By adding one third of bad suppliers, the choice space of the trust approaches reduced substantially and the utility of CF clients could be seriously jeopardized due to this overfitting-based generation of failure tendencies. However, the results obtained have shown that CF was less penalized than CS in this harsh scenario. Once again, the ability of the approach in dynamically rectifying the extracted tendencies every time there is a new evidence has shown to be a positive characteristic of the approach.

Finally, there are cases where the CF approach is not able to extract any failure tendency. This happens, for example, in the presence of RBad suppliers, which have quite an erratic and context less behavior. In these cases, our trust approach assesses the trustworthiness of the agent in evaluation taking into consideration only the value computed by the aggregator component (Equation 1). In the experiments, we used the same aggregator as the SA approach, which embeds the asymmetry property that confers more strength to negative outcomes and less strength to positive outcomes. Although both CF and SA performed better than CS in the presence of RBad suppliers, we need further work in order to

support our belief about the benefits of separating the aggregation part from the situational part of trust models.

5 Conclusions and Future Work

In this paper, we stated our strong belief that open and dynamic environments, such as global marketplaces where business agents risk new partnership seeking to increase their profit, need robust and situation-aware trust systems that better assist the selection of partners' decision. Then, we presented Contextual Fitness, our innovative situation-aware technique that can be used with any trust aggregation system and that is based on the dynamic extraction of tendencies of behavior of the agents in assessment. This is a different approach from the existing situation-aware trust models that rely on the predefinition of distance metrics to evaluate the similarity between contexts. We evaluated our approach and compared it with both a situation-less and a situation-aware approach that predefines the notion of similarity between contexts. The results have shown that our approach is more adequate than the other two in supporting the process of selection of partners when the candidate partners present some characteristic behavior (even if the behavior changes at some point of the agents' life) and when the evidences available about them are scarce. The results have also shown that the capacity of CF in rapidly correcting a bad decision supported by the incremental extraction of behavior tendencies allows the agents to be more exploratory without jeopardizing their overall utility.

As future work, we intend to further develop the method for extraction of behavior tendencies in order to allow the use of heterogeneous trust evidences. We also intend to further test the trust approach with different populations and different number of contextual dimensions.

Acknowledgments. This research is funded by FCT (Fundação para a Ciência e a Tecnologia) project PTDC/EIA-EIA/104420/2008. The first author enjoys a PhD grant with reference SFRH/BD/39070/2007 from FCT.

References

1. Sabater, J., Paolucci, M.: On Representation and Aggregation of Social Evaluations, in Computational Trust and Reputation Models. *Int. J. Approx. Reasoning* (2007)
2. Tavakolifard, M.: Situation-aware Trust Management, pp. 413–416 (2009)
3. Fabregues, A., Madrenas-Ciurana, J.: SRM: a tool for supplier performance. In: *AAMAS 2009*, pp. 1375–1376 (2009)
4. Hermoso, R., Billhardt, H., Ossowski, S.: Dynamic evolution of role taxonomies through multidimensional clustering in multiagent organizations. In: Yang, J.-J., Yokoo, M., Ito, T., Jin, Z., Scerri, P. (eds.) *PRIMA 2009. LNCS*, vol. 5925, pp. 587–594. Springer, Heidelberg (2009)

5. Rehak, M., Gregor, M., Pechoucek, M.: Multidimensional context representations for situational trust. In: DIS 2006: Proceedings of IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications, pp. 315–320, 383–388 (2006)
6. Nguyen, C., Camp, O.: Using Context Information to Improve Computation of Trust in Ad Hoc Networks. In: 2008 IEEE International Conference on Wireless & Mobile Computing, Networking & Communication (2008)
7. Neisse, R., Wegdam, M., van Sinderen, M., Lenzini, G.: Trust management model and architecture for context-aware service platforms. In: Chung, S. (ed.) OTM 2007, Part II. LNCS, vol. 4804, pp. 1803–1820. Springer, Heidelberg (2007)
8. Stolle, D.: In Trust in Society. In: Cook, K.S. (ed.) Russell Sage Foundation, pp. 202–244 (2001)
9. Macy, M.W., Sato, Y.: Trust, Cooperation, and Market Formation in the U.S. and Japan. Proceedings of the National Academy of Sciences of the United States of America 99(10)(suppl. 3) , 7214–7220 (2002)
10. Marsh, S.P.: Formalising Trust as a Computational Concept. PhD Thesis, University of Stirling (1994)
11. Dasgupta, P.: Trust as a Commodity. In: Gambeta, D. (ed.) Making and Breaking Cooperative Relations, Basil Blackwell (1990)
12. Dimitrakos, T.: System Models, e-Risk and e-Trust. Towards Bridging the Gap? In: Towards the E-Society: E-Business, E-Commerce, and E-Government, Kluwer, Dordrecht (2001)
13. Christianson, B., Harbison, W.S.: Why isn't trust transitive? In: Proceedings of the International Workshop on Security Protocols, London, UK, pp. 171–176 (1997)
14. Tavakolifard, M., Herrmann, P., Ozturk, P.: Analogical trust reasoning. In: Trust Management III, ch. 10, pp. 149–163 (2009)
15. Toivonen, S., Denker, G.: The Impact of Context on the Trustworthiness of Communication: An Ontological Approach. In: Proceedings of the ISWC Workshop on Trust, Security, and Reputation on the Semantic Web (2004)
16. Golbeck, J., Parsia, B., Hendler, J.: Trust networks on the semantic web. In: Klusch, M., Omicini, A., Ossowski, S., Laamanen, H. (eds.) CIA 2003. LNCS (LNAI), vol. 2782, pp. 238–249. Springer, Heidelberg (2003)
17. Tavakolifard, M., Knapskog, S.J., Herrmann, P.: Trust transferability among similar contexts. In: 4th ACM Symposium on QoS and Security for Wireless and Mobile Networks, New York, USA, pp. 91–97 (2008)
18. Strang, T., Linnhoff-Popien, C., Frank, K.: CoOL: A context ontology language to enable contextual interoperability. In: Stefani, J.-B., Demeure, I., Zhang, J. (eds.) DAIS 2003. LNCS, vol. 2893, pp. 236–247. Springer, Heidelberg (2003)
19. Jeh, G., Widom, J.: SimRank: a measure of structural-context similarity. In: Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 538–543. ACM Press, New York (2002)
20. Nakatsuji, M., Fujiwara, Y., Tanaka, A., Uchiyama, T., Ishida, T.: Recommendations Over Domain Specific User Graphs. In: 19th European Conference on Artificial Intelligence (ECAI 2010), Lisbon, Portugal (2010)
21. Jung, J.J.: Ontology-based context synchronization for ad hoc social collaborations. Knowledge-Based Systems 21(7), 573–580 (2008)
22. Reháč, M., Pěchouček, M.: Trust modeling with context representation and generalized identities. In: Klusch, M., Hindriks, K.V., Papazoglou, M.P., Sterling, L. (eds.) CIA 2007. LNCS (LNAI), vol. 4676, pp. 298–312. Springer, Heidelberg (2007)

23. Reháč, M., Pěchouček, M., Grill, M., Bartos, K.: Trust-based classifier combination for network anomaly detection. In: Klusch, M., Pěchouček, M., Polleres, A. (eds.) CIA 2008. LNCS (LNAI), vol. 5180, pp. 116–130. Springer, Heidelberg (2008)
24. Jonker, C.M., Treur, J.: Formal analysis of models for the dynamics of trust based on experiences. In: Garijo, F.J., Boman, M. (eds.) MAAMAW 1999. LNCS, vol. 1647, pp. 221–231. Springer, Heidelberg (1999)
25. Urbano, J., Rocha, A.P., Oliveira, E.: Computing confidence values: Does trust dynamics matter? In: Lopes, L.S., Lau, N., Mariano, P., Rocha, L.M. (eds.) EPIA 2009. LNCS, vol. 5816, pp. 520–531. Springer, Heidelberg (2009)
26. Danek, A., Urbano, J., Rocha, A.P., Oliveira, E.: Engaging the dynamics of trust in computational trust and reputation systems. In: Jędrzejowicz, P., Nguyen, N.T., Howlet, R.J., Jain, L.C. (eds.) KES-AMSTA 2010. LNCS, vol. 6070, pp. 22–31. Springer, Heidelberg (2010)
27. Thurman, R.: Purchasing contracts in the textile industry: how can purchasing contracts decrease subcontracting risks and reclamations in a textile industry company?, Tampereen ammattikorkeakoulu, Thesis Report (2007)
28. Urbano, J., Rocha, A.P., Oliveira, E.: Trust estimation using contextual fitness. In: Jędrzejowicz, P., Nguyen, N.T., Howlet, R.J., Jain, L.C. (eds.) KES-AMSTA 2010. LNCS, vol. 6070, pp. 42–51. Springer, Heidelberg (2010)
29. Urbano, J., Rocha, A.P., Oliveira, E.: Trustworthiness Tendency Incremental Extraction Using Information Gain. In: IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, wi-iat 2010, vol. 2, pp. 411–414 (2010)
30. Quinlan, J.R.: Induction of Decision Trees. *Machine Learning* 1(1), 81–106 (1986)
31. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The WEKA Data Mining Software: An Update. *SIGKDD Explorations* 11(1) (2009)