

REFINING THE TRUSTWORTHINESS ASSESSMENT OF SUPPLIERS THROUGH EXTRACTION OF STEREOTYPES

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*

Keywords: Situation-aware Trust, Dynamics of Trust, Multi-agent Systems.

Abstract: Trust management is nowadays considered a promising enabler technology to extend the automation of the supply chain to the search, evaluation and selection of suppliers located world-wide. Current agent-based Computational Trust and Reputation (CTR) systems concern the representation, dissemination and aggregation of trust evidences for trustworthiness assessment, and some recent proposals are moving towards situation-aware solutions that allow the estimation of trust when the information about a given supplier is scarce or even null. However, these enhanced, situation-aware proposals rely on ontology-like techniques that are not fine grained enough to detect light, but relevant, tendencies on supplier's behaviour. In this paper, we propose a technique that allows the extraction of positive and negative tendencies of suppliers in the fulfilment of established contracts. This technique can be used with any of the existing "traditional" CTR systems, improving their ability in selectively selecting a partner based on the characteristics of the situation in evaluation. In this paper, we test our proposal using an aggregation engine that embeds important properties of the dynamics of trust building.

1 INTRODUCTION

Several technologies are being studied and applied in the general process of computerized supply chain management. Computational trust management is one such technology that will allow extending electronic sourcing to world-wide located, non registered and probably unknown business partners. With this technology, a business entity will be able to search the suppliers offer space and to filter the ones that are fitted to the entity current needs, in a scale of the size of the Internet.

The first generation of CTR systems addressed the representation and the aggregation of trust evidences into trustworthiness scores for evaluating trustees, and most of these proposals are based on some sort of statistical aggregation methods (e.g. Ramchurn, Sierra, Godo and Jennings (2004), Sabater (2003), Jøsang and Ismail (2002), Zacharia and Maes (2000), Erete, Ferguson and Sen (2008), and Huynh, Jennings and Shadbolt (2006)). Other works proposed more sophisticated engines that considers the dynamics of trust in the computation of confidence scores, in theoretical and practical terms (e.g. Elofson (1998), Falcone and

Castelfranchi (1998), Jonker and Treur (1999), Marsh and Briggs (2008), and Melaye and Demazeau (2005)). However, none of the current computational trust approaches are mature enough to be themselves trusted by real managers.

Trying to cope with this question, trust community is moving towards a second generation of models that explore the situation of the trust assessment in order to improve its credibility, also allowing for the estimation of trustworthiness values when trust evidences on the trustee partner are scarce or even null. However, few proposals have been made on this specific area (see Tavakolifard (2009), Neisse, Wegdam, Sinderen and Lenzi (2009), Rehak, Gregor and Pechoucek (2006), Fabregues and Madrenas-Ciurana (2009), and Hermoso, Billhardt and Ossowski (2009)).

The purpose of this paper is two-folded. First, we describe our proposal for an aggregation engine that embeds three fundamental dynamics of trust properties – asymmetry, maturity, and distinguishable past –, and present our conclusions about the relevance of the inclusion of such properties in trust aggregation engines. Then, we propose a situation-aware technique that allows the

extraction of tendencies in the behaviour of agents. This technique allows, for instance, to detect whether a given supplier has a tendency to fail or to succeed contracts that are *similar* to the current business need (e.g. in terms of good, quantity and delivery time conditions). We performed experiments that show that this technique enhances traditional CTR systems by bringing context into the loop; i.e. it not only concerns if a given supplier is generally trusted good or bad, but if it is trusted good or bad *in the specific contractual situation*. Also, this approach differs from other recently proposed situation-aware proposals in the way that it does not imply the use of hierarchical-based structures (e.g. ontology) and is able to detect fine-grain subtle dissimilarities in *related* situations.

Although we contextualized the use of our trust system in the sourcing/procurement part of the supply chain, agent-based trust and reputation systems are of general interest in many other domains (for instance, general business, psychology, social simulation, system resources' management, etc), and apply to all social and business areas of the society where trust is deemed of vital importance.

The remaining of this paper is structured as follows: Section 2 describes our study about the relevance of considering properties of the dynamics of trust in the aggregation engine of CTR systems. Section 3 describes the stereotype-based technique we developed in order to complement traditional CTR engines with situation-aware functionality. Section 4 presents the experiments we run in order to evaluate the proposed situation-aware technique, and Section 5 concludes the paper.

2 USING TRUST DYNAMICS IN THE AGGREGATION ENGINE

In previous work, we described an S-like aggregation curve (see Figure 1) that allows for an expressive representation of the dynamics of trust, particularly, implementing the following properties:

- *Asymmetry* property, that stipulates that trust is hard to gain and easy to lose;
- *Maturity* property, that measures the maturity phase of the partner considering its trustworthiness, where the slope of growth can be different in different stages of the partner trustworthiness;
- *Distinguishably* property, that distinguishes between possible different patterns of past behaviour.

The trustworthiness estimation of a given supplier agent using this curve implies a slow growth upon positive results when the partner is not yet trustable, an acceleration when it is acquiring confidence, and a slow decay when the partner is considered trustable (i.e., in the top right third of the curve), allowing for the definition of three different trust maturity phases (the Maturity property). The decrease movement upon negative results follows the same logic, although the mathematical formula subjacent to the curve includes parameter λ that permits that trust grows slower and decays faster (the Asymmetry property).

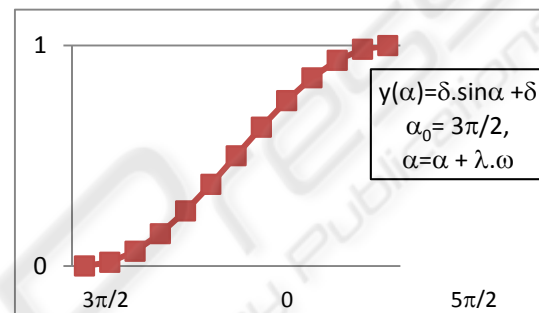


Figure 1: The S-like curve.

One can argue that we could use other S-like curves instead of a sin-based one, such as the Sigmoid curve. However, we intuitively feel that a Sigmoid curve permits a probably too soft penalisation of partners that proved to be trustable but that failed the last n contracts. This can happen accidentally (e.g. due to an unexpected shortage of good or to distribution problems), but it is also described in the literature as a typical behaviour of deceptive provider agents, who tend to build up a trustworthy image using simple contracts and then violate bigger contracts exploring the acquired trustworthiness.

2.1 Evaluation of Trust Properties

Previous work provides a detailed description of the S-like curve, as well as an experimental evaluation of its behaviour. In this section, we summarize the main conclusions we obtained when we experimentally compared it (thereafter called the S approach) to a weighted mean by recency approach (that we named WMEAN), a common approach seen in literature for traditional CTR aggregating engines (cf. Huynh, Jennings and Shadbolt (2006)).

In this work, we explored three different scenarios. In the first scenario, we wanted to compare the capacity of both approaches in differentiating between different types of supplier

agents, namely, the capacity of primarily choosing ‘good’ suppliers that with a high probability do not violate a contract. In such a scenario, we observed that the S approach outperforms the WMEAN approach in its capacity of selecting ‘good’ partner agents, in one hand, and in avoiding ‘bad’ partners, in the other hand. One difference between both approaches resides in the fact that in S all the historical path is taken into account in the process of trust construction, and partners have to accumulate several good experiences in the past until they are able to get an average to high trust score (the maturity property). In opposition, the WMEAN approach allows the selection of partners with fewer past events. For instance, analyzing the traces of the experiments, we verified that some bad choices of WMEAN happened when the algorithm selected partners with rather few contractual past evidences (e.g. the pattern of the previous evidences to the time of selection where $V-F-F-V-F-F$, where V means a violated contract and F a fulfilled contract).

Another difference between the two approaches is due to the *asymmetry* property of S. This seems to be particularly important when identifying and acting upon partners that show intermittent behaviour (e.g. $F-V-V-F-V-F-V-F-F-F-F$). This last pattern of behaviour is indeed severely punished by the S approach, where violations weight more than fulfillments (therefore penalizing undesirable intermittent patterns), and where the last few positive evidences are not sufficient to ‘push’ the confidence level of the partner to the second third of the curve.

In the second scenario, we intended to study how S and WMEAN react in the presence of extreme partners that have a bursty-like behaviour (i.e. that switch between sequences of good and deceptive behaviour). By analysis of the traces of the experiments, we realized that both approaches act quite differently as they tend to select different partners in similar conditions. The main point to consider here is that WMEAN, by privileging recency, actually assigns high trust levels to candidate partners that systematically behaved deceptively in the past, had no classification for a long time, and then got one positive classification in the present. I.e., WMEAN-like approaches can *forgive too fast* in certain temporal scenarios. One could argue here that this forgiveness issue is solved by increasing the size of the window used (i.e. the number of the last past evidences considered); however, in our experiments we found it hard to select the optimal window size, as it deeply depends on the frequency of the contracts (historical

evidences) made in the past. The forgiveness question does not apply to S, due to the action of the *maturity* property; however, we realized that S has a somewhat bigger tendency to enter a burst of deceptive behaviour and that it can be slower in penalizing good partners immediately after they invert their behaviour.

Finally, the last scenario intended to study the *abuse of prior information* scenario defined in (Zacharia and Maes, 2000), where ‘good’ partners definitely invert their behaviour after a given number of iterations. The results that we obtained showed that S outperforms WMEAN in detecting and penalizing the change of behaviour of originally ‘good’ partners, while WMEAN showed a significantly higher tendency to choose ‘bad’ partners than S.

2.2 Remarks about the S Curve

Taking into account all the experiments performed, we can conclude that the three *properties* of the dynamics of trust embedded in S are effective in distinguishing between different types of target agents, therefore in detecting and acting upon undesirable agents’ behaviours. Namely, the asymmetry property penalizes intermittent behaviour, the maturity property avoids selection of partners who did not prove to be trustable enough, and the distinguishable past property avoids the phenomenon of forgiveness described above. Considering this last property, we have a somewhat different view than the one presented in (Sabater, Paolucci and Conte, 2006), where the authors state that the aggregation of evaluations shall not depend on the order in which these evaluations are aggregated.

In these experiments, we could not evaluate, however, the potential full benefits of the curve shape against simpler curves that do show similar trust dynamics properties (e.g. curves with linear shape). In fact, as stated previously, S considers different growth/decay slopes in different stages of the trustworthiness acquisition of a target agent, and it also presents a sigmoid-like shape. The choice of this shape was based on the concept of *the hysteresis of trust and betrayal*, from Straker (2008). In this work, the author proposes a path in the form of a hysteresis curve where trust and betrayal happens in the balance between the trustworthiness of a self and the trust placed on the self. The S curve simplifies the hysteresis approach by using just one curve for both trust and betrayal representation and considering three different growth/decay stages:

Creating Trust (first third of the curve), Trust is Given (second third of the curve), and Taking Advantage (last third of the curve).

Performance tests of the S representation against a simpler curve were performed. This new simpler curve uses λ and ω parameters from S (cf. Figure 1) to update the trustworthiness value of target agents, but it lacks the softness round curve at Creating Trust and Taking Advantage extremes. The results of these experiments show similar performance of both curves in the tested scenarios. Therefore, we conclude that we need different, much more complex models of target population to further study the impact of the sigmoid-like shape of S on its capability of distinguishing between partners. We leave this topic for future work.

3 THE PROPOSED SYSTEM

3.1 Motivation for Situational Trust

Computational trust estimations help the trustier agent to predict how well a given candidate partner will execute a task and to compare between several candidate partners. However, there are some questions that a real-world manager would pose before making a decision that cannot be answered by simply aggregating available trust evidences into trust and reputation values. These questions involve somehow a certain level of *intuition*. We propose to first analyze three scenarios that might occur in real world business and that would help to understand this concept.

In the first scenario, an agent may decide to exclude from selection a candidate partner with which it had never entailed business before but that it knows that rarely fails a contract, just because the agent intuitively fears that this partnership would not be successful. For example, a high tech company may fear to select a partner from a country of origin without high technology tradition, even though this partner has proved high quality work in the desired task in the recent past. We call this situation the *intuitive fear*. For this scenario, it would be desirable that the selector agent could reason taking into account additional contextual information about the characteristics of the entity represented by the candidate agent. For instance, the presence of key figures such as the annual turnover or the number of employees of the entity would allow the selector agent to better know the entity. Also, the establishment of argumentation between both parties is a real-world procedure that could be automated

into the computational decision process. We address the *intuitive fear* situation in future work.

In the second scenario, the agent may decide to exclude from selection a candidate partner that is currently entering the business, for which there is not trust and/or reputation information yet. This scenario deals with the problem of newcomers, for which there is no information about prior performance, and we name it *absence of knowledge*. (Huynh et al., 2006) suggest that in these cases the use of recommendations and institutional roles could be useful to start considering newcomers in the selection process. Although we do not address this situation in this paper, we propose here to use conceptual clustering of entities' characteristics in order to generate profiles of business entities. In a second step, the profile of the newcomer is compared with the profiles of business entities for which there is some trust information and an estimation of the newcomer trustworthiness is inferred. This approach implies that the characteristics of the business entities are available, which is a reasonable assumption for centralized virtual market places and virtual organizations built upon electronic institutions, and might also be applied to more decentralized approaches by transmission of this kind of entities' knowledge between communicating agents.

Finally, in the third scenario, the selector entity knows that a candidate partner is well reputed in fulfilling agreements in a given role and context (e.g. selling cotton zippers to European countries), but it is afraid that the candidate is not able to provide high quantities of the material in a short period of time. We name this situation the *contextual ignorance*. In this scenario, the evaluator agent knows that the candidate partner is trustworthy in a particular business scenario, or even that is generally trustworthy, but needs to know how well it would adapt to a different type of business. In this section, we address this question by presenting a description of our situation-aware technique, a component complementary to the CTR aggregation engine that is intended to give extra information to the trustier agent by computing a value of *how well the candidate partner fits in the selector current needs*, as defined in the issued call for proposals (CFP).

3.2 The Situation-aware Technique

Every time a client issues a CFP, it may receive several proposals from suppliers. In order to select the best proposal, the client (trustier) computes a general trustworthiness score for each supplier/

proposal. This score evaluates not only the general behaviour of the supplier (trustee), but also the adequacy of its *business behavioural profile* to the CFP in question. The algorithm for computing the trustworthiness value of a trustee is given next:

```

computeTrust (CFP, evds)
begin
    negEvds ← getNegEvds (evds)
    strtype ← getStrtype (negEvds)
    trustScr ← compare (CFP, strtype)
    if trustScr is 1
        trustScr ← computeTrust (evds)
    end
end
    
```

The algorithm above starts by evaluating the fitness of the received proposal in relation to current CFP. At line 3, all the evidences of the trustee that have a negative outcome are put in one class. This negative outcome can represent, for example, the past transactions of the trustee that triggered relevant contractual sanctions, although the meaning of such outcome can be established by each individual client agent. Then, at line 4, a stereotype is extracted for this class, which means that the *most significant contractual characteristics* of the evidences in the class are going to be extracted. Depending on the degree of the required extent of frequency increase (cf. parameter α in equation 1) and on the evidence set of the trustee, it is possible that the algorithm does not return any stereotype. We must refer here that this is an online process that is repeated every time a new trust assessment is performed, which allows to capture the variability of the behavior of trustee agents at any time.

At line 5, the stereotype extracted (if any) is compared to the current business need (CFP). A match between stereotype and CFP attributes means that the supplier/trustee has a tendency to fail this type of contracts, and therefore the comparison function returns a zero value, that would be the final trustworthiness score for the trustee proposal. Otherwise, there is no evident signal that the supplier is inapt to perform the current transaction, and its final trustworthiness score is computed using the S approach described previously, or any other CTR ‘traditional’ system (lines 6-7).

As can be seen above, we simplified the proposed situation-aware technique by using (by now) just a negative class of the evidences of the trustee. The use of the positive class and the use of distinct degrees of fitness will allow refining our algorithm and this constitutes ongoing work.

We further describe the management of stereotypes in the next sections. First, in order to clarify the overall process, we describe how

contractual information (CFP and trustee evidences) is represented in current implementation.

3.2.1 Representation of Information

In current implementation, contractual information is represented by the tuple $\langle A_c, A_s, At_1..At_n, t, o \rangle$, where:

- $A_c \in C$ is an agent from the set C of clients’ agents (i.e., the trustier agent);
- $A_s \in S$ is an agent from the set S of suppliers’ agents (i.e., the trustee agent);
- $At_i \in AT$ is the value of an attribute from the set AT of n contract attributes (e.g. good, quantity and delivery time);
- t is the timestamp of the transaction. Although we are not using this value in current implementation, it is needed in aggregation systems that weights evidences by their recency;
- $o \in \{T, F\}$ is the outcome of the contract, either representing successful (*true*) or violated (*false*) contracts by the supplier.

For instances, the evidence $\langle A_i, A_j, \text{cotton}, 360000, 7, t, \text{false} \rangle$ means that agent A_i contractualized with A_j at time t the acquisition of 360000 meters of cotton to be delivered within 7 days, and that A_j failed to deliver the product in the aforementioned conditions.

We must refer here that prior to stereotype extraction, all evidence attributes are quantified to categories or quantitative values using a fuzzy approach. In current implementation, both quantity and delivery time values are quantified to *low*, *medium* and *high* categories.

Concerning the representation of trust, we represent the trustworthiness score of trustee agent A_j relative to the current business situation as $\text{trust}(A_j, \text{context}_k) \in [0, 1]$, where context_k is an instance of the context space Ctx constituted by all possible combinations of the fuzzyfied values of the attributes in set AT .

3.2.2 Stereotype Management

The extraction of stereotypes from the evidences contained in each class is done using the metric of equation 1 that measures the increase in the frequency of an attribute within the community. This metric was proposed in (Paliouras, Karkaletsis, Papatheodorou and Pyropoulos, 1999).

$$\alpha = \left(\frac{\#InstAttClass}{\#InstClass} \right)^2 - \left(\frac{\#InstAttTotal}{\#InstTotal} \right)^2 \quad (1)$$

In the equation above, $\#InstAttClass$ is the number of times that a given attribute appears in the class, $\#InstClass$ is the total number of evidences in the class, $\#InstAttTotal$ is the number of times that the attribute appears in all classes, and $\#InstTotal$ is the total number of evidences kept for the trustee. As mentioned before, parameter α is the degree of the required extent of frequency increase, and determines the granularity of stereotype extraction.

Finally, the comparison between a stereotype and the current CFP is done attribute by attribute. Figure 2 illustrates a CFP request $\langle A_i, ?, chiffon, 1080000, 7, t, ? \rangle$ from agent A_i , whose quantity and delivery time values are quantified into *high* and *low* values, respectively, and the negative stereotype extracted for an hypothetic proposal of agent X . The stereotype means that, whatever fabric and quantity is considered, agent X has a tendency to fail contracts with low delivery time. Therefore, a match is detected and the trustworthiness score of X for the current proposal is set to zero.

Stereotype: Agent X, null, null, low, false
CFP: chiffon, 1080000, 7

Figure 2: Examples of a CFP and a stereotype.

4 EXPERIMENTS

In order to evaluate the benefits of the proposed situation-aware technique (hereafter named SAT), we run a series of experiments where a traditional CTR aggregation engine – represented by the S approach – was compared to the global solution constituted by both S and SAT components.

4.1 Experimental Testbed and Methodology

We run all the experiments in the Repast tool (<http://repast.sourceforge.net>). The experiments simulated a virtual textile marketplace, where at every round every client agents post buying leads (in the form of call for proposals) discriminating a fabric to buy and correspondent quantity and delivery time, and supplier agents propose in response to these leads if they have the described quantity of the fabric. Table 1 presents the configuration options for the experiments.

In these experiments, we wanted to evaluate if the situation-aware technique would improve the ability of the trust

Table 1: Configuration of experiments.

Fabrics	{Chiffon, Cotton; Voile}
Quantities	{Low, Medium, High}
Delivery Time	{Low, Medium, Big}
# buyers	20
# of sellers	50
Types of sellers	Chosen upon a uniform distribution over the types {"S _{HQT} ", "S _{HDT} ", "S _{HFB} "}
# rounds / # runs	100 / 40
α threshold	0.25

system in selecting partners taking into account the current business needs. Therefore, we run the same experiment using, first, just the S component, and then the global solution of S plus SAT. We used the utility criterion to compare both approaches: in each round, the utility of a client agent was 1 if the contract done in this round is successful and 0 if the contract was violated. Therefore, the best approach is the one that gets the higher average utility of all clients in all rounds, i.e. the one that is more efficient in selecting the best partners for every CFP attributes at any time.

We also used a specific population of suppliers constituted by three different types of suppliers, each one showing some kind of handicap in fulfilling a contract, as shown in Table 2.

Table 2: Different types of Suppliers.

Type	Description
S _{HQT}	Probabilistically succeeds 95% of the established contracts, except the ones that involve the delivery of <i>high quantities</i> , which probabilistic fails 95% of the time
S _{HDT}	Probabilistically succeeds 95% of the established contracts, except the ones where the <i>delivery time is low</i> , which probabilistic fails 95% of the time
S _{HFB}	Probabilistically succeeds 95% of the established contracts, except the ones that involves the delivery of a <i>given fabric</i> , which probabilistic fails 95% of the time

For example, a S_{HQT} supplier would have a handicap in providing the service if the quantity to provide is high. Therefore, the best approach in evaluation is the one that is more capable of detecting, and reacting to, these types of handicaps.

4.2 Results

In every experiment, we measured the number of successful contracts per type of target agents and per approach, and averaged this number over the total

number of rounds. In the best case possible, each client is able to identify the handicap of every provider and to select the best proposal, leading to an average of 95% of successful contracts (Table 2).

The results obtained show that the SAT approach gets, in average, 85.21% of utility, which means that it is less than 10% away from the theoretical best result. Also, the traditional approach gets in average 77.82% of utility, performing relevantly poorly than its situation-aware counterpart.

Figure 3 shows the average number of successful contracts per round, including the trendlines for the traditional approach (above) and for the situation-aware approach (below).

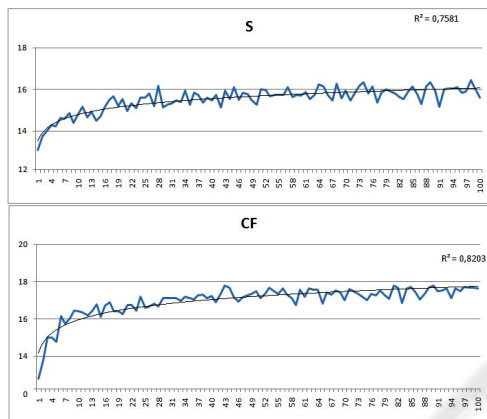


Figure 3: Average number of successful contracts per round for S (above) and SAT (below) approaches.

We can observe from the figures above that, although simple, the SAT algorithm is able to extract correct stereotypes for each trustee agent with a few number of past contractual evidences, for the experimented population. This is a very important issue in several domains, such as in the textile industry, where direct or even indirect evaluations of a given supplier might be scarce.

Also, by analyzing the traces of the experiments, we observed that the learning curve for the situation-aware solution is consistently more evident than the one of the traditional approach. Figure 4 plots the relative number of unsuccessful contracts per type of supplier for S (above) and SAT (below), obtained in one run of the experiments. From several of these observations, we verified that, no matter what the first choices were concerning the initial selection of partners, the situation-aware solution often succeeds in stabilizing the selection of all three types of suppliers with low values of violated contracts.

On the other hand, the traditional approach keeps selecting suppliers based only on the aggregate trust

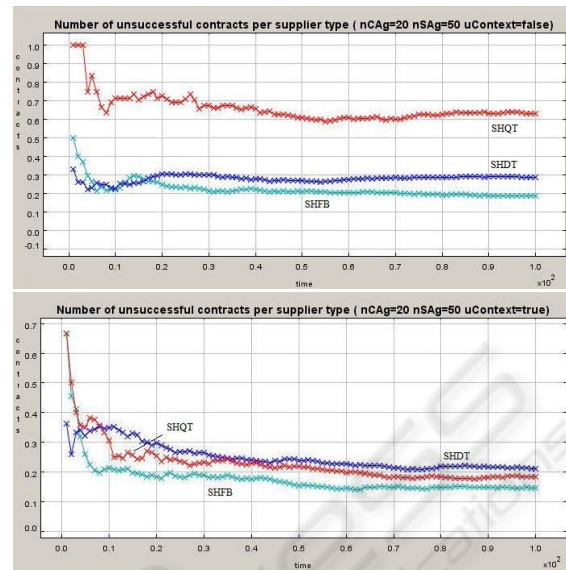


Figure 4: Relative number of unsuccessful contracts per supplier type without (above) and with (below) the situation-aware technique.

score. As the three types of suppliers have equal probability of failing (and therefore approximate trustworthiness) if the analysis of the context is excluded, they can be equally chosen for a given CFP independently of their specific handicap.

In fact, looking at Figure 4, we observe that suppliers with a handicap on quantity suffered from a cold start, most probably because they were initially selected to provide high quantities of material. As S is not able to capture the handicaps – and as suppliers of this type would tend to succeed on all other contracts they are engaged to, therefore maintaining some level of trustworthiness – the algorithm will continue to select suppliers with quantity handicap to provide high quantities of material.

4.3 Interpretation of the Results

The results obtained show that the traditional approach of aggregating trust, even when enhanced with heuristics based on the dynamics of trust, presents limitations when the suppliers in evaluation present tendencies of failure. In these situations, the extraction of behaviour stereotypes is effective in discriminating the best suppliers to be chosen for any particular business need. Also, this technique showed to be effective since the first rounds of the experiments, meaning that it is adequate to situations where the available number of trust evidences about the supplier in evaluation is scarce.

5 CONCLUSIONS

In this paper, we presented a simple situation-aware technique (SAT) based on the extraction of stereotypes of agents' behaviour that can be used with any traditional CTR system in order to enhance the estimation of trustworthiness scores. Although other situation-aware approaches are now being proposed in the trust management field, the SAT technique presents some benefits: i) it is simple and can be used with any of the existing CTR 'traditional' aggregation engines; ii) it is an online process, meaning that it captures the variability in the trustee behaviour as it happens; iii) it does not rely on ontology-based situation representation, and therefore the extraction of the similarity between the situation in assessment and the past evidences of trustee agent does not require specific, domain-based similarity functions; also, it allows for fine-grain *dissimilarity* detection (e.g. it distinguishes between the *similar though different* situations of providing one container of cotton in 7 or in 14 days).

The SAT approach was evaluated using a traditional aggregation engine approach enhanced by the inclusion of properties of the dynamics of trust. Although these properties showed to be beneficial, we conclude that the study of the benefits of a sinusoidal like shape that follows Straker (2008) work on the area of Psychology needs proper data/models concerning the behaviour of real-world organizations; therefore, we will address the acquisition of such data sets in future work.

ACKNOWLEDGEMENTS

The first author enjoys a PhD grant with reference SFRH/BD/39070/2007 from the Portuguese Fundação para a Ciência e a Tecnologia.

REFERENCES

- Elofson, G., 1998. Developing Trust with Intelligent Agents: An Exploratory Study. In *Proceedings of the First International Workshop on Trust*, pp. 125-139.
- Erete, I., Ferguson, E., Sen, S., 2008. Learning task-specific trust decisions. In *Procs. 7th Int. J. Conf. on Autonomous Agents and Multiagent Systems*, vol. 3.
- Fabregues, A., Madrenas-Ciurana, J., 2009. SRM: a tool for supplier performance. In *AAMAS'09*, 1375-1376.
- Falcone, R., Castelfranchi, C., 1998. Principles of trust for MAS: cognitive anatomy, social importance, and quantification. In *Procs. Int. Conference on Multi-Agent Systems*.
- Hermoso, R., Billhardt, H., Ossowski, S., 2009. Dynamic evolution of role taxonomies through multidimensional clustering in multiagent organizations. In *Principles of Practice in Multi-Agent Systems*, vol. 5925, chapter 45, pp. 587-594.
- Huynh, T. D., Jennings, N.R., Shadbolt, N.R., 2006. An integrated trust and reputation model for open multi-agent systems. In *Autonomous Agents and Multi-Agent Systems*, Vol. 13, N. 2, September 2006, pp. 119-15.
- Jonker, C. M., Treur, J., 1999. Formal Analysis of Models for the Dynamics of Trust Based on Experiences. In *Procs. of the 9th European Workshop on Modelling Autonomous Agents in Multi-Agent World: Multiagent System Engineering*. F. J. Garijo and M. Boman. LNCS, vol. 1647. Springer-Verlag, London, 221-231.
- Josang, A., Ismail, R., 2002. The Beta Reputation System. In *Proceedings of the 15th Bled Electronic Commerce Conference*, Sloven.
- Marsh, S., Briggs, P., 2008. Examining Trust, Forgiveness and Regret as Computational Concepts. *Computing with Social Trust*. Springer, ed. J. Golbeck, pp. 9-43
- Melaye, D., Demazeau, Y., 2005. Bayesian Dynamic Trust Model. *CEEMAS 2005*: 480-489.
- Neisse, R., Wegdam, M., Sinderen, M., Lenzini, G., 2009. Trust management model and architecture for context-aware service platforms. In *On the Move to Meaningful Internet Systems*. LNCS, pp. 1803-1820.
- Paliouras, G., Karkaletsis V., Papatheodorou, C., Pyropoulos, C. D., 1999. Exploiting Learning Techniques for the Acquisition of User Stereotypes and Communities. In *Procs. of UM99*.
- Ramchurn, S., Sierra, C., Godo, L., Jennings, N.R., 2004. Devising a trust model for multi-agent interactions using confidence and reputation. In *Int. J. Applied Artificial Intelligence* (18) 833-852.
- Rehak, M., Gregor, M., Pechoucek, M., 2006. Multidimensional context representations for situational trust. In *IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications*, pp. 315-320.
- Sabater, J., 2003. Trust and Reputation for Agent Societies. Number 20 in *Monografies de l'institut d'investigacio en intelligenca artificial*. IIIA-CSIC.
- Sabater, J., Paolucci, M., Conte, R., 2006. Repage: Reputation and image among limited autonomous partners. In *Journal of artificial societies and social simulation*, Vol. 9, pp. 3.
- Straker, D., 2008. *Changing Minds: in Detail*. Syque Press.
- Tavakolifard, M., 2009. Situation-aware trust management. In *RecSys '09: Proceedings of the third ACM conference on Recommender systems*, 413-416.
- Zacharia, G., Maes, P., 2000. Trust management through reputation mechanisms. In *Applied Artificial Intelligence*, 14(9), 881-908.