Engaging the Dynamics of Trust in Computational Trust and Reputation Systems

Agnieszka Danek, Joana Urbano, Ana Paula Rocha, Eugénio Oliveira

Faculdade de Engenharia da Universidade do Porto — DEI, Rua Dr. Roberto Frias, 4200-465 Porto, Portugal
LIACC — Laboratory of Artificial Intelligence and Computer Science at University of Porto
{adanek, joana.urbano, arocha, eco}@fe.up.pt

Abstract. Computational Trust and Reputation (CTR) systems are essential in electronic commerce to encourage interactions and suppress deceptive behaviours. This paper focus on comparing two different kinds of approaches to evaluate the trustworthiness of suppliers. One is based on calculating the weighted mean of past results. The second one applies basic properties of the dynamics of trust. Different scenarios are investigated, including a more problematic one that results from introducing newcomers during the simulation. Experimental results presented in this paper prove the benefits of engaging properties of the dynamics of trust in CRT systems, as it noticeably improves the process of business partners' selection and increases the utility.

1 Introduction

Computational Trust and Reputation (CTR) systems, which are capable of collecting trust information about candidate partners and of computing the trust scores for each of them, are a necessity in virtual marketplaces. They provide a way to discourage suppliers from being deceptive, if only the punishment for cheating is severe enough.

Even though some simple practical examples of CTR systems do already exist (e.g. eBay.com, Amazon.com), there are still many doubts and open issues in this area. One aspect is to understand better the social component of trust and reputation building process. Another great field of research is focused on finding the most suitable way to represent and aggregate social evaluation and existing evidences in a form of trust and/or reputation scores that could be used as a recommendation in selection of partners in virtual business. These models range from arithmetic and weighted means ([1], [2], [3]) to Beta ([4]) and Dirichlet distributions ([5]), Bayesian approaches ([6], [7]), and trust learning approaches ([8], [9]). One of the recent branches of researches is the consideration whether engaging some properties of the dynamics of trust (that follows the psychological and sociological research in the field of trust acquiring/loosing) in the computation of the trust score can lead to the improvement in the estimation of trustworthiness of business partners.

The motivation of this paper is to compare different types of trust aggregation engines proposed in the literature. We are particularly interested in two important kinds of aggregation approaches. The first considers properties of the dynamics of trust in the process of trust building ([10], [11]). Although few proposals present practical implementation of these properties, we think that their consideration allows to improve the trustworthiness assessment of target agents. The second kind of aggregation engines is based on a weighted mean of a number of the past trust evidences. The reason why we consider this simpler class of approaches is that it is employed in numerous trust and reputation models. The weights are based on different sources, e.g. coherence of reputation ([4]) or confidence in the trust evidence ([2], ([12]). The most common way used, investigated also in this paper, is to give more significance to the more recent evidences ([2], [12]). Behaviour of these different approaches is here evaluated in different scenarios, including those when newcomers appear during experiments.

The remaining of this paper is structured as follows: section 2 presents the approaches evaluated together with short comments on them, their properties and parameters. Section 3 presents the experiments performed in order to evaluate and compare the described models. Finally, section 4 concludes the paper and presents future work.

2 CTR Systems Analysed

In the experiments described in this paper, five different approaches for evaluation of partners trustworthiness were explored. Three of them, *SinAlpha*, *Line* and *JaT* engage the dynamics of trust in the trust aggregation formula. The *WMean* approach is based on calculation of weighted mean of past results. The last approach, *Random* selects the partners randomly, without using any trust estimation mechanism. It was used to explore the suppliers population.

2.1 Approaches Using The Dynamics of Trust

The *SinAlpha* approach, characterised in [10], uses trigonometric formula (recency factor is omitted) presented in (1) to estimate a truth score T for each available supplier after n trust evidences.

$$T(\alpha_n) = 0.5 \cdot sin(\alpha_n) + 0.5,\ \alpha_n = \alpha_{n-1} + \lambda \cdot \omega,\ \alpha_0 = \frac{3}{2}\pi\,. \tag{1}$$

The shape of *SinAlpha* curve was introduced to reflect some dynamics of trust. It applies different growth/decay slopes in different stages of the trustworthiness acquisition of a target agent and was chosen to follow the idea of *the hysteresis of trust and betrayal* [13]. The *SinAlpha* curve reflects some properties of the dynamics of trust: the *asymmetry* property (trust is hard to gain and easy to lose, modelled with positive and negative increments, λ_+ and λ_-), the *maturity* property (avoiding selection with few evidences, i.e. trust gaining is slow, as results from the size of pace of growth ω) and the *distinguishability* property (past pattern of behaviour has an influence on the future trust). Hence described properties are also present in a simple line, the question is whether using a line curve (which is an easier model) instead of the *SinAlpha* curve will lead to the same or similar results.

The *Line* approach, designed particularly for the purpose of this paper, is a representation of such a simplified case, where trigonometric curve is replaced by a line with modified, more straightforward range of argument values (see (2)). Its parameters were chosen to be consistent with parameters of the *SinAlpha*.

$$T_{n+1} = T_n + \lambda \cdot \omega,\ T_0 = 0\,. \tag{2}$$

The *JaT* approach uses the trust update schema, as described in [11]. As for the dynamics of trust, similarly to the previous models it implements the *asymmetry*, *maturity* and *distinguishability* properties. Although it does not actually follow the described *hysteresis of trust and betrayal*, it applies, however, different growth/decay slopes in different stages of the trust acquisition (see (3), δ^+ and δ^- are positive and negative impact factors, related by the endowment coefficient e, where $\delta^+ = \delta^- \cdot e$).

$$T_{n+1} = \begin{cases} (1-\delta^+) \cdot T_n + \delta^+ & \text{if the contract was fulfilled,} \\ (1-\delta^-) \cdot T_n & \text{if the contract was violated.} \end{cases} \tag{3}$$

2.2 Approach Not Using The Dynamics of Trust

The *WMean* approach uses an aggregation engine that computes the mean of last N results weighted by the recency of these results. It lacks the properties of dynamics of trust. It is based on a recency function used by *FIRE* model [12], being an enhancement of *Regret*'s method of calculating rating recency [2]. Equation (4) presents its formula, where c is a contract, $res(c)$ is its outcome and $w(c)$ is a weight for that outcome. Despite of introduced changes, still the recency scaling factor λ, because of its design, has to be adjusted in each testing environment to the currently used time unit. It is a big disadvantage of this approach. For example, to obtain the same behaviour when the time of contract is measured in days and when it is measured in months, different values of λ should be used.

$$T(N) = \frac{\sum_{i=0}^{N-1} w(c_i) \cdot res(c_i)}{\sum_{i=0}^{N-1} w(c_i)},\ w(c_i) = e^{-\frac{\Delta t(c_i)}{\lambda}},\ \lambda = -\frac{d}{ln(0.5)}\,. \tag{4}$$

2.3 Analysis of the Parameters of Different Models

The behaviour of different approaches described above depends largely on the values of their parameters. Extensive experiments were performed in order to find the best, most suitable values of these parameters that would lead to the most consistent, good results in different scenarios. The best values of the parameters were used in the main experiments presented in this paper and can be found in Table 1.

The obtained results revealed that in the *SinAlpha* model violated contracts should be penalised little more than described in previous works. It reveals the importance of penalties for deceptive suppliers. The values of the parameters of the *Line* model were not investigated and chosen to be in consistency with the ones for the *SinAlpha*. Concerning the *JaT* approach, we realised that tuning of negative impact factor δ_- is more relevant than tuning of endowment coefficient e. It demonstrates how important is to set the proper penalty for violated contract. Indeed, in the performed simulations, small values of δ_- lead to better results. It is because large values of δ_- cause the model to react too rapidly to one violated contract, decreasing the trust too largely, not in agree with the common sense. Finally, as for the *WMean* approach, it is very sensitive for changes of λ. A small value of parameter d, characterising how old contracts will have meaningful weights in calculation of the trust score, causes that very recent contracts have a large influence on the trust score. This, in turn, can promote intermittent behaviour and forgive fast about supplier's previous failures. On the other hand, large value of parameter d makes *WMean* approach very close to calculating a simple mean of the recent outcomes, what also has its disadvantages — not take into account the trend of the supplier's behaviour. Overall, it turned out that in tested scenarios larger values of parameter d, characterising how old contracts will have meaningful weights in calculation of the trust score, lead to better results. Tests performed also showed that the size of the time window does not have much influence on the results and it is hard to find the optimal size (working consistently good in different scenarios), as it largely depends on the frequency of the historical evidences.

3 Experiments

To evaluate described approaches for partners' trustworthiness estimation, a series of experiments was run where all models were compared to each other.

3.1 Experimental Testbed and Methodology

All simulations were run using Repast agent simulator [14]. In the experiments, a virtual marketplace were simulated, where customer agents post call for proposals defining desired goods to purchase and supplier agents propose in response to these calls. Table 1 presents complete configuration options for the experiments, including parameters of different approaches. In presented simulations all suppliers were able to provide to all customers. In order to evaluate customers (different trust models) behaviour three distinct population of suppliers: A, B and C were used. The populations vary by types of suppliers available, as presented in Table 2. The specific number of suppliers of the particular type in the single experiment is a result of a uniform distribution over the types available in the population. A capability of supplier in fulfilling the contract is modelled by a Markovian process with two states (1 for contract fulfilment, 0 for contract violation) and transition probabilities $P11$ (fulfilment-to-fulfilment) and

$P01$ (violation-to-fulfilment). Probabilities for *GoodBad* supplier change after half of the experiment (two values separated by an arrow).

Table 1. Configuration of the experiments

Parameter	Value
#customers	2 *JaT*, 2 *SinAlpha*, 2 *Line*, 2 *WMean*, 2 *Random*
#suppliers	16 (uniform distribution over all types available in the population)
#rounds	100
#runs	100
JaT parameters	$\delta^{-} = 0.1$, $e = 0.9$
SinAlpha parameters	$\lambda_{+} = 1$, $\lambda_{-} = -3$, $\omega = \frac{\pi}{12}$, $\alpha \in < \frac{3}{2}\pi, \frac{5}{2}\pi >$
Line parameters	$\lambda_{+} = 1$, $\lambda_{-} = -3$, $\omega = \frac{1}{12}$, $T \in < 0, 1 >$
WMean parameters	$\lambda = -\frac{d}{ln(0.5)}$ $(d = 100)$, $N = 10$

Table 2. Types of suppliers together with values of their transition probabilities

Supplier	Probability	Value	Pop A	Pop B	Pop C
Good	P11	0.9	√	–	√
	P01	1.0			
Fair	P11	0.8	√	√	√
	P01	0.75			
Bad	P11	0.5	√	√	√
	P01	0.5			
GoodBad	P11	0.9→0.5	–	√	√
	P01	1.0→0.5			
Bursty	P11	0.9	–	–	√
	P01	0.2			

Additionally, another group of experiments were performed, where apart from suppliers from the considered population a set of newcomers were introduces in the middle of the experiment. This set consisted of two suppliers of each of the types: *NewGood*, *NewFair* and *NewBad*, which transition probabilities are the same as of suppliers *Good*, *Fair* and *Bad*, respectively. All other configurations of the experiments remained unchanged. The aim of these experiments was to explore the behaviour of different models when some new suppliers (i.e. without any prior trust information) appear during the simulation.

3.2 Results

Two different metrics were used to evaluate the performance of the approaches examined. The first one is the utility gained by each type of customer (trust model). It is the average of utilities of all customer of the considered type, where utility of a customer is measured as the ratio of all contracts enacted by this customer that where successful over all contracts enacted by that customer. The second one is the percentage of each type of supplier chosen throughout the experiment by each trust model. All presented results are the averages from all runs of experiments.

The results of the first group of experiments are depicted in Fig. 1. They include the utility gained and the percentage of different types of suppliers chosen by each trust model, depending on the employed population of suppliers (A, B or C). It can be observed that the results obtained with *JaT*, *SinAlpha* and *Line* are similarly good in all three populations, with *JaT* performing slightly better than other two and all noticeably better than *WMean*, in both metrics. In fact, all three models are efficient in distinguishing between different types of suppliers, avoiding picking up many *Bad* suppliers and choosing noticeably more *Good* suppliers than the remaining two approaches. The *SinAlpha* and the *Line* approach in population C choose little more *Bursty* suppliers than the *JaT* approach, but much less than the *WMean* model, which chooses them more than two times more often.

The results of the second group of experiments, with three different populations of suppliers employed and newcomers introduced during the simulation, are depicted in Fig. 2. In particular, the results contain the utility gained and the percentage of different types of suppliers chosen (including newcomers) by each trust model, depending on the employed population of suppliers (A, B or C). As in the previous experiments, *JaT*, *SinAlpha* and *Line* approaches lead to similar, good results. In the populations A and C, the *JaT* is slightly better than the other two, while in the population B the *SinAlpha* results in the best utility gained. All three approaches outperforms the *WMean* model in all populations. It can be also observed that the *WMean* approach is significantly worse than the first three models mentioned in distinguishing between different types of suppliers and, especially, in avoiding the bad ones.

Concerning newcomers, they have the least probability to be chosen with the *JaT* approach and little greater with *SinAlpha* and *Line* approaches. The *WMean* model gives noticeably more chances to newcomers to start interactions than the three previously mentioned models. It was observed during the experiments and is visible in the results — the percentage of newcomers chosen by the *WMean* model is higher than in case of other examined models, for all populations employed (see Fig. 2, for suppliers types *NewGood*, *NewFair* and *NewBad*). The decision whether or not giving many chances to newcomers is a good strategy depends on the specific marketplace and its application.

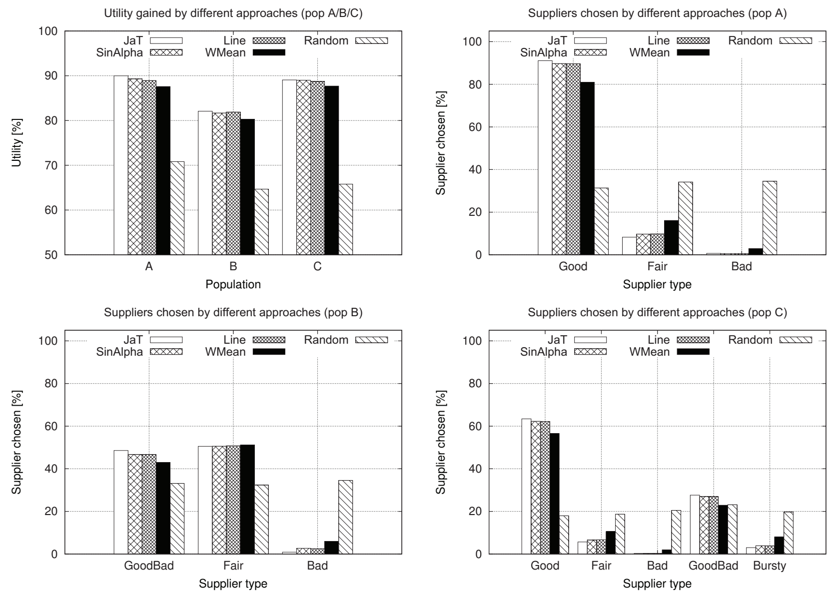

Fig. 1. Results of simulations performed in the populations A, B and C

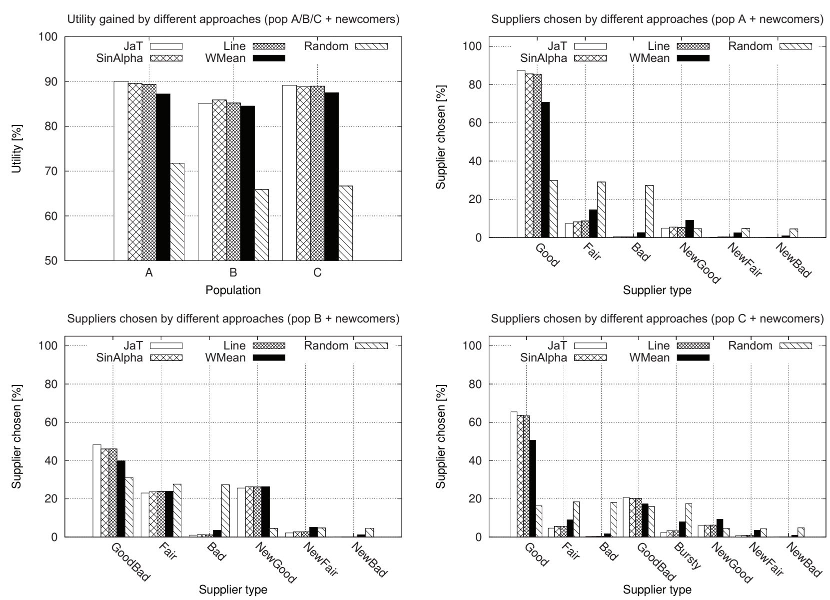

Fig. 2. Results of simulations performed in the populations A, B and C with newcomers

3.3 Interpretation of the Results

The results show that in all tested scenarios trust aggregation engine based on the weighted mean is significantly less efficient in the utility gained, in distinguishing between different suppliers' types and in avoiding choosing the *Bad* ones, than the models that engage the dynamics of trust.

The main limitation of the *WMean* in the simulations performed is the lack of the *maturity* feature, what cause this approach to often choose the suppliers basing on a small amount of evidences. For example, a newcomer with one fulfilled contract has a maximum value of the trust score, while an experienced, good supplier with many contracts fulfilled in the past will have less trust score if it violated even one of the previous N contracts. On the other hand, this feature of the *WMean* gives more chances for newcomers to start working in an already established population. Nevertheless, it does not result in a better utility, as it is connected with the risk of trusting inexperienced suppliers. Moreover, absence of the *asymmetry* property in the *WMean* approach allows intermittent behaviour, not penalising enough for past violated contracts. This cause many *Bursty* supplier types to be chosen by this approach in the population C. Furthermore the *WMean* approach does not have the *distinguishability* property, so it cannot recognise different patterns of past behaviour, relying only on the time of enacting of the last contracts. It can lead to assigning high trust score to the supplier with many violated contracts in the past and only one contract fulfilled very recently.

Differences in results of the three aggregation engines that use the aforementioned properties of the dynamics of trust are very small, much smaller than standard deviations. The JaT approach chooses slightly more *Good* suppliers, giving less chances to contract to other types of suppliers in comparison to *SinAlpha* and *Line* models. It is also visible in case of newcomers, which are chosen least frequently by the JaT approach. It is because in the picked configuration of the experiments, the *SinAlpha* and *Line* approach requires less fulfilled contracts in a row to give a high value of trust than the JaT and penalise more for a violated contract. Thus, it is a possibility for *Fair* suppliers and the newcomers to gather the trust score higher than the mature suppliers with large trust scores, in case of their occasional failure. In most of the scenarios, the JaT approach leads to the best utilities gained. However, in case of population B with newcomers, giving more chances to the inexperienced suppliers is an advantage of *SinAlpha* and *Line* (as suppliers *NewGood* are the best available in the second half of the simulation) and results in better utilities of these models than the utility gained by the JaT.

Obtained results revealed also that in the performed simulations the *SinAlpha* approach does not take any advantage of the shape of its curve and its behaviour is close to the behaviour of much simpler *Line* model, based on the line shape. The thorough study of the slope of the *SinAlpha* and its influence on the obtained results would require a different, more complex type of experimental scenarios.

4 Concluding Remarks and Future Work

This paper presents an evaluation of different kinds of aggregation engines used to estimate the trust in business partners. Performed experiments revealed that, in different scenarios, approaches that incorporate fundamental properties of the dynamics of trust are noticeably better in distinguishing between different types of suppliers and especially avoiding *Bad* and *Bursty* ones than the approach based on the weighted mean of the past results. This behaviour is reflected in higher utilities gained by aggregation engines of this type. It proves that properties of the dynamics of trust are actually significant in the process of evaluating the trust score in the CTR systems.

As for newcomers, they have the biggest chances to come into being with the *WMean* approach, as it is more liable to trust based on the few past evidences, what in fact promotes the new suppliers with little positive past. In examined scenarios this behaviour leads to lower utility than the ones obtained with other models. Extra experiments performed showed that changing the number of newcomers does not have much influence on the utilities gained by different approaches.

Differences in results obtained with *JaT*, *SinAlpha* and *Line* approaches are small and a special shape of the *SinAlpha* curve does not seem to give any benefits in the described experimental scenarios. In the future we plan to investigate further the properties of *the hysteresis of trust and betrayal*, to take advantage of this psychological background of trust transitions.

Overall, this paper reveals that engaging the dynamics of trust in CRT systems improves the process of business partners' selection and increases the utility gained. Several models that use the weighted mean with recency can take advantage of applying a very easy approach that incorporates the dynamics of trust, i.e. *JaT*. This change should both reduce the required computations and improve the obtained results (utilities).

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

References

1. Ramchurn, S., Sierra, C., Godo, L., Jennings, N.R.: Devising a trust model for multi-agent interactions using confidence and reputation. Int. J. Applied Artificial Intelligence 18, 833852 (2004)
2. Sabater, J.: Trust and Reputation for Agent Societies. Number 20 in Monografies de linstitut dinvestigacio en intelligència artificial. IIIA-CSIC (2003)
3. Carbo, J., Molina, J., Davila, J.: A BDI Agent Architecture for Reasoning about Reputation. In: IEEE Int. Conf. on Systems, Man, and Cybernetics (2001)
4. Jøsang, A., Ismail, R.: The Beta Reputation System. In: 15th Bled Electronic Commerce Conference, Sloven (2002)

5. Reece, S., Rogers, A., Roberts, S., Jennings, N.R.: A Multi-Dimensional Trust Model for Heterogeneous Contract Observations. In: 22nd AAAI Conf. on Artificial Intelligence (2007)
6. Zacharia, G., Maes, P.: Trust management through reputation mechanisms. Applied Artificial Intelligence 14(9), 881908 (2000)
7. Haller, J.: A Bayesian Reputation System for Virtual Organizations. In: Dagstuhl Seminar Proceedings 06461, Negotiation and Market Engineering (2006)
8. Erete, I., Ferguson, E., Sen, S.: Learning task-specific trust decisions. In: 7th Int. J. Conf. on Autonomous Agents and Multiagent Systems, Portugal, vol. 3 (2008)
9. Smith, M., desJardins, M.: Learning to trust in the Competence and Commitment of Agents. Autonomous Agent Multi-Agent Systems 18, 3682 (2009)
10. Urbano, J., Rocha, A.P., Oliveira, E.: Computing Confidence Values: Does Trust Dynamics Matter? In: L. Seabra Lopes, N. Lau, P. Mariano & L. M. Rocha (eds.), Progress in Artificial Intelligence — 14th Portuguese Conference on Artificial Intelligence (EPIA 2009), Springer, pp. 520-531, Aveiro, Portugal (2009)
11. Bosse, T., Jonker, C.M., Treur, J., and Tykhonov, D: Formal Analysis of Trust Dynamics in Human and Software Agent Experiments. In: 11th Int. Workshop CIA 2007 on Cooperative Information Agents, Springer-Verlag, Lecture Notes in AI, vol. 4676, pp.343-359, Delft, The Netherlands (2007)
12. Huynh, T.D., Jennings, N.R., Shadbolt, N.R.: An integrated trust and reputation model for open multi-agent systems. Autonomous Agents and Multi-Agent Systems 13(2), 119-154 (2006)
13. Straker, D.: Changing Minds: in Detail. Syque Press (2008)
14. Repast, `http://repast.sourceforge.net/repast_3`