Multidimensional Context Representations for Situational Trust

Martin Rehák, Miloš Gregor and Michal Pěchouček
Department of Cybernetics and Center for Applied Cybernetics, Czech Technical University in Prague
Technická 2, Prague, 166 27, Czech Republic
mrehak@labe.felk.cvut.cz, pechouc@labe.felk.cvut.cz

Abstract

We propose a generic model for situation representation in agent trust models. As we argue that general trust models are inappropriate in complex environments and their use for advanced decision-making is limited due to the lack of detail, we introduce a generic framework for context representation in trust model. Presented method is independent of the trustfulness representation formalism in the underlying trust model. It doesn't prescribe any particular form for the set of reference contexts that holds the trustworthiness information in the context space: we define and evaluate the baseline approach with regular grid and discuss future extension towards more adaptive approaches using octant trees and fuzzy clustering techniques.

1 Introduction

This contribution addresses the need to efficiently represent complex contexts in the trust model, both for the observations (or impressions) and the current decision. The need to represent the context is common for all agent models that attempt to model a situational trust: trust in an agent in a particular situation [5]. To successfully model situational trust, we have to model the situation – represented by context – first.

A crucial problem to address is a selection of the mapping between the *situation* in the environment and the *context*: a formal representation of the situation in our model. Obvious problem to address is a selection of relevant situation attributes to include into the context, as well as issues of scale and precision. Obviously, successful situation model has to contain all the relevant attributes that influence the trusting decision. On the other hand, adding more characteristics than necessary makes it difficult to determine the similarity between two situations, assuming that the trusting situations are rarely completely identical. Our context model addresses these concerns by resorting to standard approaches from mathematics and artificial intelligence, that are well adopted for the task. The main modelling problems we want to address are:

- **computational** and **memory efficiency**: we need to devise a model that doesn't require storage of all observations and their processing for each trusting decision;
- **learning efficiency**: the model must use the observations efficiently, to allow fast bootstrapping even in very big context spaces.

2 Formal Model

We denote agents in the community $A, B, C, ...$, forming a set $Agents$. General trustfulness (without taking the context into account) of agent B as perceived by agent A is denoted $\Theta_A(B)$. Context space $\mathbb{C}$ is a metric space of contexts with distance function $d(c_1, c_2)$. The context of each trust observation $\tau_A(X|c_i)$ or trusting decision $\delta_A(X|c_i)$ (where A is the observing/deciding agent, X denotes the evaluated agent and c_i is a context of this observation/decision) represented by exactly one point in this space, but several different observations may fall into the same point c_i in $\mathbb{C}$.

2.1 Context Space

In general, we define the metric space $\mathbb{C}$ in several steps:

1. Identify all relevant features of the environment.

2. Define the Q-dimensional context space where each dimension q matches a relevant feature of the trusting environment.

3. For each dimension q, define its quantification (either discreet or continuous) and appropriate distance metric d^q that correctly represents the feature.

4. Define a joint metric d on the full space $\mathbb{C}$, taking into considerations the domain characteristics and marginal metrics d^i[1].

There are 4 basic properties of any distance function $d : \mathbb{C} \times \mathbb{C} \rightarrow R$ that we need to respect while we define our own, domain dependent distance metrics:

[1]Alternatively, we may proceed without defining a distance metric for each dimension, if we define the overall metric directly.

1. non-negativity:
$$d(c_1, c_2) \geq 0 \quad (1)$$
2. symmetry:
$$d(c_1, c_2) = d(c_2, c_1) \quad (2)$$
3. zero distance ⇔ identity:
$$d(c_1, c_2) = 0 \Leftrightarrow c_1 = c_2 \quad (3)$$
4. triangle inequality:
$$d(c_1, c_3) \leq d(c_1, c_2) + d(c_2, c_3) \quad (4)$$

Obviously, all marginal distance metrics d^q shall respect these requirements as well.

To combine the marginal distances into the d function, we will typically chose one of the special types of Minkowski distance:

$$d(c_1, c_2) = (\sum_{q=1}^{Q} |c_1^q - c_2^q|^p)^{\frac{1}{p}} \quad (5)$$

For many practical purposes, we chose the values of p to be 1, defining so called Manhattan distance that adds the marginal distances of each dimension, or 2 to defines an Euclidean distance, or we pose $p \to \infty$, obtaining Chebyshev distance defined as a maximum of marginal distances.

For the purpose of determining the weights in our model, we need to define a (possibly domain dependent) function $f : \mathbb{R}^{0+} \to [0, 1]$, where the $\mathbb{R}^{0+}$ is an output range of the distance metric function. Naturally, we require f to be *non-increasing* on this interval – points that are farther away shall have smaller or at most identical weight compared to the points in proximity.

2.1.1 Example

To illustrate the abstract notions of metric space $\mathbb{C}$, we introduce an example of such space for a specific logistics scenario. We model the trust reasoning of a humanitarian aid organization agent that acquires transportation services from several local transporters after major disaster.

In our (simplified) domain, we model each trusting situation (either observation or decision) by three parameters: cargo type, cargo size and road quality. Cargo *type* defines the product we transport: medical supplies, food or durable goods. Each cargo type has specific handling requirements – medical supplies are the most sensitive to carry, while the durables require less care. *Size* of the transport is simply a quantity to carry, while the *road quality* represents the quality of the roads to use for transport. It is interesting to note that *type* dimension is discrete, while the *size* and *road quality* are real-valued variables, but varying greatly in their scale: one has an absolute scale (size), while the other will be close to 1.

The context space $\mathbb{C}$ is therefore three dimensional, with one discreet dimension and two continuous ones. The next step is a definition of marginal distances d^q for each dimension. In the *type* domain, we place our products on a "sensitivity" scale: medical supplies require most attention: 5, with the food in the middle: 1 and the durables as least sensitive ones, with 0.2 value[2]. Our type distance metrics is defined as follows, using the product values defined above:

$$d^{type}(c_1, c_2) = |ln(type_1) - ln(type_2)| \quad (6)$$

In the size domain, the metric shall describe the similarity between two contracts in terms of their relative size. We propose a measure

$$d^{size}(c_1, c_2) = |ln(size_1) - ln(size_2)| \quad (7)$$

The logarithmic relation captures an intuitive notion of ratio: 10 tons difference between two 20 and 30 ton transports is much more important than the same difference between two shipments of thousands of tons.

We apply the same reasoning for the road quality dimension:

$$d^{road}(c_1, c_2) = |ln(qual_1) - ln(qual_2)| \quad (8)$$

Then we combine the above metrics using a slightly modified (weighted) "Manhattan distance":

$$d(c_1, c_2) = \alpha_1 d^{type}(c_1, c_2) + \alpha_2 d^{size}(c_1, c_2) + \alpha_3 d^{road}(c_1, c_2) \quad (9)$$

2.2 General Algorithm

In the context space, we define one or more *reference contexts* r_i, forming a set $\mathcal{R}$ in $\mathbb{C}$. For each point r_i and each partner agent X we (agent A) keep trustfulness estimate denoted $\Theta_A(X|r_i)$. This value can be a result of application of any relevant trust model – for example Regret [9], FIRE [4] or other [7, 1, 8], provided that the inputs to this model are weighted or selected to reflect the performance of agent X in the specific context r_i.

Therefore, each observation $\tau_A(X|c_o)$ is used to update the trustfulness of reference contexts $r_i$with the weights determined using the general formula:

$$w_i = f(d(c_d, r_i)) \quad (10)$$

, where f is a non-increasing function on $[0, +\infty)$ as defined above. This function represents the decay of the observation usefulness with increasing distance d of the particular reference context r_i – obviously, it is most useful when its distance $d(c_d, r_i)$ from the reference context is zero and

[2] Inverting the scale will not change the result thanks to the distance symmetry stated in Eq. 2.

it shall decrease with increasing distance. This function, together with the metric, is a part of the domain description. For example, in our simple experiment presented in Section 3.1 we use a simple form of weight function defined as $w_i = e^{-d(c_d, r_i)}$.

Generally speaking, we integrate the new observation $\tau_A(X|c_o)$ into the apriori trustfulness evaluation $\Theta_A^p(X|r_i)$ (where p is the number of previous observations, with aggregate weight $\sum_{j<=p} w_i^p$) for each r_i (where the corresponding w_i is non-zero) using the weighted aggregation formula:

$$\Theta_A^{p+1}(X|r_i) = WeiAggr((\Theta_A^p(X|r_i), \quad \sum_{j<=p} \quad w_i^p), (\tau_A(X|c_o), w_i^{p+1})) \tag{11}$$

Exact form of the $WeiAggr()$ operator depends entirely on the model used to represent $\Theta_A^{p+1}(X|r_i)$. Assuming for the moment that the $\Theta_A(X|r_i)$ is just a w_i weighted average of all p previous observations, we obtain the update relationship:

$$\Theta_A^{p+1}(X|r_i) = \frac{(\sum_{j<=p} w_i^j)\Theta_A^p(X|r_i) + w_i^{p+1}\tau_A(X|c_o)}{(\sum_{j<=p} w_i^j) + w_i^{p+1}} \tag{12}$$

In the decision time, when we take a trusting decision, current context is determined and the trustfulness is deduced as a weighted combination of trustworthiness of reference contexts.

$$\Theta_A(X|c_d) = WeiAggr_{r_i \in \mathcal{R}}((\Theta_A(X|r_i, w_i)) \tag{13}$$

By assuming the weighted average case again, we obtain:

$$\Theta_A(X|c_d) = \frac{\sum_{r_i \in \mathcal{R}} w_i \Theta_A(X|r_i)}{\sum_{r_i \in \mathcal{R}} w_i} \tag{14}$$

The general update approach as presented in Eq. 11 ensures that several trustfulness relative to different contexts r_i $\Theta_A(X|r_i)$ are updated simultaneously – this is a critical feature for any trust model, as it reduces the time before the model can provide meaningful results. Similarly, the decision-making Formula 13 gathers the data from all relevant contexts, increasing the quantity of the information the trusting decision is based on. On the other hand, this model characteristic can turn in our disadvantage as it can be exploited by informed adversary – therefore, in our future work, we intend extend the Formula 11 to reflect the amount of the information we already have and to reduce the weight of distant reference contexts accordingly.

In this contribution, two different classes of approaches to reference contexts definition will be examined. In the first approach, the reference points will be placed regularly through the metric space (see Section 3), while in the second approach we will enhance the behavior by introduction of adaptive techniques, as shown in Section 4.

Note that while the context space $\mathbb{C}$ properties (e.g. metrics and dimensions) are the same for all agents modelled by agent A, the actual instances of the trustfulness and reference points r_i positions are separate for each evaluated agent – while they may coincide in the reference grid approach, this is no longer true for adaptive approaches mentioned below.

3 Regular Grid Approach

In the first approach to the problem, the reference context set $\mathcal{R} \in \mathbb{C}$ is defined as a regular grid covering the space $\mathbb{C}$ in each dimension q, where the regularity is defined by the distance metric d^q for each dimension. The density of $\mathbb{C}$ sampling is therefore defined a-priori, in design time, before we know what will be the real distribution of samples c_i.

3.1 Experimental Evaluation

To evaluate the regular grid approach, we have conducted a series of experiments using the context space defined in Section 2.1.1. In our simple scenario, several humanitarian organization agents who acquire transportation services from 5 local for-profit transporters (providers). Trustfulness Θ of each provider is strongly dependent on context and influences the *real price* pr_r of delivery that includes not only the amount billed, but also the **price of the lost or damaged goods**. Θ is context dependent and is defined as a function of three parameters that define the context of the operation in the space $\mathbb{C}$: *size* of the contract, *type* of the goods to transport and *road quality*: the status of the roads to use.

In our provider model, the transporters answer the call for proposals with *bid prices* pr_b based on the transportation cost and their margins. Then, at transport time, the real price (e.g. the price of the cargo lost during transport) is derived from the bid price and transporter *real trustworthiness* Θ defined by the position of each particular transport in $\mathbb{C}$. Real price pr_r is determined as $pr_r = \frac{pr_b}{\Theta}$, where

$$\Theta = \Theta_{type} \cdot atan'(price) \cdot atan'(supply) \tag{15}$$

The function $atan'(x)$, used as a sigmoid approximation, is defined as a normalized $arctan$: its range is (x_{inf}, x_{sup}) (both x_{inf}, x_{sup} are in the range set) and x coordinate of its flection point is defined by parameter x_{center}. x_{slope} determines the first derivation - speed of the growth on the domain.

$$atan'(x) = \frac{1 - x_{inf}}{\pi} \cdot arctan(\frac{x_{center} - x}{x_{slope}}) \tag{16}$$

On the humanitarian agent side side, each agent A maintains its trust model. While the context space $\mathbb{C}$ properties

Method	*Aggregate Loss*	*Last 10 Rounds Loss*
General Trust	2,24	2.22
Context Trust	1,83	1.53
Best Case	1	1

Table 1. Loss ratios for General and Context-Based trust relative to the optimum selection (Best Case). Lower values are better.

and grid-defined $\mathcal{R}$ are identical for all transporters, each transporter X is modelled by its own trustworthiness values $\Theta_A(X|r_i)$ in each point r_i. These values are used to discount the bids of suppliers when they answer a particular CFP (CFP content defines the point c_d in $\mathbb{C}$.) Discounting is implemented as a bid price modification – agent actually uses X's trustworthiness to estimate the real price using the relation:

$$\hat{pr}_r(X|c_d) = \frac{pr_b}{(\Theta_A(X|c_d))^{\kappa}} \tag{17}$$

where $\kappa \geq 1$ is a discounting coefficient that penalizes inappropriate behavior. For our experiment, we impose $\kappa = 2$. Once the discounted prices are known by A, it simply selects the agent X with lowest $\hat{pr}_r(X|c_d)$ as a supplier for a particular case. Upon delivery, agents use Formula 11 to update its trust model regarding the X's performance.

The fundamental variable in the experimental results is the difference between the bid price and the real price due to the lost cargo. In the experiment we have performed, the providers were quite specialized, each excelling in one speciality and lacking the skill in other situations. In such context, general trust models are not appropriate as we can see in Table 3.1 and Fig 1, where we present the respective loss of both methods and compare them with the best possible result.

The difference of the outcome of both models is perfectly visible when we visualize the ratio of provider selection for one type of transport over time for both approaches, as we see in Figures 2 and 3, where the application of situational reasoning ensures that the best provider is selected in most cases and dominates the market with increasing experience of transporters over time.

On the other hand, in the general trust case (Fig 2), the mediocre quality providers share the market by capitalizing on their good results from other, unrelated situations.

3.2 Limitations of the Approach

Even if the grid approach suggested above satisfies the requirements for the representation of multi-dimensional context in trusting situation, it suffers of scalability concerns.

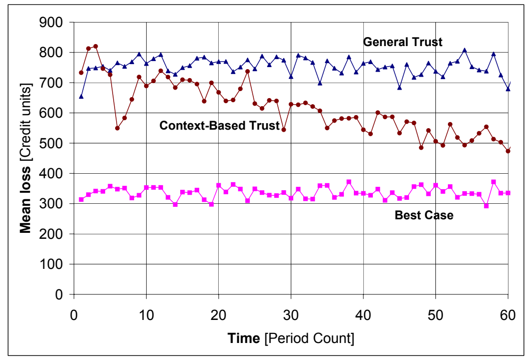

Figure 1. Comparison of the mean loss per round using General and Context-Based trust.

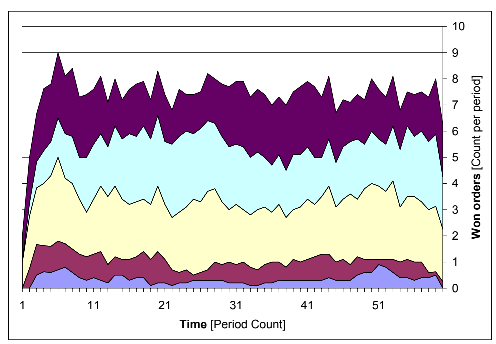

Figure 2. Number of orders captured per provider in a similar particular situation over time – using General Trust approach.

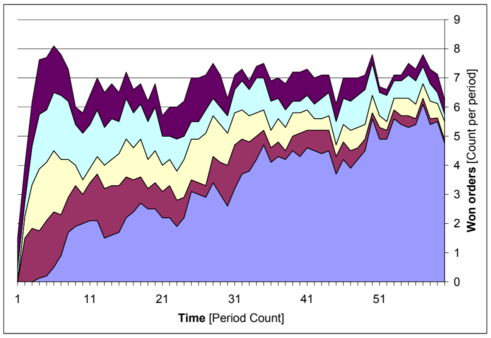

Figure 3. Number of orders captured per provider in a similar particular situation over time – using Context-Based Trust approach.

The main limitation is its granularity: the grid is defined by system designer and is difficult to change once the code is deployed, especially in case of situations when the decision must be taken rapidly or when we have to address hard system constraints. As in the majority of deployment cases the points will be not be spread uniformly, but rather concentrated in several regions of $\mathbb{C}$, most of the reference contexts will be almost useless. On the other hand, there will be relatively few reference contexts in the regions with high concentration of trust situations. Efficiency of the trust update (Eq. 11) decreases significantly with growing density of the reference context grid that each agent maintains for each partner. Therefore, designers face tough decisions: performance optimization can seriously affect the quality of the trust model as it imposes the reduction of grid density. In the remainder of the paper, we are going to examine possible solutions of this problem.

4 Towards Adaptive Approach

The first alternative approach is the introduction of grid with adaptive density using the octant tree-like approach [3]: it shall automatically add new reference contexts in the areas with high density of *diverse* observations. Such regions are easy to detect if we apply trust models that explicitly represent the uncertainty of the information (As most of the current models does – see [6]). Once we detect that some area features a high number of diverse observations, we split the grid cells (multidimensional) into $2^{dim(\mathbb{C})}$ subcells and introduce new reference contexts into the set $\mathcal{R}$. To compute the initial trustfulness values $\Theta_A(X|r_i$ for the new reference point r_i, we apply the Formula 14 – the same one used for evaluation of trustfulness in the decision-making phase. Note that while the initial set $\mathcal{R}$ used by trusting agent A for all partners is equivalent, use of any adaptive method will modify the sets used to model the trustworthiness of other agents individually, as the variability of behavior in different contexts can vary from agent to agent, as well as typical cooperation contexts. Along similar lines, we may reduce the number of reference contexts in the areas with sparse interaction or similar behavior.

A complete extension of this principle is a situation when we initially model each agent with a single reference context r_0; this is equivalent to general trust[3]. Once we detect that the model quality doesn't satisfy our needs, we split the space $\mathbb{C}$ into $2^{dim(\mathbb{C})}$ subspaces and continue with the above algorithm until the results are satisfactory.

Another alternative approach we propose addresses the limitations of the regular grid approach by leveraging classic classification techniques. We propose use a fuzzy c-means (k-means) clustering algorithm [2] to define the set $\mathcal{R}$ – the center of each fuzzy cluster will define a reference point and the corresponding trustfulness will be derived from the observations forming the cluster. The advantage of using the fuzzy variant of this unsupervised-learning algorithm is the fact that each observation contributes its information to several reference contexts: learning is faster, even if arguably more precise. We shall also note that in our particular case, we don't optimize to perfectly separate the clusters in the space $\mathbb{C}$, but rather to define a representative set $\mathcal{R}$ to hold the trustworthiness information.

Fuzzy c-means minimizes following objective function:

$$J_m = \sum_{i=1}^{p} \sum_{j=1}^{|\mathcal{R}|} \mu_{ij}^m d(c_i, r_j) \tag{18}$$

, where μ_{ij} is a membership of the sample context c_i in the cluster j, defined around the reference context r_j. p is the number of available observations.

Position of each reference context is defined by the center of the corresponding cluster – it may therefore move as new observations are classified:

$$r_j = \frac{\sum_{i=1}^{p} c_i \mu_{ij}^m}{\sum_{i=1}^{p} \mu_{ij}^m} \tag{19}$$

the membership coefficient is determined as follows:

$$\mu_{ij} = \left(\left(\frac{d(c_i, r_j)^{\frac{2}{m-1}}}{d_{\mathcal{R}}}\right)\right)^{-1} \tag{20}$$

where the $d_{\mathcal{R}}$ is defined by the relation

$$d_{\mathcal{R}} = \sum_{k=1}^{|\mathcal{R}|} (d(c_i, r_k))^{\frac{-2}{m-1}} \tag{21}$$

[3]With the assumption that the function f used to determine weights is initially constant: $w_i = 1$.

The main problem with the use of fuzzy k-means is the initial phase, when we have o determine how many clusters to create and with what initial center positions. It shall be noted that our situation is somewhat different from the classic clustering – we don't have all the points available from the beginning, but we obtain them one after another, depending on the observation rate and time.

5 Conclusion

In this contribution, we have addressed a very specific problem relevant to trusting decisions in complex environments. While the simple, specialized agents can successfully rely on general trust models [8], once we want to consider the tradeoffs of particular decision or use the trustfulness estimates to draft and evaluate several alternative coalition plans for the same goal, these models can typically no longer provide relevant outputs to support such decisions. We have detailed the general model of *context space* and *reference contexts* where we represent important features of the situation. We have also specified a generic trust update method and trustfulness aggregation method using the general set of reference contexts. This method is independent on the exact form of the set $\mathcal{R}$ and the trustfulness representation formalism used in individual reference contexts.

In the remainder of the article, we have introduced three distinct forms of the set R: regular grid, adaptive octant tree and unsupervised fuzzy k-means clustering. We have performed experiments using the regular rid and determined that it compares favorably when compared with the general trust model that uses the same formalism.

In our future research, we plan to correctly validate and benchmark all above mentioned approaches of reference set representation and determine their mapping to various types of trusting problems and computing environments - we intend not only to evaluate them with respect to trust model quality, but also to describe their computational complexity and other relevant properties. Integration of this method with advanced decision-making algorithms is also a promising area of research.

Acknowledgment

Effort sponsored by the Air Force Office of Scientific Research, Air Force Material Command, USAF, under grant number FA8655-04-1-3044. The U.S. Government is authorized to reproduce and distribute reprints for Government purpose notwithstanding any copyright notation thereon.

The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the U.S. Government.

References

[1] C. Castelfranchi and R. Falcone. Principles of trust for mas: Cognitive anatomy, social importance, and quantification. In *Proceedings of the 3rd International Conference on Multi Agent Systems*, page 72. IEEE Computer Society, 1998.

[2] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. John Wiley & Sons, New York, NY, USA, 2nd edition, 2001.

[3] S. Frisken and R. Perry. Simple and efficient traversal methods for quadtrees and octrees, 2002.

[4] D. Huynh, N. R. Jennings, and N. R. Shadbolt. Developing an integrated trust and reputation model for open multi-agent systems. In *Proc. 7th Int Workshop on Trust in Agent Societies*, pages 65–74, 2004.

[5] S. Marsh. Formalising trust as a computational concept, 1994.

[6] S. Ramchurn, D. Huynh, and N. R. Jennings. Trust in multiagent systems. *The Knowledge Engineering Review*, 19(1), 2004.

[7] S. Ramchurn, N. Jennings, C. Sierra, and L. Godo. Devising a trust model for multi-agent interactions using confidence and reputation. *Applied Artificial Intelligence*, 18(9-10):833 – 852, 2004.

[8] M. Rehák, Lukáš Foltýn, M. Pěchouček, and P. Benda. Trust model for open ubiquitous agent systems. In *Intelligent Agent Technology, 2005 IEEE/WIC/ACM International Conference*, number PR2416 in IEEE, 2005.

[9] J. Sabater and C. Sierra. Regret: reputation in gregarious societies. In *AGENTS '01: Proceedings of the fifth international conference on Autonomous agents*, pages 194–195. ACM Press, 2001.