

Trust Model for Open Ubiquitous Agent Systems

Martin Reháč, Lukáš Foltýn, Michal Pěchouček, Petr Benda

*Gerstner Laboratory, Department of Cybernetics and Center for Applied Cybernetics
Czech Technical University in Prague*

Technická 2, Prague 6, 166 27 Czech Republic

{rehakm1, lfoltyn, pechouc, bendap1}@labe.felk.cvut.cz

Abstract

Trust management model that we present is adapted for ubiquitous devices cooperation, rather than for classic client-supplier relationship. We use fuzzy numbers to represent trust, to capture both the trust value and its uncertainty. The model contains the trust representation part, decision-making part and a learning part. In our representation, we define the trusted agents as a type-2 fuzzy set. In a decision-making part, we use the methods from the fuzzy rule computation and fuzzy control domain to take trusting decision. For trust learning, we use a strictly iterative approach, well adapted to constrained environments. We verify our model in a multi-agent simulation where the agents in the community learn to identify defecting members and progressively refuse to cooperate with them. Our simulation contains significant background noise to validate model robustness.

1. Introduction

The problem of trust in multi-agent systems is already a relatively well defined one, with many important contributions in the area. In our contribution, we extend the existing trust models [3, 13, 8, 14] with the use of fuzzy numbers for trust representation. Our model represents general trust - trust in an individual [10], suitable for embedded devices with specialized range of functions and goals.

We extend the above cited approaches by using fuzzy numbers and arithmetics – a mathematical apparatus known from the field of fuzzy set theory [11] and fuzzy control – to define, implement, and validate the trust model that is iterative, computationally efficient and in the same time robust with respect to considerable environment noise. Other crucial features are (i) automatic identification of the acceptable trustfulness using self-trust, (ii) domain independence and (iii) coalition cooperation representation, instead of simple client-supplier relationship.

After the brief presentation of fuzzy numbers, we will define our formal model in Section 2. Section 3 is then dedicated to model evaluation, and Section 4 summarizes the results and presents our current and future work.

Fuzzy numbers are extension of normal, crisp numbers in the same fashion as the fuzzy sets are the extension of crisp sets. We may define (see [5]) the fuzzy number as a *normal convex fuzzy set on the real line*, where *normality* means that the height of the set is 1; i.e. set has a non-empty core. The set is said to be *convex* iff $\forall(x, y, z) \in (R^3) | x \leq y \leq z$ holds that $\mu(y) \geq \min(\mu(x), \mu(z))$. Example of fuzzy numbers currently used in our model can be found on fig. 1.

Informally, the *core* of the fuzzy set is defined as a subset of the fuzzy set containing the elements x whose membership $\mu(x) = 1$. In some definitions, it is required that the fuzzy number core shall be a single value and when this condition is not fulfilled, the term *fuzzy interval* is used. In this text, we will only use the term *fuzzy interval* to emphasize the cases when the fuzzy number represents a range, rather than value, but all results relevant for fuzzy numbers are valuable for fuzzy intervals as well.

In theory, all the agent’s qualitative reasoning about the cooperation could have been done using the fuzzy numbers and appropriate fuzzy arithmetics operations. In practice, the complexity of this approach could not be justified in most cases. Therefore, **defuzzification** of trust values is necessary for selected operations. An obvious choice for defuzzified value is the core of the fuzzy number. In the case of the fuzzy interval, the center of the core is chosen. Ordering and ranking in general is a crucial operation necessary for partner selection. More advanced approaches have been studied for example by Fortemps and Roubens in [6] or in [2]. The authors provide a method that is very easy to implement with the limitations we have adopted and gives us the results that are intuitive enough. The most notable difference from the real numbers is that the antisymmetry does not hold for this relation on fuzzy numbers.

As we represent the trust, we restrict the support of the set to the $[0, 1]$ interval. Moreover, we limit ourselves to

the fuzzy numbers defined by piecewise linear membership functions on the above specified interval, to speed-up the computation and inference process.

2. Formal Model

In our formal model, we extend the existing trust representations using the fuzzy set theory. To do so, for each agent A we define a set of agents trusted by agent A , denoted Θ_A . We denote $\Theta_A(X)$ the membership function of this set defined on the set of all agents known to the agent A .

Whether Θ_A is a fuzzy set or not depends on the value range and type used for trust definition. Binary trust mentioned above results in a normal, crisp set - membership function takes only two values, $\Theta_A : Agents \rightarrow \{0, 1\}$ - agent is either trusted completely or not at all. Use of the real value in the $[0, 1]$ (or $[-1, 1]$ with transformation) interval defines a standard fuzzy set, $\Theta_A : Agents \rightarrow [0, 1]$.

Use of the fuzzy numbers to represent trust makes the set Θ_A a type-2 fuzzy set, as the membership function itself is a fuzzy set, albeit a simple one (see for example [11]). This does not pose any serious problem to us, as the mathematical concepts necessary to work with fuzzy values are already well developed in the fuzzy control field.

The set Θ_A represents the agent's A trust in other agents as a mental state. Besides this representation, we need to address the following problems: (i) deriving the trust values from the experience, (ii) updating the trust in agents using these values and (iii) using the trust values to make cooperation decisions. We will address these issues in the following sections.

2.1. Deriving Trust Observations from Coalition Cooperation Results

In this section, we will propose a general method how to evaluate the trustfulness of the coalition partners in a specific coalition C as a function of the utility generated by the cooperation. Using this method, each coalition member A can obtain a single value in the $[0, 1]$ interval representing the trust observation τ for each coalition member $Agent$, denoted $\tau_{C,Agent}^A$ or simply $\tau_{C,Agent}$ when no confusion is possible.

We have decided to use a completely peer-to-peer approach that can be applied in most environments where the agents cooperate to achieve their goals. As we try to keep our algorithm as domain independent as possible, we start by normalizing the cooperation result into $[0, 1]$ interval. The simplest way would be to use the minimum utility (maximum loss) u_{min} , expected (maximum) utility u_{max} and real, obtained utility u to measure the *success ratio* u_n

using the formula

$$u_n = \frac{u - u_{min}}{u_{max} - u_{min}} \quad (1)$$

In theory, we could stop here. In practice, this linear relation is rarely appropriate, as shown by many experiments [12]. Therefore, we will typically replace this relationship by subjective loss function with u_{min} and u_{max} as parameters and u as an input to model the perceptions of gain or loss by the agent. The result of this transformation is called *subjective utility* u_s^A (or simply u_s).

As a simple example of such function, our agents do perceive the losses worse than linearly. Therefore, they obtain their final *subjective utility* by raising the value u_n to power of two, obtaining the value $u_s^A = u_n^2$ used as an input for the suite of the process.

Raising to the power of two is an arbitrary choice, modelling the fact that the losses are perceived worse than their real value, but may have another signification - attribution of the blame to an individual, rather than to some stochastic process [7].

Each coalition member calculates its value u_s^A and uses this value to obtain the values $\tau_{C,Agent}^A$ for all other coalition members. Different strategies may be used to do so, analogously to profit distribution in coalitions[9]. The cases we consider in the scope of the current work are:

- **Equally** - the value u_s^A is used as an input to update trust in each coalition member.
- **Proportionally** - the observation value depends on the defuzzified apriori trust the agent has in the coalition member and the u_s^A . Currently, we use the relation:

$$\tau_{C,Agent}^A = \frac{defuzzy(\Theta_A(Agent)) \times u_s}{Avg_{Agent_i \in C}(defuzzy(\Theta_A(Agent_i)))} \quad (2)$$

where $\tau_{C,Agent}$ denotes the trust observation (real number) we derive from the agent's participation in the coalition C . $defuzzy(\Theta_A(Agent))$ denotes the center of the apriori trust membership function, formally core of agent's membership function in Θ_A .

The value u_n^2 (or $\tau_{C,Agent}$ respectively) is then used as an input for updating the trust in $Agent$ as described in the following section.

2.2. Iterative Learning of Trust Values

When we try to represent the trust in agent B $\Theta_A(B)$ as a single fuzzy number, we must be able to find an optimum form of this number to represent the past experience. Formal restrictions on the fuzzy number are not very strict, but

we limit ourselves to *piecewise linear fuzzy numbers* and *iterative learning* in order to keep the processing lightweight and well adopted for potential embedded solutions.

To simplify the notation, we will denote τ_B^A or τ_B all trust observations of agent A about the agent B - suite of n_B real values in $[0, 1]$. Note that these values are not kept in agent's memory, as the learning is iterative.

Strictly speaking, iterative learning requires that the new value for the $\Theta_A(B)$ is obtained only using the apriori value of the $\Theta_A(B)$ and the observation $\tau_{C,B}$. In order to make our algorithm easy to understand, we shall maintain some supplementary simple data - like n_B , number of previous observations used to establish the trust in agent B and the $Avg\{\tau_B^2\}$ used to estimate the variance of the data as shown further.

Trust *average* $Avg\{\tau_B\}$ is a good candidate to define the center of our fuzzy number - it is easy to understand, it is consistent with non-fuzzy approaches and it may be computed iteratively, provided that we keep both the current average and n_B - number of observations used to obtain this value.

But besides the center, we need to represent the uncertainty at least by specifying the left and right sides of the fuzzy number membership function. The simplest way to do so is to use triangular fuzzy number and to define left and right bounds as $\min\{\tau_B\}$ and $\max\{\tau_B\}$. This corresponds indeed with the intuitive interpretation using the possibility theory (min and max values provide us with an interval where the whole past experience falls). Moreover, the values are trivial to maintain iteratively. On the other hand, two major drawbacks outweigh the advantages of this method - (i) both values tend to have only limited information content in the real environments, as they are often very close to 0 or 1 for most agents B and (ii) the uncertainty of the trust is actually non-decreasing with growing number of observations n_B , what is being counterintuitive.

Another value, at least partially solving the issues with minimum and maximum approach, is *variance* $\sigma_A^2\{\tau_B\}$ or *standard deviation* $\sigma_A\{\tau_B\}$. It also has an advantage of being well understood and even if it cannot be maintained iteratively as easily as average, it is possible to estimate it rather exactly using the well known relation:

$$\sigma^2\{\tau_B\} \leq \widehat{\sigma^2\{\tau_B\}} = Avg\{\tau_B^2\} - Avg^2\{\tau_B\} \quad (3)$$

With growing number of observations, the right-hand side of the inequality approximates the variance fairly well.¹

The representation we propose, (see fig 1), uses all the notions presented above. The left boundary is defined as $\min\{\tau_B\}$, right boundary as $\max\{\tau_B\}$, both




Figure 1. Trust membership function example.

with the membership = 0. Core is defined by the average value $defuzzy(\Theta_A(B)) = Avg\{\tau_B\}$. Two points are added on each side of the center, both with membership = $\frac{1}{n_B+1}$ that decreases with increasing number of data, with positions $max\{\min\{\tau_B\}, Avg\{\tau_B\} - \widehat{\sigma_A\{\tau_B\}}\}$ and $min\{\max\{\tau_B\}, Avg\{\tau_B\} + \widehat{\sigma_A\{\tau_B\}}\}$. At first, min and max values are prominent. With increasing number of observations, the relative importance of the extremes decreases and the bulk of the data represented by the central triangle becomes dominant. This behavior corresponds well with the human approach to the same problem.

2.3. Self-Trust as a Parameter for Trusting Decisions

In the paragraph above, the agent has never excluded itself from the group of agents between whom we distribute the collaboration trust value. It means that the agent actually estimates the trust in itself: $\Theta_A(A)$. There are two good reasons for such behavior.

First, an agent does not necessarily trust itself - we may easily imagine a situation when an agent runs on a hardware with malicious intruding software and is almost never able to protect its private data and communications from platform-originating intrusion [1]. When the agent observes its self-trust and detects a significant decrease, it may decide to migrate, to interrupt communication with others or even to terminate itself in order to protect the cooperators.

The other reason why we measure the self-trust is environmental adaptation. In many cases, it is difficult or even impossible to estimate correctly what is the expected payoff of the cooperation in the given environment. In our approach, we don't take this factor into account during the

¹ We assume that the trust observations are conform to Gaussian distribution.




Figure 2. Example of the trust decision. Height of the intersection with low trust (left segment) and high trust (right segment). As the incidence with the high trust is bigger, agent is considered to be trustful.

evaluation of the cooperation success - we rather integrate this information into the cooperation rules derived from the self-trust data.

We define two linguistic variables [4] on the trust membership support $([0, 1])$. First of them is a *low trust* domain, denoted LT^A while the other is *high-trust* domain, HT^A . The sum of their membership functions is equal to 1 on the whole interval $[0, 1]$ - they form a partitioning of unity.

We use the self-trust data to establish the fuzzy intervals HT^A and LT^A as follows. First, we define that $HT^A(1) = 1$, a natural assumption as the complete trust is undoubtedly high. Then, we say that agent A considers itself as trusted. (We use the self-trust for environmental adaptation, rather than for intrusion detection.) Therefore, we say that $HT^A(\text{defuzzy}(\Theta_A(A))) = 1$. From this value on, we decrease the trust linearly until we reach the 0 membership for the trust = $\max\{\min\{\tau_A\}, \text{defuzzy}(\Theta_A(A)) - \widehat{\sigma_A\{\tau_A\}}\}$. LT^A is complementary - it is equal to 1 between 0 and $\max\{\min\{\tau_A\}, \text{defuzzy}(\Theta_A(A)) - \widehat{\sigma_A\{\tau_A\}}\}$, where it starts to decrease linearly to finally reach zero at $\text{defuzzy}(\Theta_A(A))$. The use of linguistic variables for the inference process is shown in fig. 2.

2.4. The Decision to Cooperate and Partner Selection

A set Θ_A and the fuzzy intervals HT^A and LT^A represent the mental state of the agent.

When an agent proposes a coalition or is invited to participate in one, it needs to take a trusting decision; it has to decide which other agents are admissible as partners and

order the admissible partners by trust to minimize the risk. While using the fuzzy rules to take trusting decisions, we use the fuzzy numbers as inputs for these rules. Therefore, we must be able to determine the **inference** with fuzzy intervals representing the rules/decisions. We have opted for the use of Mamdani inference, defined as

$$D_T(X^*, A_i) = \text{hgt}(X^* \cap_T A_i) \quad (4)$$

where T is a selected *t-norm* (see [11]). We have selected the Standard (Gödel, Zadeh) t-norm, defined as

$$T(A, B) = \min(A, B) \quad (5)$$

To establish whether an agent B is trusted, we use the formula 4 to calculate the incidence of the trust in the agent B with the intervals HT^A and LT^A .

$$D_{\min}(\Theta_A(B), HT^A) = \text{hgt}(\Theta_A(B) \cap_{\min} HT^A) \quad (6)$$

$$D_{\min}(\Theta_A(B), LT^A) = \text{hgt}(\Theta_A(B) \cap_{\min} LT^A) \quad (7)$$

Agent B is considered to be *trusted* iff $D_{\min}(\Theta_A(B), HT^A) \geq D_{\min}(\Theta_A(B), LT^A)$.

When an agent A needs to organize a coalition, it identifies a subset of trusted agents. Then, it calculates the *usefulness* (see [3]) of these agents for the coalition using the social knowledge in its acquaintance model. The usefulness of each agent is then multiplied by the trustworthiness (defuzzified) of this agent, to account for the *willingness*² and the candidates are ordered by this value. Suitable subset of acceptable candidates is then invited to form a coalition. In our current work, we exploit the fuzzy ranking techniques to include the risk notion into the coalition planning.

In the opposite case, when the agent A is invited to participate in a coalition, it evaluates its trust in the members of the coalition. When all members are considered to be trustful, it agrees to take part in the coalition.

We discuss the future development of this mechanism in section 4.

3. Experiments

In order to compare various trust representation methods, we have conducted the experiments using a fully-fledged multi-agent simulation using a logistics management scenario in the non-collaborative, self interested environment - ACROSS developed using \mathcal{A} -globe [15] agent platform.

² In this work, we consider usefulness and willingness as perceptions. Therefore, we model the environment effects in the willingness rather than in the usefulness part, that is established using the information provided by agents themselves.

3.1. Scenario Description

In this section, we will briefly present the experimental scenario we have used to conduct the experiments. Following types of agents form the scenario:

Location Agents: Location agents represent population (cities) and natural resources. They create, transform or consume resources and they also acquire the transport of goods from the ad-hoc coalitions of transporter agents.

Transporter Agents: Transporter Agents are the principal agents in our scenario. They use their resources - vehicles, driven by *Driver agents* - to transport the cargo as requested by location agents. As a normal request exceeds the size that may be handled by a single transporter, transporters must form one-time coalitions in order to increase the coverage – the principal evaluation parameter in the auction bids. To form these coalitions, they use the trust model presented in the previous section.

Driver Agents: Driver Agents drive the vehicles owned by Transporter agents. They handle path-planning, loading, unloading, status reporting and other driver duties.

There are two causes for the cooperation failure - random attacks and member defection. **Random attack** on drivers, when they become victims of the bandit agents, are the "standard" reason for the failure. Selection of victims is a pseudo-random process with probability p_n (see details below).

Directed attacks on drivers are a result of coalition **member defection**. In this case, one or more members of the coalition inform the bandit simulation component (a service running on the system infrastructure simulating the actual attacks) about the coalition goal and members, increasing significantly the likelihood that the drivers working on this task will be attacked (to p_d value). In exchange, the defecting agent's drivers obtain the immunity during this operation and defector receives side payment from the bandit.

When the cargo is lost during the transport, the payments to all coalition members are reduced, as the coalition is paid only for the completed deliveries. The losses are shared between the members. Transporters are therefore motivated not to cooperate with defectors or frauds and they use the above described reasoning to select their coalition partners.

The problem we solve is analogous to iterated prisoners dilemma (with more than two players) with significant background noise, produced by the following phenomena: (i) The attack probability depends on the length of the transport path and the speed of the car, as the random experiment that may result in attack is done in each step the vehicle travels with the cargo - this fact was intentionally omitted in the

	p_n	p_d	r_{attack}	$r_{success}$
Scenario A	0	0.04	–	–
Scenario B	0.0005	0.04	1 : 1.8	1 : 0.85
Scenario C	0.002	0.04	1 : 7	1 : 2.5
Scenario D	0.004	0.02	1 : 12	1 : 4

Table 1. Scenario settings.

trust reasoning. (ii) The success of the cooperation is determined by the ratio of the delivered cargo - and each vehicle has different capacity and load. Therefore, a single attack on a fully loaded car can have greater impact than many attacks on the cars with small lots. (iii) The drivers working for the defector are protected. In the very special case when they carry a bulk of the coalitions' cargo, defection may actually increase the success of the coalition. This is also true in cases of high background crime – when the probability of undirected attack approaches one, the drivers working for the defector are still protected and are often the only ones to deliver the cargo.

We have conducted four series of experiments, with variable parameters regulating the background noise. Simulation parameters - attack probability without defection (p_n) and with defection (p_d) - are shown in table 1. Variable r_{attack} denotes the average ratio of informed attack attempts to uninformed ones and $r_{success}$ the average ratio of informed successful attacks (when some goods was stolen) to uninformed ones. The group of agents for which we present the results consisted of 9 fair agents and 1 defector and we present the results aggregated from 10 runs.

In our measurements we have focused on a trust evolution during time in a noisy environment. Data plotted in the graphs shows for how many agents the particular agent is trustworthy. For sake of simplicity we have chosen only values of the defector and the average value throughout the whole alliance, including the defector. Size of the gap between the two values in the given moment t grows with the number of agents that have detected the defector by the time t .

The scenario A is considered to be a reference case as there are no background attacks. We can observe how the trust of alliance members into defective agent falls – figure 3. We should mention that average value of trusting agents is lower than in next scenarios because each agent has a high self-trust and is therefore cautious trusting someone else. The actual numerical trustfulness values $\Theta_A(B)$ are lower than in the other scenarios, but the self-trust value correctly identifies the environment as a low-risk one.

In the second and third scenario (B,C), a new source of the background noise was introduced by the activation of random attacks. This made the trust learning harder, as the agents need more data to distinguish the coalitions with and




Figure 3. Trust evolution in Scenario A (limited background noise).

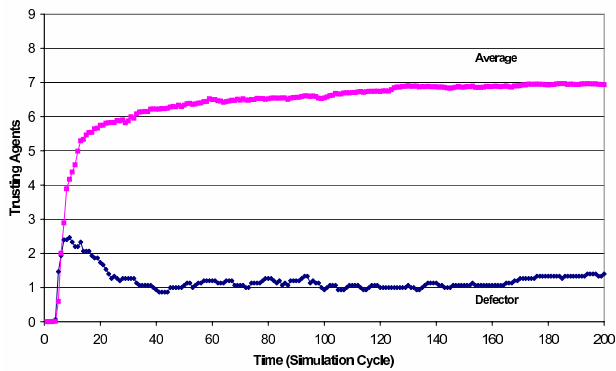


Figure 4. Trust evolution in Scenario B (low background noise).

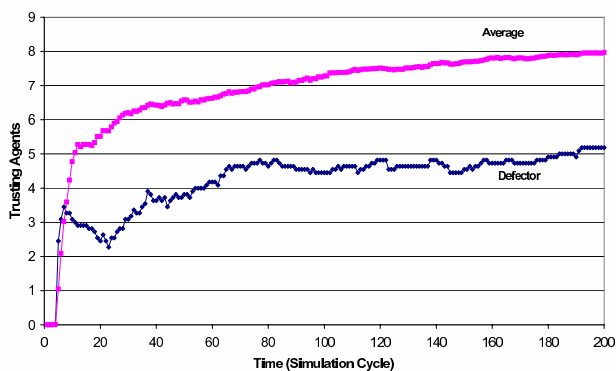


Figure 5. Trust evolution in Scenario C (medium background noise).

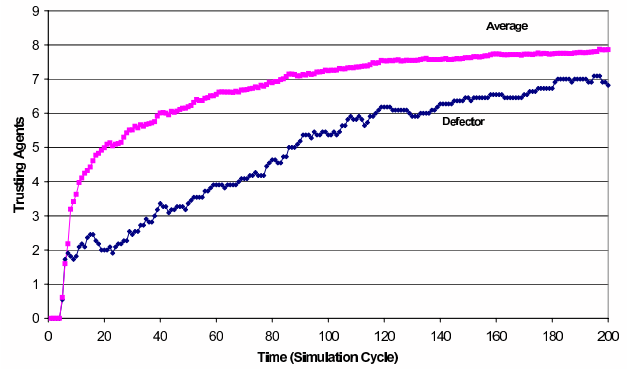


Figure 6. Trust evolution in Scenario D (high background noise).

without defector. Due to this issue, some of the agents trust the defector and are ready to cooperate with it until later in the runs – figure 5. However, agent is successfully detected and refused as a coalition member by the majority of agents, especially in fig. 4, where the noise is comparable with the signal. We should emphasize that agents were able to reveal defector from data (C) where only 29% of the attacks were caused by member defection, while the 71% of attacks are attributed to background noise.

In the scenario D we have further increased the background noise and simultaneously decreased the probability of attack with defection, making the environment even harder. As graph in figure 6 indicates, settings for Scenario D are on verge of reasonable limits. Although we can still detect the defector, the difference between defector’s trustfulness and average trustfulness value in alliance is low - about 10%. In such setups, we discover the limits of our approach - only about 20% of the observations correspond to the real agent defection. In the real applications, we can easily improve algorithm performance by inclusion of context information (for instance the transport path length in our case) and improve the algorithm performance in the well defined context.

4. Conclusions and Future Work

In this paper, we have presented a general trust model that addresses properties that are crucial for most ubiquitous systems: *iterativity*, *computational efficiency*, *self-adaptation* and *team cooperation* representation. Another major feature of our model is *robustness* in a high-noise environment, where we have shown that the model works even with limited knowledge about the mechanisms of the environment, using only the cooperation result as an input. This also allows easy model integration with existing devices and algorithms, where it provides both the quantitative (trustful-

ness) and qualitative (decision) output.

The experiments conducted in the simulation environment with various (both systematic and random) noise influences have proved that the fuzzy-number approach is justified and is robust with respect to more than significant environment noise. We have shown that a general, non-situational trust can be sufficient even for complex ubiquitous environments, when the agents are highly specialized - a typical case in embedded control or sensor networks. We have also shown that our mechanism autonomously adapts to the environment using the *self-trust* concept.

However, we expect to develop the presented model further to cover many issues not considered so far. The most important issue is the time variability of agent's behavior. The model will fail to detect reasonably soon the agents who start to defect after a big number of observations, and the use of the current decision-making process makes the exclusion of agent from the community an irreversible process, leaving no option for re-integration of the agent into the community. Therefore, two other extensions are possible - emphasizing the data acquisition time and extension with reputation management mechanism.

Situational trust is another obvious development path, where fuzzy rules and fuzzy control provide the well-developed apparatus for efficient decision-making. Fuzzy number representation of the trustfulness values also allows more advanced planning in the non-cooperative environment, as the risk management and game-theoretic approaches can benefit from the conveniently represented information.

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] J. Ametller, S. Robles, and J. A. Ortega-Ruiz. An implementation of self-protected mobile agents. In *ECBS*, pages 544–549, 2004.
- [2] G. Bortolan and R. Degani. A review of some methods for ranking fuzzy subsets. *Fuzzy Sets and Systems*, 15:1–19, 1985.
- [3] 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.
- [4] D. Driankov, H. Hellendoorn, and M. Reinfrank. *An Introduction to Fuzzy Control*. Springer-Verlag, New York, 1993.
- [5] D. Dubois and H. Prade. Fuzzy real algebra:some results. *Fuzzy Sets and Systems*, 2(4):327–348, 1979.
- [6] P. Fortemps and M. Roubens. Ranking and defuzzification methods based on area compensation. *Fuzzy Sets Syst.*, 82(3):319–330, 1996.
- [7] R. J. Heuer. *Psychology Of Intelligence Analysis*. Center for the Study of Intelligence, US National Technical Information Service, 1999.
- [8] 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.
- [9] S. Kraus, O. Shehory, and G. Taase. The advantages of compromising in coalition formation with incomplete information. In *AAMAS*, pages 588–595, 2004.
- [10] S. Marsh. Formalising trust as a computational concept, 1994.
- [11] W. Pedrycz and F. Gomide. *An Introduction to Fuzzy Sets : Analysis and Design*. Complex Adaptive Systems. Bradford Book, Cambridge, London, 1998.
- [12] M. Rabin. Risk aversion and expected-utility theory: A calibration theorem, 2000.
- [13] 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.
- [14] 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.
- [15] D. Šišlák, M. Rollo, and M. Pěchouček. A-globe: Agent platform with inaccessibility and mobility support. In M. Klusch, S. Ossowski, V. Kashyap, and R. Unland, editors, *Cooperative Information Agents VIII*, number 3191 in LNAI. Springer-Verlag, Heidelberg, September 2004.