IEEE International Conference on Social Computing / IEEE International Conference on Privacy, Security, Risk and Trust

A Trust Model for Intervehicular Communication Based on Belief Theory

Walter Bamberger, Josef Schlittenlacher, and Klaus Diepold
Department of Electrical Engineering and Information Technology
Technische Universität München, Munich, Germany
w.bamberger@ibbamberger.de, josef.schlittenlacher@mytum.de, kldi@tum.de

***Abstract*—Vehicles will exchange much information in the future in order to efficiently maintain their inner model of the environment. Before they can belief received pieces of information, they must evaluate their reliability. Trust is a mechanism to estimate this reliability based on the sender. As cars often drive the same route, they meet each other again and again. They can establish friendship-like relations and thus are embedded in a social structure. A trust model depends on this social structure. For this reason, we simulate the driving pattern of a small town. Within this simulation, all cars are equipped with a trust model that continuously monitors the experiences made with others. The developed model focuses on direct experiences of the individual and not on a system-wide reputation which would depend on a central unit. It continuously evaluates the performance and reputation of other cars and includes a feedback loop to faster adapt to changes in the other's behaviour. To make a decision out of the collected data, the model uses the capacity of the binary error and erasure channel from information theory. This capacity provides a better decision criterion than the traditional expectation value. The proposed trust model is an individual-level model; nonetheless it can be connected to a system-wide reputation mechanism.**

***Index Terms*—Belief theory, information theory, probabilistic logic, reputation, social structure, traffic, trust, uncertainty, vehicular ad hoc network**

I. Introduction

Consider a scenario in which cars continuously communicate with each other while driving around. These future cars can perceive, reason, learn, plan, and act in a way that they understand the surrounding traffic scene while still controlled by a driver (as in the CoTeSys satellite project MuCAR-3 in Munich). In the literature, these cars are sometimes called Cooperative Cognitive Automobiles [1]. They are expected to improve the traffic efficiency and safety. These cars need a vehicular network to efficiently maintain the inner model of their environment. With this network, they exchange all kinds of model information about structural alterations of the road network, traffic signs, points of interest (like petrol stations or hotels), hazardous locations (e.g. oil on the road), free parking space, traffic congestions, green light timing, collision warning, and so on [2–4].

A. The Problem Statement

Despite the cryptographic measures already proposed for the information security in vehicular networks, the received data may still be wrong [5, especially p. 65]. Reasons for this might be bad recognition capabilities or defect sensors at the sender, but also computer virus infections or manipulations of the sending car. So, a car must try to handle the uncertainty associated with received data. It does so in two steps: It evaluates the sender over several interactions and it evaluates the received data with regard to the current situation. This is where the trust model proposed in this paper comes in. It is a mechanism to evaluate the sender over several interactions. Models for the second step (that is, evaluate the data with regard to the current situation) are, for example, introduced in [6] and [7].

Thus, this paper addresses the following problem: How can a car estimate a belief, whether another car will send correct information? The authors of this paper call a mechanism that can do this *inter-agent trust*. So, the proposed model is a trust model. More specifically, it focuses only on the processing within the car; it is only the individual-level side of a trust-enabled network according to Ramchurn et al. [8].

With the help of this belief, the car should be able to distinguish good from bad information sources. In the end, it should have a more accurate inner model of the environment and should better decide what information should when be exchanged with whom (i.e. what information is really necessary and who would give it). The following scenarios make clearer in what situations trust is necessary and how a trust model must be designed.

B. Trust Scenarios

Cars have differently restricted sensory capabilities. Therefore, they sometimes miss to report an event (like a new sign) or spread out wrong information (e.g. a wrong degree of free parking space). The error rate depends on the car and on the involved sensors.

The trusting car must know how correct the values of the other car usually are. The authors call this property the competence of the other car (see Sect. IV-A). Received values must be processed with an uncertainty that is related to the sender's competence.

A car randomly generates fake messages about nonexistent events. This may be, because it has been infected by a software virus or because the manipulator does this just for fun.

In this scenario, the trusting car must estimate how good it can predict the other's behaviour. If this is not possible (because the other's behaviour is sometimes very good and

978-0-7695-4211-9/10 $26.00 © 2010 IEEE
DOI 10.1109/SocialCom.2010.20

sometimes very bad), it should increase the disbelief in the other's trustworthiness as shown in Sect. IV-B.

Every Wednesday, a car reports wrong parking space values to ensure it can park on the desired car park. While it drives to the car park, it claims it has been there 30 minutes ago and there was no free parking space. It reports bad recommendations about other cars of the opposite direction that have reported free parking space. Besides that Wednesday, the car always acts trustworthy.

As in the previous scenario, a quasi-random misbehaviour should be considered as untrustworthy. In addition, this scenario shows that reputation can be manipulated (to increase and to decrease it). Because of this, the proposed trust model takes recommendations with care as explained in Sect. IV-C.

C. Overview and Terms

So far, we clarified the problem this paper tackles and its surrounding scenarios. Next, we give an overview of the paper and the specifics of the proposed trust model.

First, we relate our model to selected other trust models from different fields. Then, we briefly introduce subjective logic, the framework we use for artificial reasoning.

In the subsequent Sect. IV, the model itself is presented in detail. In contrast to reputation systems, which are already in discussion for intervehicular communication [e.g. 5], the proposed model focuses on the individual car and its experiences. The opinion of others – that is, reputation – is still considered, with lower priority though. This paper further features the capacity of the binary error and erasure channel as a better indicator for trust when uncertainty is there (see Sect. IV-E). While trust is often related to something like a mean error in the performance of the other agent, the model proposed here additionally incorporates how good the car can decide based on the information received from the other car. The authors call this property the other's *predictability* (detailed in Sect. IV-B). It helps to faster adapt to changes in the other's behaviour.

As pointed out so far, the authors understand trust as a mechanism that relies on a social structure. The "social structure" of cars comes from their periodical trips. So, these trips must be part of an evaluation environment for trust. Sect. V introduces such an environment and proposes it as the right environment to investigate a trust model for intervehicular communication. The results section then shows that a car indeed meets some other cars regularly. So, a car can create an image of other cars on its own – during "social interaction".

The following terminology is used throughout this paper: *Reputation* is considered as the opinion about someone by people in general – it represents the common opinion. It can be used as an indication when judging about the trustworthiness of another party, but it is different from trust. The proposed trust model tries to build up a local view on the reputation by combining several received *recommendations* of others about a third car.

When a car drives around, it observes properties of its environment (*observation*). It may pass these properties to other cars. This received information is called a *report*. Information that helps to judge about the trustworthiness of another car is *evidence*. So, a report that has been verified by the car becomes evidence.

II. Related Work

When people from vehicular network security talk about trust, they usually refer to trust relations in a public key infrastructure. This mechanism from information security mainly provides identification of entities but not trust in the sense of this paper. Raya and Hubaux [5] give a good overview about security aspects and mechanisms in vehicular networks. For data verification, they first considered reputation [5, p. 65], but then developed their own data-centric trust model [7]. While their model focuses only on the data (disregarding the sender), the proposed model focuses only on the sender. So, both models complete each other. Reference [6] also verifies the data, this time based on a model of the network.

Vehicular ad hoc networks are similar to wireless sensor networks. In both networks mainly sensory data is exchanged. The framework of Zhang et al. [9] served as a starting point for our work. In contrast to their scenarios, a car receives less reports about the same subject, but it can get to specific places (if the driver does so) to verify the received information itself.

In the area of electronic marketplaces, reputation systems are well known (e.g. eBay's feedback system [10], the beta reputation system [11], or even Google's PageRank [12]). Here, the trust comes from human beings. The reputation system then combines these statements to build a global view on an entity – the reputation. To get this global view, reputation systems work on the network level. In contrast, the trust directly comes from artificial agents in purely electronic market places of multi-agent systems. Ramchurn et al. [8] and Sabater and Sierra [13] give a good overview of the trust and reputation models in this area.

III. Subjective Logic

Subjective Logic [14, 15] is a framework for artificial reasoning under uncertainty. It has a foundation in Bayesian statistics and set theory. This chapter shortly introduces the main concepts used in this paper. The reader finds more details in the referenced papers.

A. Opinions

An agent can have some evidences that support a statement and some that oppose it. More generally, the agent could argue about more than two possible situations (more than just pro and contra). Let $X = \{x_1, \ldots, x_N\}$ be the set of considered situations. In the evidence representation of an opinion, the strength of evidences for each situation is expressed in the variables $r_i \in [0; \infty[$. These variables are collected in the evidence vector $r = (r_1, \ldots, r_N)^T$. In addition, an agent may have a subjective opinion about the situations without any direct evidence. The base rate vector a expresses this opinion. Its influence decreases with an increasing strength of evidences. The tuple $\omega_X^M = (r, a)$ then describes the opinion of the agent M about the set of situations X. The more evidences the agent has in total ($\sum_i r_i$), the more certain it is about its opinion.

The same opinion could also be expressed with probabilities $p = (p_1, \ldots, p_N)^T$ for every possible situation and a degree of uncertainty u about these probabilities. The tuple $\omega_X^M = (p, u, a)$ is a probabilistic representation of the opinion.

Finally, the opinion could also be expressed with the beliefs $b = (b_1, \ldots, b_n)^T$. Beliefs [16] are subjective ratings of the situations with the constraint $\sum_i b_i + u = 1$. So, the belief representation of an opinion consists of the tuple $\omega_X^M = (b, u, a)$.

There exist mappings between these three representations of an opinion, based on the Dirichlet distribution. Only the mapping from the evidence to the belief representation

$$\begin{aligned} b_i &= \frac{r_i}{W + \sum_{i=1}^N r_i} \\ u &= \frac{W}{W + \sum_{i=1}^N r_i} \end{aligned} \tag{1}$$

is explicitly used in this paper. W denotes the weight of the base rate vector a.

B. Operations on Opinions

When opinions from two or more different sources exist about a set of possible situations X, they can be fused into one [15]. If the evidences are independent (like two rolls of the same dice), then the *cumulative fusion* $\omega_X^A \oplus \omega_X^B$ is the right fusion. For example, this can be used to combine the opinions about the outcomes of several interactions for trust development. If the evidences are dependent (like two concurrent, maybe conflicting observations of one dice roll), then the *average fusion* $\omega_X^A \underline{\oplus} \omega_X^B$ must be applied. This operation is appropriate to combine the opinions about several properties of the outcome of a single interaction.

When a car B has the opinion ω_X^B about the proposition that there is a new traffic sign at a certain location, then A only beliefs this as far as it considers B as trustworthy in providing such information. So, if A has the opinion ω_B^A about B's trustworthiness, then both opinions can be combined with the *uncertainty favouring discounting* [17] into the opinion $\omega_X^{A:B} = \omega_B^A \otimes \omega_X^B$. The discounted opinion states how much A beliefs the proposition based on B's opinion.

When an opinion should have a lower weight in the cumulative fusion operation (i.e. is less important), the uncertainty must be increased. For example, this could be necessary because its evidence is older than that of the other opinions. The certainty can simple be scaled with the factor w by scaling the evidence vector r with this factor: $w\,\omega_X^M = (w\,r, u, a)^T$. The authors call this operation the *certainty scaling*.

IV. The Trust Model

As described in the introduction, cars exchange their knowledge about the environment in form of reports. Such a report contains the opinion ω_v^A of the sender A about the possible values v_i of a model attribute v. This is, for example, the belief ($v_+ =$ yes) and disbelief ($v_- =$ no) whether there is a new traffic sign at a certain location. The receiver M uses

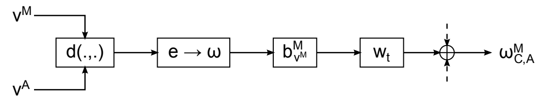

Fig. 1. Signal flow for the competence evaluation. The difference between what the car M thinks is right and what the car A has sent to M determines the competence of A from M's point of view.

this opinion (and all the opinions received from other cars) to build its own opinion ω_v^M about v. But M takes A's opinion also to develop trust in A. For this, it must judge how good the opinion is. However, M can do this only, if it knows the correct value of v. So, it must wait until it is very certain about this value. Then, it can evaluate all reports regarding v. The reports then become evidences for the trust development. This section describes how a car evaluates those reports to update its opinion ω_A^M about the trustworthiness of A. The figures 1 to 4 guide through this process.

The proposed trust model incorporates three components: M's opinion $\omega_{C,A}^M$ about A's competence reflects the mean error of all available evidences. This is the main criteria whether reports of A are good. The opinion $\omega_{P,A}^M$ about A's predictability indicates, whether M is able to make a right decision based on A's opinion and its trust in A. It is sensitive to outliers. As its name indicates, its computation includes a predictor. Based on the predicted value, feedback control adjusts the predictability opinion. Finally, the cars exchange their opinions about other cars in addition to the reports about their environment. These opinions are called recommendations and build up M's opinion $\omega_{R,A}^M$ about A's reputation. In fact, this is not the opinion of all cars that know A but only M's view on the reputation of A. So we call it a local reputation of A.

A. Evaluation of the Competence

Trust and reputation are often derived from the number of good and bad experiences in the past or from the degree how good or bad these experiences were [e.g. 9–11]. The authors of this paper call this notion of trust the competence of the other car.

A car is more competent, if the information it provides is more precise. Thus, competence is inverse to the mean error of all evidences. Fig. 1 shows the steps, how a car M can obtain an opinion about the *competence* of another car A.

For each evidence i, it first computes a distance between the value v_i^A proposed by A and the value v_i^M which M thinks is right. This operation can use any metric that maps to the closed range $[0; 1]$. In many cases, the relative error $e_{\text{rel},i} = \frac{e_i}{e_{max,i}}$ might be appropriate.

This relative error supports the opinion about A's incompetence. With it, the opinion about A's competence in a single situation i can be calculated in the evidence space as

$$\begin{aligned} r_- &= e_{\text{rel},i} \\ r_+ &= 1 - e_{\text{rel},i} \ . \end{aligned}$$

r_- reflects A's incompetence and r_+ A's competence.

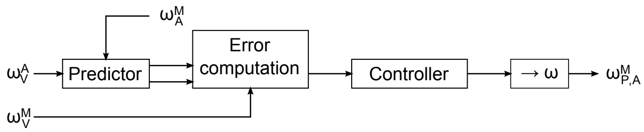

Fig. 2. Signal flow for the predictability evaluation. First, only based on the opinion of A, a decision is made as described in section V-C. Then, this decision is compared with the car's own opinion about v. If both are different, an error measure is passed to the controller, which, in turn, adjusts $\omega^M_{P,A}$.

M cannot be completely sure, whether its reference value v^M_i is correct. So, it can also not be sure about the relative error. The belief $b^M_{v^M_i}$ in its assumed correct value v^M_i reflects the belief in the correctness of the relative error. As a consequence, M must scale the above computed opinion with $b^M_{v^M_i}$.

Finally, the opinions about A's competence in all situations are combined to an overall opinion about A's competence. For this operation, the cumulative fusion operation of subjective logic is appropriate. To have an opinion about A's present competence, old evidences must be weighted lower in this operation. Therefore, all opinions are scaled with a time-dependent factor $w_{t,i}$. (It is further discussed in Sect. IV-F.) So, in the end, the opinion about A's competence is

$$\begin{aligned} r_+ &= \sum_{i=1}^{N} w_{t,i}\, b^M_{v^M_i}\, (1 - e_{\mathrm{rel},i}) \\ r_- &= \sum_{i=1}^{N} w_{t,i}\, b^M_{v^M_i}\, e_{\mathrm{rel},i} \ . \end{aligned} \tag{2}$$

So far, we have shown how several evidences with the car A are transformed in one opinion about A's competence. The remaining paragraph shows that this transformation has a reasonable interpretation even in belief representation: The disbelief in A's competence results from (1) and (2) as

$$b_- = \frac{\sum_{i=1}^{N} w_{t,i}\, b^M_{v^M_i}\, e_{\mathrm{rel},i}}{\sum_{i=1}^{N} w_{t,i}\, b^M_{v^M_i} + W} \ .$$

W is the weighting factor for the base rate vector (the a priori probability) in subjective logic. When only the observed evidences are taken into account without a priori considerations ($W = 0$), then b_- is a weighted mean of the relative error. So, $\omega_{C,A} = (b, u)$ reflects the competence as proposed in the beginning of this section: competence is inverse proportional to the mean error. The weighting factors are due to the age of the evidence and M's own certainty about each computed error. Both are reasonable.

B. Evaluation of the Predictability

The competence measure given above describes the mean behaviour of the other car. It does not account for outliers or quick changes in the other's behaviour. As a consequence, this paper proposes to include a feedback control that is sensitive to outliers. The idea is, that the car M behaves as if it received three times the same message from three other cars that have about the same trust value. Would it decide right, then? So, it predicts its decision based on the message of the other car A and finally verifies whether the predicted decision was right.

Fig. 2 shows the computation schema with the predictor and the controller. There, ω^M_A is the total trust opinion about A. It is taken from the output of the trust system and fed back here to control the prediction. So, the predictability component of the trust model implements a feedback loop together with a controller to benefit from the methods of control theory.

The prediction goes in the same way as the decision making (for example, as described in Sect. V-C. In general, a decision making is represented by an operator D from a set of opinions $\{\omega^1_x, \ldots, \omega^N_x\}$ about a certain statement x to a tuple d, which contains the decision whether to belief x_i and the associated opinion ω^M_x of the car M. More specifically, if the car M applies the operator D on three times the opinion of A (ω^A_v), the result can be to accept A's opinion or note. This can then be compared with M's real opinion about v.

If the decision to accept or not accept the value proposed by A was right, then A was predictable. So, $\omega^M_{P,A}$ is right; there was no prediction error ($e = 0$). Otherwise, if the decision to accept or not accept was wrong, M must compute a prediction error e and readjust the relation $c = \frac{b_+}{b_-}$ between the belief and disbelief in $\omega^M_{P,A}$ to obtain a new $\omega^M_{P,A}$ (see Fig. 2).

The prediction error is not simply 0 or 1 if the decision was right or wrong. Instead the degree of the error depends on the predicted opinion and M's own opinion about v:

$$e = \hat{b}^M_v - b^M_v \ .$$

For the readjustment, this paper proposes a PI controller (proportional and integral controller) with the output

$$y_k = K_P\, e_k + K_I \sum_{i=1}^{k} e_i \ ,$$

where k is the step counter. It is increased with every piece of evidence. (So, it represents the number of evaluated reports.) With the controller output, a new c for the next evidence can be computed by

$$c_{k+1} = \begin{cases} (1 + y_k)\, c_k & \text{for } y_k \geq 0 \\ \frac{1}{1 - y_k}\, c_k & \text{for } y_k < 0 \ . \end{cases}$$

To compute $b^M_{P,A}$ from c, the uncertainty $u^M_{P,A}$ is necessary. It reflects the number of evidences; so it is known and always equal to $u^M_{C,A}$. With both together, the new predictability opinion can be obtained with

$$\omega^M_{P,A} = \left(\left(\frac{c\,(1 - u^M_{P,A})}{1 + c},\ \frac{1 - u^M_{P,A}}{1 + c} \right),\ u^M_{P,A} \right) \ .$$

In the proposed trust model, the predictability opinion should make the model more sensitive to outliers in the other car's behaviour. So, the control must be more sensitive to wrongly accepted reports than to wrongly rejected reports. Therefore, the simulation for this paper has been performed with $K_P = 0.6$ and $K_I = 0.2$ if the report has been wrongly accepted. In the other case, $K_P = 0.2$ and $K_I = 0.2$.

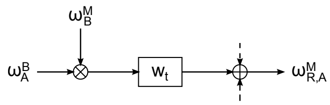

Fig. 3. Signal flow for the reputation computation. To aggregate the reputation of A, M discounts all reputations and accumulates them then (weighted by their age).

C. Evaluation of the Local Reputation

Reputation reflects what people say about another person or object. It condenses the general opinion of a group or society – not only of one or two individuals. So in technical systems, it is mostly realised as a feature of the network, which collects all recommendations to compute a reputation from them. For example, eBay [10] and the beta reputation system [11] do so. This paper focuses on individual-level trust, though. Here, the individual collects recommendations from others to build an opinion about another cars reputation. This individual reputation opinion may be different from the "real" global reputation. So, the authors call it *local reputation*, when it is helpful to distinguish both.

Reputation systems have to face the problem of statistically dependent recommendations [8, 17]. This is especially necessary in decentralised systems. The proposed reputation mechanism takes three measures to handle this: Firstly, only the most recent recommendation of one source about a third car is saved. Secondly, a car should only spread a recommendation about cars, with which it had at least one direct interaction. And thirdly – which is the most important measure –, reputation plays only a subordinate role in the opinion formation as detailed in the next section.

With these measures, the reputation can be computed as in other reputation systems (see Fig. 3): Every recommendation is discounted with the opinion about the source of the recommendation. So, the more the car M trusts the car B, the more weight B's opinion gains in the fusion process. This is done by the uncertainty favouring discounting as discussed in Jøsang et al. [17]. This operation weights the opinion of trustworthy cars higher in the subsequent opinion fusion. Further, the reputation opinions are weighted according to their age as described in Sect. IV-F. This whole process finally ends in a single estimate $\omega_{R,A}$ about A's reputation.

D. Combining Everything Together in a Trust Value

With the previous three sections, the car M has opinions about three properties of A: its competence, its predictability and its reputation. This section describes how these opinions are transformed into a trust value which quantifies the sender-related uncertainty of values received from A.

When looking at the three properties of the previous sections, competence and predictability are obtained by the car itself from direct interactions. As both are based on the same evidences, they must be combined with the averaging fusion. This is shown in the upper branch of Fig. 4.

However, only few evidences exist to evaluate competence and predictability. In contrast, recommendations of many cars are usually available resulting in an opinion about another car's reputation with small uncertainty. But reputation has the disadvantage that it is prone to attacks of others (see the requirements in Sect. I-B). That is why a fusion method is needed that takes advantage of the dense network offered by the reputation mechanism and of the reliability that comes along with the competence and predictability opinions.

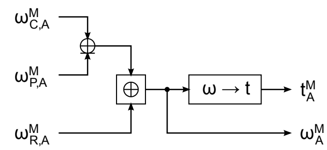

Fig. 4. The three components competence, predictability, and reputation are combined in an overall opinion. This opinion represents the trust in the car A. For deciding, it is transformed in the one-dimensional trust value t_A^M.

The authors propose a new operator for subjective logic, the *priority fusion*. It is defined as

$$u = \frac{u_P \cdot u_D}{u_P + u_D - u_P \cdot u_D}$$
$$b_{x_i} = b_{P,x_i} + (u_P - u)\,\frac{b_{D,x_i}}{b_{D,x_i} + d_{D,x_i}}\ ,$$

where the subscript P denotes the prioritised opinion and D the discriminated one. This way the prioritised opinion – gained by the car itself as a combination of competence and predictability – is not influenced by the recommendation of other cars. Even though, the consideration of reputation as the discriminated opinion ensures that all available information is still used. Furthermore, it allows an opinion about another car if no or only few own evidences exist.

At this point, the car M has computed its opinion about the trustworthiness of the other car A. An opinion is a two-dimensional magnitude, though. Sometimes – for example, to make a decision – a one-dimensional trust value is preferable. So, the final computation step in Fig. 4 transforms the trust opinion ω_A^M in a trust value t_A^M. Such a transformation already exists: for example, the expectation value of subjective logic. Another one with roots in information theory is proposed in the following section.

E. Opinions and the Binary Error and Erasure Channel

The opinion of the car M about the trustworthiness of another car A means whether M expects that A sends reliable information or not. The belief b_+ expresses the degree whether M expects to receive correct information, the disbelief b_- whether M expects to receive wrong information. The uncertainty u comes from a lack of knowledge; so, it is no expectation at all. This setting is similar to the model of the binary error and erasure channel shown in Fig. 5 on the following page [18]. Its channel capacity

$$C = (1-u)\left(1 - H_b\left(\frac{b_+}{1-u}\right)\right)$$
$$\text{with } H_b(x) = -x\,\log_2(x) - (1-x)\,\log_2(1-x)$$

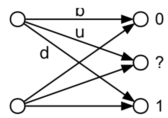

Fig. 5. The binary error and erasure channel. The disbelief d in the trustworthiness of another car A can be interpreted as an error rate of a transmission channel, the uncertainty u as an erasure rate.

expresses how much information can be correctly transmitted. It can be considered as a trust value t_A^M derived from the opinion $\omega_A^M = ((b_+, b_-), u)$. If the disbelief is higher than the belief, a negative sign is placed in front. Then, the trust value has a range from -1 to 1.

What is the advantage of using the channel capacity over the (commonly used) expectation value of opinions? The ratio between both (R is the trust value based on the expectation value as given in [11])

$$\frac{C}{R} = \frac{b_+ + b_-}{b_+ - b_-} \left(1 - H_b\left(\frac{b_+}{1-u}\right)\right)$$

has the following properties: It remains constant if the ration $\frac{b_+}{b_-}$ remains constant, that is, if only the uncertainty is varied. Otherwise the channel capacity is more sensitive to higher disbelief values; it punishes high disbelieves and rewards low disbelieves stronger than the expectation value. As Sect. VI shows this property makes it easier to separate agents with different error rates. So, it helps in the decision process.

F. Temporal Weighting of Opinions

Recent evidences should influence a present decision more than old evidences. This lets a system better adapt to changes in the environment: Trustworthy cars can become untrustworthy, and untrustworthy cars must have a chance to become trustworthy. So, a car must weight new opinions higher than older ones, when it fuses them in an overall opinion. This weight must depend on the age of the opinion. For the weighting operation, the certainty scaling fits well when the opinions are fused with the cumulative fusion.

The mapping from the age of an opinion to its weight should have the following properties:

1) At age $t = 0$, the mapping should be 1.
2) For $t \to \infty$, it should tend to zero.
3) In between, it should be monotonic.

Three mappings with the proposed properties are common: From the set of exponential functions, only $w_t = b^{\frac{t}{1\,\text{year}}}$ (and its equivalent formulations) have the required properties. The rational polynomials $w_t = \frac{a}{t+a}$ and the piecewise functions

$$w_t = \begin{cases} m\,t + 1 & \text{for } t < -\frac{1}{m} \\ 0 & \text{else} \end{cases}$$

with $m < 0$ fit here well, too.

Figure 6 compares these weighting functions. It shows that the rational polynomial decreases faster than the exponential functions in the beginning, but becomes flatter, then, and tends more slowly to zero in the end. The linear function with $m = -1/3$ is a good approximation for the exponential functions. The described properties are already significant within the lifetime of a car.

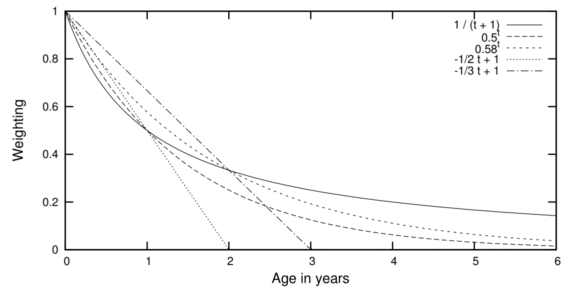

Fig. 6. Examples of temporal weighting functions. The exponential functions (dashed lines) decreases slower in the beginning, but show lower weighting values in the end compared to the rational polynomial (solid line). The linear functions approximate the exponential functions from a global view.

All in all, the rational polynomial does not suppress old evidences sufficiently. The exponential functions show a suitable form. In addition, they can be used in a recursively updating algorithm, because $w(t) = w(\Delta t_1) \cdot w(\Delta t_2)$ for $t = \Delta t_1 + \Delta t_2$. (A recursive update algorithm can help to save memory.) On the downside, they are computationally expensive. In contrast, the linear functions can be computed quickly and can be adjusted to have a sufficiently good form. They cannot be updated recursively, though.

V. Method of the Model Evaluation

In order to test the performance of the trust model the authors have created a simulation environment with respect to the characteristics of the cars' driving and communication behaviour. The simulation code is publicly available at www.ldv.ei.tum.de/Members/wbam/fidens for extension or verification by the research community.

A. The Social Structure in the Simulation

As pointed out in the introduction, a simulation environment for trust models must reflect the underlying social structure. Here, this structure is constituted by the movements of the cars and their communication. Because these movements are not random, they lead to regular meetings of cars. For example, some cars could use the same parking site and have similar trip times. The regular meetings are important so that trust can develop.

We developed a suitable traffic scenario, which simulates an urban population of a small town with about 15 000 inhabitants. It only contains trips to work and back home; however, these incorporate the majority of the regular ones and also allow random meetings. The duration of 23 weeks is pretty long compared to other simulations but it is necessary because developing trust takes some time. The simulation is performed with the microscopic traffic simulator SUMO [19].

The cars' communication is then simulated with the network simulator Shawn [20]. It is very efficient and suitable for our case, because it simulates the effects caused by a

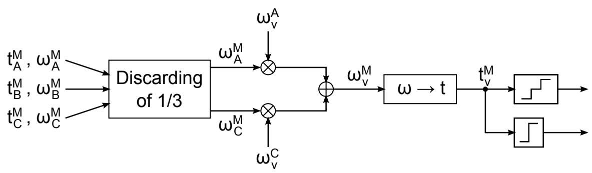

Fig. 7. Signal flow of the decision process. First, only the most trustworthy cars are extracted. Then, their opinions are discounted and combined. This also includes M's own observation. The result is M's opinion about v. It can be transformed in the trust value t_v^M for decision making.

phenomenon, not the phenomenon itself. When cars meet, they only exchange own observations about the environment or recommendations that origin from own experiences. Information of a third party is not yet relayed.

Both together, the cars' movements and their communication, represent the social structure.

B. The Information Model in the Simulation

As described in Sect. I, there is much information vehicles can exchange. Changes in the environment are modelled as events at randomly chosen locations in the simulation. These events are implemented as special network nodes. They send out binary values, which represent for example that there is a sign or not, there is free parking space or not, there is a point of interest or not, etc. Ten such events occur per day and each exists for one day. This frequency corresponds approximately to that of fuel prices or free parking lots. In reality, more events can be expected; this would then increase the trust model's performance.

C. The Cars' Behaviour in the Simulation

A trust model should help to maintain a more correct knowledge base. The process to determine knowledge from received information must therefore incorporate the trustworthiness of the senders. A car does this in a way that illustrated in Fig. 7. It is similar to the method in [9]. First, the car takes all messages regarding a specific event and discards the third of them that comes from the senders with the lowest trust values. Afterwards the remaining opinions are discounted with the opinion about its sender's trustworthiness. This has the effect that good known and well behaving vehicles get a higher weight. The discounting must be uncertainty favouring discounting as given in Sect. III-B. Finally the opinions about the event are combined into one using cumulative fusion.

A car must also decide, when to process a received report in the way described in trust model. In general, it can do this as soon as its own opinion about the included information is certain enough to reliably judge about the report. In the simulation, the car does this every time it observes the related event on its own. Whether this observed value is right depends on the error rate of the involved sensors, though. (Artificial reasoning has the property to be non-monotonic [21]. This means that new evidences can come in and lead to a completely different judgement of the situation. So, it is mostly a good idea to collect as much evidences as possible before judging about the trustworthiness of another car. But certainly, a trust model is only useful when it is used. So, a car must judge some when.)

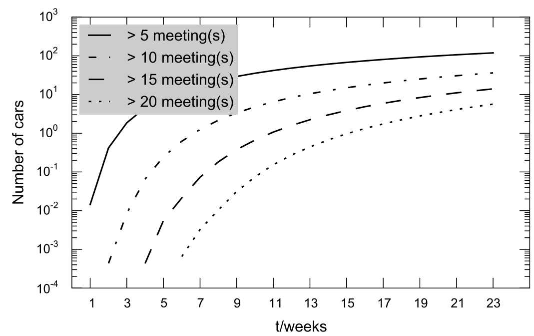

Fig. 8. This graph shows how often a prototypical car meets other cars. For example, on average a car meets about two cars at least once a week.

VI. Evaluation Results and Discussion

The simulation gives insights of the social structure and the trust model. Fig. 8 illustrates how often cars meet each other. The values are mean values over all cars; they represent a prototypical car. For example, after 23 weeks this prototypical car has met about five other cars at least 20 times. These are something like "good friends" that it meets once a week. The figure also visualises how the system settles down over time.

Fig. 9 on the following page visualises the competence property of the trust model. In the simulation, cars have different error rates between 0.0 and 0.2. The figure shows that a car assigns different trust values to other cars depending on this error rate. So, the competence indeed reflects the mean error of the other car and it is a good indicator for the reliability of the information source, when the car fuses information from several sources and decides about them. The figure also shows that a car with a higher error rate cannot estimate the competence of other cars as good as cars with a lower error rate. This is, because their reference values are more often defective.

The trust value in the upper diagram has been computed with the formula of the channel capacity. The solid line spans a value range of about 0.45. In contrast, in lower diagram the trust value is calculated with the formula of the expectation value. Here, the value range is only 0.3. This is the reason, why the channel capacity is a better transformation for the trust value computation than the expectation value. It makes it possible to better separate cars with different error rates.

The predictability is hard to visualise, because it cannot be presented by averages of the trust value. The feedback control has an effect only if the competence does not reflect the other's behaviour anymore. Then, it helps to more quickly adapt to the new situation than the competence alone. Exemplary numbers from the simulation show this: If a car with a good competence opinion $\omega_C^A = (0.84, 0.08, 0.08)$ and a predictability opinion $\omega_P^A = (0.32, 0.60, 0.08)$ $(c = b_P/d_P = 0,53)$ sends a report with a high belief error of 1, the predictability opinion may

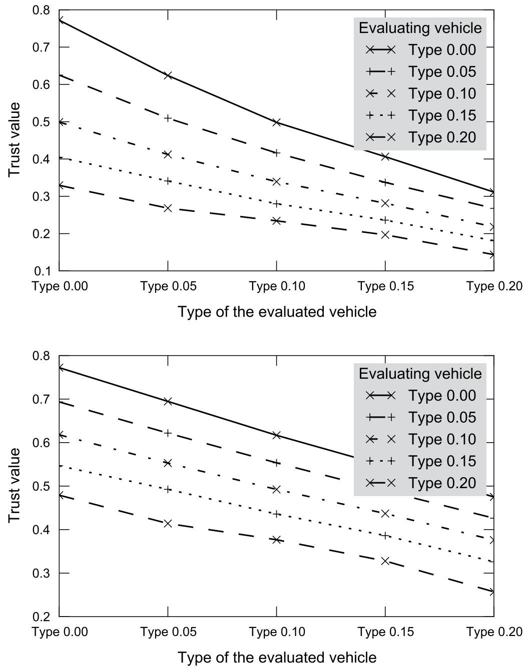

Fig. 9. The trust value of the competence depending on the error rate of the car. Cars of all types (different lines) can clearly distinguish the error rate of other cars (x-axis) based on the competence opinion. In the top diagram, the trust value is based on the formula of channel capacity. In the bottom diagram, the trust value is based on the formula of the expectation value. In this graphic, the value range is smaller than in the graphic above.

immediately fall to $(0.23, 0.69, 0.08)$ ($c = 0.33$). In contrast, the competence changes only slightly to $(0.80, 0.12, 0.08)$.

VII. Conclusion and Future Work

In this paper, we systematically modelled trust as a social mechanism to handle the uncertainty in the information exchange. It integrates different components to make it suitable for different situations. It also features the capacity formula of the binary error and erasure channel as an appropriate transformation of a two-dimensional trust opinion in a one-dimensional trust value. We also proposed a traffic simulation that partly represents the "social structure" between the cars. Thus, it is suited to investigate a social phenomenon like trust.

In the future, we plan to apply learning algorithms to the trust development to make the model even better adapt to different situations. We also intend to extend the simulation environment by more types of car behaviours. There should be car types that really try to exploit the weaknesses of trust models like human beings do.

Acknowledgment

The authors would like to thank Thomas Knauer; he implemented the software that evaluates the output of the simulation. We also thank the authors of SUMO and Shawn for providing their simulators to the research community.

This work is part of the project *Fidens: Trust for Cognitive Vehicles* [22]. It has been financially supported by the DFG Cluster of Excellence "Cognition for Technical Systems" (www.CoTeSys.org).

References

[1] C. Stiller, G. Färber, and S. Kammel, "Cooperative cognitive automobiles," in *Intelligent Vehicles Symposium, 2007 IEEE*, June 2007, pp. 215–220.

[2] CAR 2 CAR Communication Consortium, "Manifesto," August 2007, technical report, version 1.1. [Online]. Available: http://www.car-to-car.org/fileadmin/downloads/C2C-CC_manifesto_v1.1.pdf

[3] K. Matheus, R. Morich, and A. Lübke, "Economic background of car-to-car communication," in *Proceedings of the 2nd Braunschweiger Symposium Informationssysteme für mobile Anwendungen (IMA 2004)*, Braunschweig, Germany, 2004.

[4] M. Kranz, "Foundations of and applications for car-to-car communication," in *RADCOM 2008 Radar, Communication and Measurement*. Martinsried, Germany: GEROTRON COMMUNICATION GmbH, 2008.

[5] M. Raya and J.-P. Hubaux, "Securing vehicular ad hoc networks," *Journal of Computer Security*, vol. 15, no. 1, pp. 39–68, 2007.

[6] P. Golle, D. Greene, and J. Staddon, "Detecting and correcting malicious data in VANETs," in *VANET '04: Proceedings of the 1st ACM international workshop on Vehicular ad hoc networks*. New York, NY, USA: ACM, 2004, pp. 29–37.

[7] M. Raya, P. Papadimitratos, V. Gligor, and J.-P. Hubaux, "On data-centric trust establishment in ephemeral ad hoc networks," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, April 2008, pp. 1238–1246.

[8] S. D. Ramchurn, D. Huynh, and N. R. Jennings, "Trust in multi-agent systems," *The Knowledge Engineering Review*, vol. 19, no. 1, pp. 1–25, 2004.

[9] W. Zhang, S. K. Das, and Y. Liu, "A trust based framework for secure data aggregation in wireless sensor networks," in *Proceedings of the 2006 3rd Annual IEEE Communications Society on Sensor and Ad Hoc Communications and Networks (SECON '06)*, S. K. Das, Ed., vol. 1. IEEE, September 2006, pp. 60–69.

[10] eBay Inc., "All about feedback," Online, 2010, visited on 2010/03/02. [Online]. Available: http://pages.ebay.com/help/feedback/allaboutfeedback.html

[11] A. Jøsang and R. Ismail, "The beta reputation system," in *15th Bled Electronic Commerce Conference*, 2002.

[12] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Computer Networks and ISDN Systems*, vol. 30, pp. 107–117, 1998, special issue with the Proceedings of the Seventh International World Wide Web Conference.

[13] J. Sabater and C. Sierra, "Review on computational trust and reputation models," *Artificial Intelligence Review*, vol. 24, no. 1, pp. 33–60, 2005.

[14] A. Jøsang, "A logic for uncertain probabilities," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 9, no. 3, pp. 279–311, June 2001.

[15] ——, "Probabilistic logic under uncertainty," in *Thirteenth Computing: The Australasian Theory Symposium (CATS2007)*, ser. CRPIT, J. Gudmundsson and B. Jay, Eds., vol. 65. Darlinghurst, Australia: Australian Computer Society, Inc., 2007, pp. 101–110.

[16] G. Shafer, *A mathematical theory of evidence*. Princeton, NJ: Princeton University Press, 1976.

[17] A. Jøsang, S. Marsh, and S. Pope, "Exploring different types of trust propagation," in *Trust Management. 4th International Conference, iTrust 2006, Pisa, Italy, May 16-19, 2006. Proceedings*. Berlin: Springer, 2006, pp. 179–192.

[18] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 2006.

[19] D. Krajzewicz, G. Hertkorn, C. Rössel, and P. Wagner, "SUMO (Simulation of Urban MObility): An open-source traffic simulation," in *Proceedings of the 4th Middle East Symposium on Simulation and Modelling (MESM 2002)*, 2002, pp. 183–187.

[20] A. Kröller, D. Pfisterer, C. Buschmann, S. P. Fekete, and S. Fischer, "Shawn: A new approach to simulating wireless sensor networks," in *Proceedings of the 2005 Design, Analysis, and Simulation of Distributed Systems Symposium (DASD'05)*, 2005, pp. 117–124.

[21] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA, USA: Morgan Kaufmann, 1988.

[22] "Fidens: Trust for cognitive vehicles," 2010. [Online]. Available: http://www.ldv.ei.tum.de/Members/wbam/fidens