

RESEARCH

Güven: Estimating Trust from Communications

Anup K Kalia^{1*}, Zhe Zhang² and Munindar P Singh¹

*Correspondence:
akkalia@ncsu.edu

¹Department of Computer
Science, North Carolina State
University, , 27695 Raleigh, USA
Full list of author information is
available at the end of the article

Abstract

The extent to which an agent trusts another naturally depends on the outcomes of their interactions. Previous computational approaches have treated the outcomes in a domain-specific way. Specifically, these approaches focus on the mathematical aspect and assume that a positive or negative experience can be identified without showing how to ground the experiences in real-world interactions, such as emails and chats. We propose *Güven*, an approach that relates trust to the domain-independent notion of commitments. We consider commitments since commitment outcomes can be associated with experiences and a large body of works exist on commitments that include commitment representation and semantics. Also, recent research shows that commitments can be extracted from interactions, such as emails and chats. Thus, we posit *Güven* can provide an useful basis to infer trust between agents from their interactions.

To evaluate *Güven*, we conducted empirical studies of two decision contexts. First, subjects read emails extracted from the Enron dataset (and augmented with some synthetic emails for completeness), and estimated trust between each pair of communicating agents. Second, the subjects played the Colored Trails game, estimating trust in their opponents. *Güven* incorporates a probabilistic model for trust based on commitment outcomes; we show how to train its parameters for each subject based on the subject's assessments. The results are promising, though imperfect. Our main contribution is to launch a research program into computing trust based on a semantically well-founded account of interpersonal interactions.

Keywords: Multiagent systems; Commitments; Trust

1 Introduction

We consider multiagent system settings where agents interact with each other. A multiagent system is an open system consisting of autonomous and heterogeneous parties or agents. By autonomy, we mean that agents can act independently. And, by heterogeneity, we mean agents have diverse internal representations, including goals and internal policies. We consider a multiagent system to be open: agents may potentially enter such a system interact with others, and leave the system. Real-world examples of such systems arise in the corporate and military sectors where agents collaborate with each other in teams. In such systems, based on their mutual interactions, an agent as a *truster* estimates (and continually revises) its trust for another agent as a *trustee*. For example, in a corporate setting, an employer (truster) can assign a task to an employee (trustee). If the employee performs the task, the employer's trust increases for its employee. Similarly, in the military, a commander (truster) can ask a subordinate (trustee) to destroy a particular target. If the subordinate success, the trust of the commander toward the subordinate presumably increases.

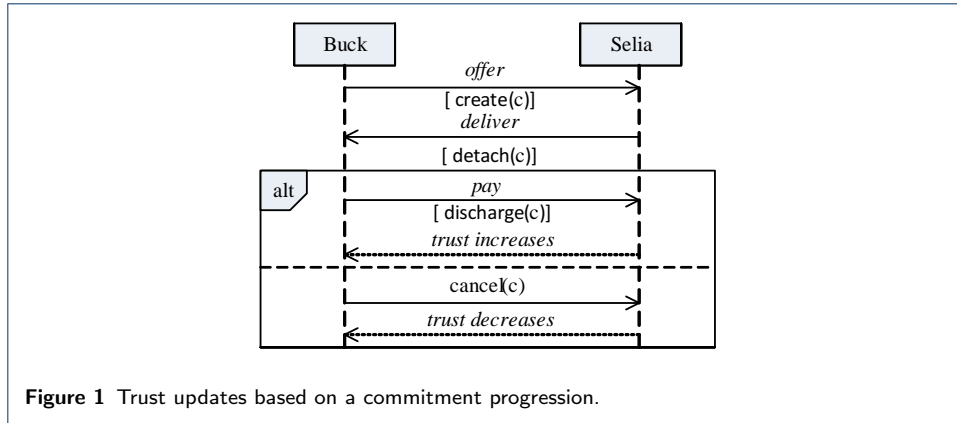
Understanding such interactions between agents and estimating trust from them is an interesting and challenging topic. Scissors *et al.* [23] exploit linguistic similarity in chat messages to estimate trust between message senders and receivers. Adalı *et al.* [1] calculate relationship strength between two users in Twitter based on social and behavioral aspects such as their numbers of friends and followers, the number of messages they exchange, and the time delay between the messages exchanged. DuBois *et al.* [8] provide an algorithm to compute trust and distrust in a social network. Wang *et al.* [31] combine positive and negative evidence to estimate trust. Teacy *et al.* [28] formulate trust as the count of fulfilling or violating obligations. Jøsang [11] represents trust as the belief measure of a truster that the trustee will cooperate. The above approaches are promising but they are limited to numerical heuristics. Such approaches have been justifiably criticized by richer approaches [4, 5, 26] for missing the essential intuitive considerations of trust, e.g., regarding the autonomy of the participants and the vulnerability of the truster to decisions by the trustee. The truster’s vulnerability refers to his or her willingness to take a risk on the trustee with the expectation that the trustee will perform the task promised to the truster [19]. However, the richer approaches have limitations since although they are formally represented, they do not lend themselves to computational techniques that could be applied in practice.

We seek to bridge the above gap between theory and practice. Specifically, we propose *Güven*,^[1] a computational model of trust founded on commitments that supports how agents determine trust in others based on their interactions. Commitments are important for trust because they can be identified from interpersonal interactions and can help us characterize the outcomes of such interactions in high-level terms. We limit *Güven* to commitments, although it can be potentially extended to related concepts such as prohibitions and authorizations.

A commitment $C(\text{debtor}, \text{creditor}, \text{antecedent}, \text{consequent})$ means that the debtor commits to bringing about the consequent for the creditor provided the antecedent holds. For example, $C(\text{Buck}, \text{Selia}, \text{deliver}, \text{pay})$ means that Buck (buyer) commits to Selia (seller) to paying a specified amount provided Selia delivers the goods. When Selia delivers, the commitment is detached. When Buck pays, the commitment is discharged or satisfied. If Selia delivers but Buck does not pay, the commitment is violated. In essence, a commitment describes a social relationship between two persons giving a high-level description of what one expects of the other. As a result, it is natural that commitments (and their satisfaction or violation) be useful as a basis for trust. In the above example, if Buck discharges the commitment, he brings a positive experience to Selia and Selia’s trust for Buck may increase; if Buck violates the commitment, he brings a negative experience to Selia and Selia’s trust for Buck may decrease. Figure 1 illustrates the intuition graphically.

Despite the apparent match, few approaches relate trust and commitments. Singh [26] and Chopra *et al.* [5] relate trust and commitments in terms of logical postulates and from an architectural perspective. In contrast, we understand trust and commitment in probabilistic terms, considering the outcomes of specific commitments and their effect on the trust relationships between the concerned parties.

^[1]From the Turkish word that brings together the concepts of trust and reliance.



We conduct two empirical evaluations, respectively, on emails (automatically analyzed using Kalia *et al.*'s approach [13]) and via the Colored Trails cooperation game [10] (analyzed manually). We show how to train the model parameters so as to capture a user model indicating each user's propensity to trust given commitment outcomes. Our evaluations yield promising, though imperfect, results on the viability of inferring trust from the commitments arising in interactions, suggesting the need for better extraction techniques. Our main contribution is to show how trust can be computed via the domain-independent concept of commitments. The contribution takes a step further to bring existing theories into practice.

This paper is organized as follows. We begin with a review of the related work. Next, we provide essential background on commitments and intuitions regarding how commitments affect trust. Then, we describe Güven as an evidence-based approach for updating trust based on commitments along with the requisite computational methods. We present our hypotheses in informal terms along with a strategy for evaluating them, followed by our evaluation and results. We conclude with a discussion and future directions.

2 Related Work

Teacy *et al.* [28] provide a trust model based on fulfilling or violating obligations. In their model, the trust of a truster (a_{tr}) toward a trustee (a_{te}) is represented as the expected probability that the trustee will fulfill its obligations ($B_{a_{tr}, a_{te}}$) given the outcome of all interactions $O_{a_{tr}, a_{te}}^{1:t}$ between the truster and the trustee. In contrast to their work, we consider representing trust as the basis of commitments (instead of obligations). We consider commitments since there has been an extensive work on capturing and formalizing the semantics of commitments in multiagent interactions [24, 25]. In addition, the lifecycle of commitments [29, 14] describe how commitments are created and progressed in multiagent interactions. Further, we evaluate our models on an email and a game dataset proving that such models are applicable to real-world settings.

Wang *et al.* [31] define trust (α) as the ratio of positive outcomes experienced (r) by the truster from the trustee to the total number of positive and negative outcomes (s), i.e., $\alpha = \frac{r}{r+s}$. Further, to denote the certainty of the truster toward a trustee, Wang *et al.* [31] define a certainty function ($c(r, s)$) that employs a beta

distribution and takes r and s as its input parameters. There are important limitations with Wang *et al.* [31] contributions: (1) the model is purely mathematical and hence, does not clearly describe how it can be applied to real-life interactions such as emails and chats. Specifically, the definitions of positive and negative outcomes are not formally captured with respect to agents' interactions; (2) the model needs the initial trust to be manually set. That is, it assumes fixed values. Specifically, Wang *et al.* mention that the initial trust is set based on the truster's prior experience with the trustee. Compared to their model, we clearly define positive and negative evidence as outcomes of commitments thereby extending their model to be applicable to real-world settings. Further, our model learns the initial trust from users' data instead of assuming fixed values.

Osman *et al.* [21] describe a model that estimates the trust of a truster based on the trustee's capability and willingness to execute a commitment. The willingness of a trustee to execute the commitment is computed using the trustee's past behavior in executing similar commitments. The capability of the trustee is computed by matching the capability needed for the current commitment with the capabilities of the trustee observed in the past. Osman *et al.*'s model suffers from two important limitations: First, the model considers commitments as a set of actions to be performed by the trustee, thus omitting a formal representation as well as an operationalization (lifecycle) of commitments. This makes the model less intuitive with respect to its applicability on real-world data such as emails. Second, similar to Wang *et al.*'s [31] model, Osman *et al.*'s trust model needs the initial trust to be manually set. For example, the model considers the initial trust as zero which may not hold in different settings.

Kastidou *et al.* [15] describe a trust model based on promises made and delivered by a trustee toward a truster. Similar to Osman *et al.* [21] model, Kastidou *et al.*'s model does not provide the semantics of promises, requires manual setting of the initial trust values, and does not consider real-world datasets for the evaluation.

Burnett and Oren [3] examine the effects of delegation using a probabilistic trust model [12] and propose an approach for weighting trust updates based on shared responsibility or delegation. Burnett and Oren do not restrict the delegation chain length. In contrast, we restrict our trust update to delegation chains of length three (the debtor, the new debtor, and the creditor). For example, if the new debtor delegates the commitment to another debtor (debtor'), trust between debtor and debtor' remains unaffected. This means if debtor' satisfies or violates the commitment, the trust of the debtor toward debtor' neither increases or decreases, respectively, since the debtor still depends on the new debtor to satisfy the commitment and may be ignorant of debtor'. We observe that such longer chains are rare in a real-world text corpus.

Adahi *et al.* [1] correlate textual features (linguistics, processes, and psychological processes) with social and behavioral features (reciprocity, assortativity, attention, and latency). Textual features are derived based on the content of messages exchanged between users whereas social and behavioral features are computed based on user's social network (nodes, edges). Both the textual and behavioral features are indicator of trust behavior between users. However, such measures are considered based on their frequency of occurrence and they do not capture the vulnerability of a truster toward a trustee. For example, textual features indicate number

of verbs, pronouns, affective processes (emotions), cognitive processes (causation, certainty), perceptual processes (see, hear, feel), and so on whereas behavioral features indicate number of number of friends, followers, messages sent, degree similarity between users, reciprocity of conversation and propagation messages, in-degree, out-degree and so on. Thus, we emphasize capturing semantic meanings of messages in terms of commitments (i.e., how they are created, satisfied, and violated) from the text content exchanged between users.

Scissors *et al.* [23] performed an empirical evaluation with 62 students and found that different forms of linguistic similarity, such as content (e.g., positive emotion words, task-related words), structure (e.g., verb tense, phrasal entrainment), and style (e.g., chat abbreviations, emoticons) reflect different level of trust between participants. In contrast, we consider commitments between participants to estimate trust between them.

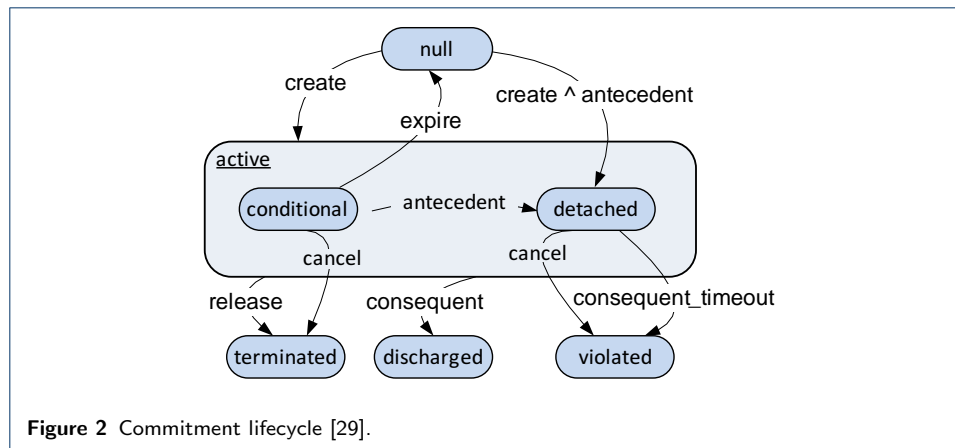
3 Background on Commitments

Figure 2 presents the commitment lifecycle we adopt. A commitment $C(\text{debtor}, \text{creditor}, \text{antecedent}, \text{consequent})$ is *created* when a debtor either voluntarily creates it (*commissive creation*) or is directed to do a certain task (*directive creation*). Given the debtor’s autonomy, the latter presumes a prior commitment on part of the debtor. A commitment is *detached* if a condition or an antecedent present for a commitment holds true; *discharged* when a debtor executes a committed task. A commitment is *terminated* when a debtor cancels the commitment before it is detached or a creditor releases the commitment. A commitment is *violated* when a debtor cancels the commitment after it is detached or when a consequent timeout occurs. Additionally, a commitment can be delegated and assigned. A commitment is *delegated* when the debtor of a commitment is replaced by a new debtor and *assigned* when the creditor of a commitment is replaced by a new creditor. We map interactions between persons to commitment operations.

Consider a commitment $C(\text{Alice}, \text{Bob}, \text{pay}, \text{ship goods})$ where Alice (buyer) commits to Bob (seller) to ship goods provided Bob pays. When Alice offers Bob to ship the goods, C gets created from Alice toward Bob and becomes conditional. When Bob makes the payments, C gets detached. If Alice ships the goods, C gets discharged. If Alice does not ship the goods despite the payment, C gets violated. If C is conditional and Alice cancels C, C gets terminated. Alice can delegate C to a new agent Charlie to create commitment $C(\text{Charlie}, \text{Bob}, \text{pay}, \text{ship goods})$ where Charlie becomes the new debtor. Similarly, Bob can assign commitment C to a new creditor John to create $C(\text{Charlie}, \text{John}, \text{pay}, \text{ship goods})$. In case of delegation and assignment, when a new commitment are created, the older commitment remains suspended and its outcome depend upon the outcome of the new commitment. If the new commitment is satisfied, so is the older commitment. If the new commitment is violated, the older commitment becomes active (conditional or detached, depending upon the truth of its antecedent) again.

4 Intuitions on Trust and Commitments

We describe some criteria for how trust values may be updated based on operations on commitments.



4.1 Effects of Commitment Operations on Trust

We describe the effect of commitment operations on trust. Before we describe the effects, let us consider some situations wherein a commitment exists from a debtor toward a creditor.

Effect of discharge. When a commitment is discharged, the creditor's trust for the debtor increases.

Effect of violation. When a commitment is violated, the creditor's trust in the debtor decreases.

Effect of delegation and discharge. When a commitment is delegated by the original debtor to a new debtor, and the new debtor satisfies it, the creditor's trust in both the original and the new debtor increases.

Effect of delegation and violation. When a commitment is delegated from the original debtor to a new debtor and the new debtor violates it, the creditor's trust in both the original and the new debtor decreases.

Effect of assignment and discharge. When the commitment is assigned from the original creditor to a new creditor and the original debtor discharges it, the trust of the original and new creditor for the original debtor increases.

Effect of assignment and violation. When the commitment is assigned from the original creditor to the new creditor and the debtor violates it, the trust of the original and the new creditor for the original debtor decreases.

We make the following assumptions regarding the increase or decrease of trust. These are simplified assumptions and could be relaxed in some settings.

- In our basic approach (baseline), the change in trust is the same for all commitment outcomes. We additionally provide an approach in which the change in trust depends upon the strength of a commitment. For example, a commitment with a strict deadline when satisfied may produce a different level of trust compared to a commitment with a flexible deadline.
- We assume commitment discharge and violation to be all or none; in our scenarios, partial success is not easy to infer.
- In case of violation, the trust of the creditor for the debtor decreases irrespective of whether the debtor was truly responsible. An agent's beliefs and goals are private and cannot be identified directly from his or her interactions.

- In case of delegation, the original creditor's trust in the original and the new debtor changes equally, reflecting the idea that the creditor has a positive experience thanks to the two debtors.

In case of a situation where the new debtor violates a commitment and the original debtor has nothing to do with it, it was still the responsibility of the original debtor to satisfy the commitment. (This is one of the patterns of delegation identified by Singh *et al.* [27].) Thus, the creditor's trust will decrease equally for both the new and the original debtor.

- In case of assignment, the new creditor's trust in the debtor changes as much as the original creditor's trust in the debtor, reflecting the intuition that both creditors' expectations are met.

4.2 Subjectivity, Memory, and Strength

Trust is modulated by features that affect how trusters judge outcomes, such as the satisfaction or violation of a commitment. First, trust assessment is subjective. Trusters differ in how they reward or penalize a trustee when a commitment is discharged or violated, respectively. Second, trust assessment depends on the truster's memory: trusters with limited memory would tend to forget all but (some varying number of) recent experiences. Recent experiences may turn out to be more predictive of future experiences (that trust is about) than past experiences. Third, the effect on trust of a commitment's outcome would be greater when the commitment is more important.

5 Güven: Model of Trust based on Commitments

We adopt Wang and Singh's [32] trust model, which represents trust as evidence $\langle r, s \rangle$. Here, $r \geq 0$ and $s \geq 0$ respectively represent the positive and negative experiences the truster has with the trustee. Both r and s are real numbers. Wang and Singh calculate trust as the probability of a positive outcome as $\alpha = \frac{r}{r+s}$. Suppose Buck and Selia transact 10 times and exactly eight transactions succeed from Selia's perspective. Then Selia's trust in Buck would be 0.8.

The basic idea is for each truster to maintain evidence $\langle r, s \rangle$ about each trustee. The initial evidence, $\langle r_{in}, s_{in} \rangle$, represents the truster's bias. An interaction may yield a positive, negative, or a neutral experience. In these cases, the evidence is updated by respectively adding $\langle i_r, 0 \rangle$, $\langle 0, i_s \rangle$, and $\langle \lambda i_r, (1 - \lambda) i_s \rangle$, where $\lambda \in [0, 1]$. In essence, we characterize each truster via five parameters $(i_r, i_s, r_{in}, s_{in}, \lambda)$.

5.1 Considering Subjectivity

To evaluate H_1 , we learn a specific truster's parameters based on positive, negative, and neutral experiences it acquires from trustees and the truster's actual trust in various trustees. For the k^{th} trustee, let α_k represent the truster's actual trust (as revealed) and $\hat{\alpha}_k$ the truster's predicted trust in k . Let E_k^+ , E_k^- , and E_k represent the numbers of positive, negative, and neutral experiences, respectively. Then,

$$\hat{\alpha}_k = \frac{r_{in} + E_k^+ i_r + \lambda \cdot E_k i_r}{r_{in} + s_{in} + E_k^+ i_r + E_k^- i_s + E_k (\lambda i_r + (1 - \lambda) i_s)} \quad (1)$$

Via nonlinear least-squares regression technique that uses trust region reflective algorithm [6], we estimate the truster’s parameters to minimize the mean absolute error (MAE) of prediction, $\sum_{k=1}^n |\hat{\alpha}_k - \alpha_k|$.

5.2 Considering Memory

We capture the effect of memory by considering a *discount window* [33], defined simply as the most recent W experiences. Let n be the total number of experiences the truster acquires from the trustee. Let $t = \min(n, W)$. Let E_t^+ , E_t , and E_t^- be the positive, neutral, and negative experiences inferred from the t transactions. The trust of a truster in the trustee depends on whether t is less than W . When $t < W$, the truster’s trust is $\langle E_t^+ i_r + \lambda E_t i_r, E_t^- i_s + (1 - \lambda) E_t i_s \rangle$; otherwise, it is $\langle r_{in} + E_t^+ i_r + \lambda E_t i_r, s_{in} + E_t^- i_s + (1 - \lambda) E_t i_s \rangle$. When $t < W$ we ignore the initial bias as the truster’s trust is based on recent W experiences, which simply means that the truster has already forgotten its initial bias.

5.3 Considering Strength

We posit that a truster acquires experiences of varying *weights* based on commitment outcomes (satisfied or violated). To calculate the weight, we identify the following features in a sentence indicating a commitment creation.

Commissive over directive. *A commissive (e.g., “I will ...”) may carry a greater weight than a directive (e.g., “Could you please ...”) because it holds even without the presumption of another commitment.*

Debtor’s type. *A single debtor may carry a greater weight than multiple debtors “we” (“We will follow up”). A single debtor, as in “I will follow up,” has clearer responsibility than multiple debtors.*

Creditor’s type. *Multiple creditors may carry a greater weight than a single creditor. Multiple creditors arise when a debtor commits to a set, e.g., when a product manager commits to his employees to review a product. The intuition is that having multiple creditors makes the debtor accountable to more parties.*

Modal verbs. *Some modal verbs (e.g., will or shall) may convey high confidence over others (e.g., can, could, may, would) [20]. The intuition is that “will” indicates that a commitment will be surely satisfied whereas “can” indicates that the commitment may not be satisfied. We learn the weights of different modal verbs based on data obtained from human subjects.*

Action verbs. *Some action verbs convey a greater level of importance than others. For example, “resolving an issue” may be more important than “reviewing a proposal.” We compute the weights of verbs using Burchardt et al.’s [2] FrameNet tool, which provides weights for words used in different senses, e.g., 1 for *resolve* and 0.383 for *review*.*

Deadlines. *Noun phrases with deadlines [7] may convey more importance than noun phrases without deadlines. For example, an explicit deadline, as in “I will repair the car by Monday,” enhances the importance of the commitment. We assume that merely mentioning a deadline increases the seriousness of a commitment. We defer to future research additional subtleties, such as the duration or urgency of a deadline and the extent to which it may be broken, since in our empirical settings durations and urgency do not arise.*

Except the feature *action verb*, we evaluate rest of the features empirically from a subject evaluation. We provide the outcome of our evaluation in Table 6. Our evaluation ranks feature values as discussed above. We map the ranks to weights (cardinal numbers, higher for higher ranked features) and sum the weights to compute a value. For example, consider two commitment creations from a trustee toward a truster: (1) “I will repair the car by Monday” and (2) “I can repair the car”. These examples have four common features: the commitment type (commissive), the debtor type (“I”), the creditor type (single), and the action verb (“repair”) and two different features: modal verbs (“will” versus “can”) and deadlines (“Monday” versus no deadline). Thus, based on the features from these two examples and the rank of those features provided in Table 6, we compute their weights as shown in Table 1. The weights computed (11 versus 8) indicate that the truster might have a different experience from the outcome of the first example than the second example.

Table 1 Computing commitment weights for two examples of commitment creations.

Features	Example 1	Example 2
1. Commissive, Directive	Commissive (2)	Commissive (2)
2. Debtor’s type	“I” (2)	“I” (2)
3. Creditor’s type	single (1)	single (1)
4. Modal Verbs	“will” (4)	can (2)
5. Action Verbs	“repair” (1)	“repair” (1)
5. Deadlines	“Monday” (1)	no deadline (0)
Total Weight	11	8

6 Evaluation Strategy

We consider two decision contexts for evaluations. The first context involves subjects (as bystanders) reading emails exchanged between other agents and assessing the levels of trust between these agents. The second context involves the subjects playing a game with each other. The game has some cooperative and some competitive elements. The subjects (as interested participants) assess the trustworthiness of their opponents.

We now present our research hypotheses in informal terms. These hypotheses are based on subjectivity, memory, and strength, as proposed in Section 5 to compute trust. Thus, these hypotheses motivate our evaluation strategy and study design. The following section on evaluation refines these hypotheses into technical claims.

H_1 (*Subjectivity*) Predicting trust values by *learning trust parameters* for each subject yields more accurate results than using *fixed trust parameters* for all subjects.

The details of the trust parameters are given in Section 5. Assuming H_1 holds, we consider learned parameters as the *baseline* approach. We check if other approaches improve accuracy beyond the baseline.

We consider H_1 since trust models [15, 21, 28, 31] described in Section 2 compute trust considering fixed parameters. That is, they consider fixed values for $\langle r, s, r_{in}, s_{in}, \lambda \rangle$. In contrast, we compute trust by learning these parameters from subjects’ assigned level of trust. Thus, we posit that learning the trust parameters would improve the trust prediction accuracy.

H_2 (*Memory*) Predicting trust values by learning a specific *discount window size* for each subject yields more accurate results than the baseline.

We consider H_2 since a subject might assign a trust level based on his or her most recent experiences [33]. Thus, we posit that learning a specific discount window size in addition to learning trust parameters for each subject can improve the trust prediction accuracy more than the baseline.

H_3 (*Strength*) Predicting trust values by inferring *strengths* of positive and negative experiences yields more accurate results than the baseline.

We consider H_3 since a subject might assign a trust level based on his or her varying experiences from different commitment outcomes. Thus, we posit that considering weights of such commitments in addition to learning trust parameters for each subject can improve the trust prediction accuracy more than the baseline.

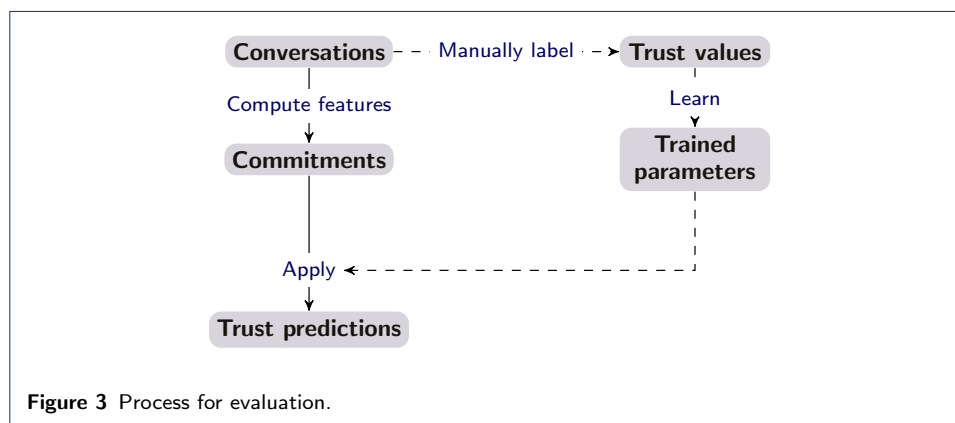
H_4 Subjects' trust assessment behavior as bystanders differs from their trust assessment behaviors as players. H_4 is measured in terms of the following sub-hypotheses, which posit that subjects' trust assessment behaviors vary across decision contexts.

H_{41} The correlation coefficients (R) between subjects' trust values and positive experiences in emails and the game, respectively, are different.

H_{42} The correlation coefficients (R) between subjects' trust values and neutral experiences in emails and the game, respectively, are different.

6.1 Evaluation Process

Figure 3 summarizes our evaluation process. Our evaluation strategy is to gather data from subjects in the two decision contexts and proceed as follows.



Step 1. Build a dataset of interpersonal interactions with trust values. For emails, subjects provide third-party assessments; for games, subjects provide their own trust assessments. Table 2 shows the examples of email and chat interactions. Based on the emails exchanged between Kim and Dorothy, the subjects assign trust values in the range of 0 and 1, from Kim toward Dorothy as well as from Dorothy toward Kim. Similarly, for chats P_1 assigns a trust value for P_4 and P_4 assigns a trust value for P_1 .

Step 2. Identify commitment operations from the interactions. For emails, using Kalia et al.'s [13] trained classifier; for games, we find these in the chat interface and analyze them manually. Table 2 shows examples of commitment operations identified in emails and chats, respectively.

Step 3. Partition the dataset into training and test datasets. Learn model parameters for each subject from the training data.

Step 4. Apply the learned model to predict trust in the test data and compute the model’s accuracy.

We repeat the process for all subjects and present our results.

7 Evaluation

We evaluated Güven via an empirical study with 30 subjects (graduate and undergraduate students from various academic departments at our university). We conducted the study in two phases. In the **first phase**, subjects read 33 emails selected from the Enron email corpus [9, 17] and provided a trust value ranging from 0 to 1 between the senders and receivers of email. The emails were selected on the basis of their containing sentences that indicate commitment creation, satisfaction, or violation—such sentences having been identified using Kalia *et al.*’s [13] method. We augmented the dataset with 28 synthetic sentences indicating commitment satisfaction or violation, which do not occur frequently in the corpus. Subjects provided trust values based on their intuitions by reading these emails. We did not disclose the commitments identified. We did not provide any additional guidelines that might restrict a subject’s individual perception of trust. Once the subjects provided their estimated trust values, we mapped commitment operations to positive, negative, and neutral experiences. Table 3 shows an example of two rows created from the first two interactions between Kim and Dorothy given in Table 2. S1, . . . , S6 in Table 3 represent the subjects who provided trust values based on the interactions between Kim and Dorothy. Based on the experiences collected from emails and trust values collected from subjects, we created 28 rows of data for each subject. Hence, for 30 subjects we obtain 28×30 or 840 rows. We provide additional details of our data, including a link to download it, in the Appendix (in Section 10.1).

Table 2 Examples of email and chat interactions.

Sender	Receiver	Email Content	Operation
<i>Kim</i>	<i>Dorothy</i>	<i>I will also check with Alliance Travel Agency . . .</i>	$\text{create}(C_1)$
<i>Kim</i>	<i>Dorothy</i>	<i>I checked with our Travel Agency . . .</i>	$\text{discharge}(C_1)$
<i>Rob</i>	<i>Kim</i>	<i>By Wednesday Aug 16 2001, please send all copies of your documentation . . .</i>	$\text{create}(C_2)$
<i>Kim</i>	<i>Rob</i>	<i>Rob, please forgive me for not sending this in by Aug 15</i>	$\text{cancel}(C_2)$
Sender	Receiver	Chat Content	Operation
P_1	P_4	<i>Can you please give 1 green tile?</i>	
P_4	P_1	<i>I can</i>	$\text{create}(C_3)$
P_4	P_1	<i>P_4 transferred P_1 1 green tile</i>	$\text{discharge}(C_3)$

Table 3 Different features and trust values from different subjects.

Trust Pairs	Experiences	S1	S2	S3	S4	S5	S6
Kim→Dorothy	2 Neutral, 0 Positive, 0 Negative	0.3	0.4	0.35	0.5	0.6	0.7
Dorothy→Kim	1 Neutral, 1 Positive, 0 Negative	0.7	0.9	0.8	0.7	0.6	0.95

Additionally, we asked subjects several questions based on the features discussed in Section 5.3. The questions were about ranking these features in order of their

perceived importance. We provide these questions in the Appendix in Section 10.2. Hypotheses H_1 , H_2 , and H_3 are about evaluating the proposed approaches with respect to their prediction accuracy.

For the **second phase**, we augmented Gal *et al.*'s [10] Colored Trails game with a chat interface. Each subject played three games of five rounds each. Figure 4 shows an instance of the game in progress. The game has a 4×4 board with a chat interface to communicate with one's opponent. In each round, subjects were allocated a fixed number of colored tiles, a starting position, and a common goal position on the board. To reach the goal, a subject must provide the requisite tiles. During the game, subjects can communicate and trade tiles via the chat interface. A subject could commit to an opponent to transfer specified tiles and could discharge or violate each commitment. After each round, subjects recorded their trust for their opponents on a five-point scale.

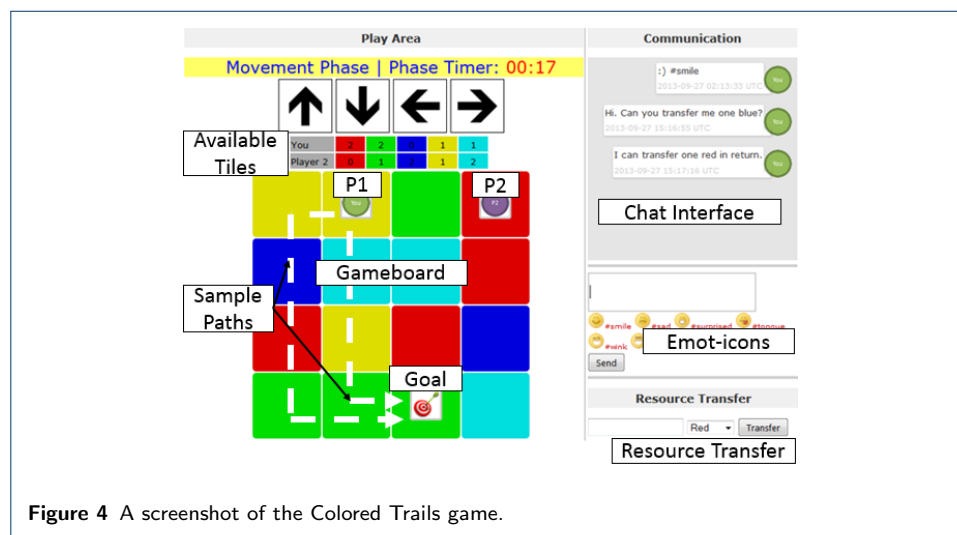


Figure 4 A screenshot of the Colored Trails game.

We randomly split our 30 subjects into five groups of six each, and each group into two subgroups of three subjects each. The players in any game sat in separate rooms and communicated through the chat tool. The subjects did not know the identities of the other subjects but knew they had the same opponent for all three rounds of each game. Hypotheses H_4 , H_{41} , and H_{42} are about evaluating whether a subject's estimation profile differs across decision-making contexts.

7.1 Results

[Verifying H_1] We collected the trust values from the subjects from the emails assigned to them. We divided the data collected from subjects into three-fold training and test data and learned trust parameters for each subject (r_{in} , s_{in} , i_r , i_s , λ) that minimize the mean absolute error (MAE) between predicted and actual trust values.

For verifying H_1 , we calculated the MAE for λ ranging from 0.1 to 0.9. Then, we calculated the MAE by learning the λ (Learned(λ)) itself. Based on the above MAEs, we obtained a customized λ (fixed or learned) for each subject. A customized λ for a subject refers to the value of λ for which the MAE is minimum. We represent the MAEs obtained using customized λ s for all subjects as Custom(λ) in

Figure 5. Finally, we arbitrarily assumed some fixed configurations of parameters ($F1 = \langle 1, 1, 1, 1, 0.5 \rangle$, $F2 = \langle 2, 1, 1, 1, 0.5 \rangle$, $F3 = \langle 1, 2, 1, 1, 0.5 \rangle$). F1 indicates no bias in the initial trust perception where as F2 and F3 indicate positive and negative biases respectively. $\lambda=0.5$ in fixed configurations indicates equal trust increments for the neutral experiences. The configurations can be changed by incrementing or decrementing different parameters. From the results, we observed that the median of Custom(λ) (0.162) is less than the medians of all other approaches. To verify if the difference is significant, we evaluated hypothesis H_1 via one-tailed t-test as shown in Table 4. From the results, we found that the difference is not statistically significant. Thus, we concluded that although the overall result (MAEs) seems to align with the hypothesis H_1 , i.e., learning the trust parameters yields more accurate results than using fixed parameters for all subjects, the significance result doesn't confirm it. We leave the evaluation of H_1 for further investigation.

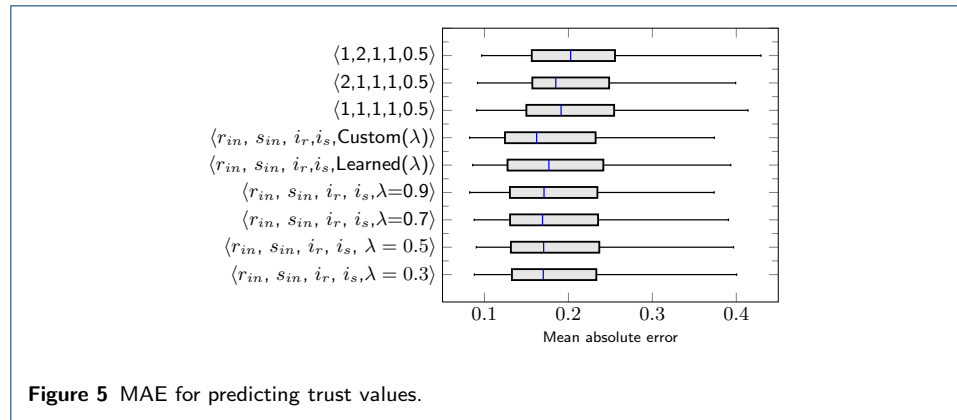


Table 4 Statistical test results for H_1 . (μ -mean of the MAEs)

$\mu_{C(\lambda)} < \mu_{others}$	$\mu_{C(\lambda)}$	μ_{others}	p-value
$\mu_{C(\lambda)} < \mu_{F3}$	0.18	0.21	0.14
$\mu_{C(\lambda)} < \mu_{F2}$	0.18	0.20	0.18
$\mu_{C(\lambda)} < \mu_{F1}$	0.18	0.21	0.07
$\mu_{C(\lambda)} < \mu_{L(\lambda)}$	0.18	0.19	0.34
$\mu_{C(\lambda)} < \mu_{\lambda=0.9}$	0.18	0.19	0.36
$\mu_{C(\lambda)} < \mu_{\lambda=0.7}$	0.18	0.19	0.34
$\mu_{C(\lambda)} < \mu_{\lambda=0.5}$	0.18	0.19	0.33
$\mu_{C(\lambda)} < \mu_{\lambda=0.3}$	0.18	0.19	0.34

[Verifying H_2] For verifying H_2 , first, we determined customized window sizes ($CW=1, 2, \dots, 12$) for each subject for all the values of λ (0.1, \dots , 0.9, learned λ). A customized window size (CW) for a subject refers to the value of CW for which the MAE is minimum. We found that if we increase the window size further (i.e., greater than 12), the MAEs for the subjects do not decrease. We obtained MAEs for all the subjects based on various values of λ and CWs , and represent them as Custom(λ)+ CW shown in Figure 6. We compared Custom(λ)+ CW with Custom(λ) obtained from H_1 and MAEs obtained from other approaches. From the results, we found that the median of Custom(λ)+ CW (0.153) is less than the median of Custom(λ) (0.162) and other approaches. From the one-tailed t-test results shown in Table 5, we found that the mean of Custom(λ)+ CW is significantly lower than the

means of approaches that consider fixed configurations (e.g., $\lambda=0.7+CW$). However, for others, the t-test results show that the differences are not significant. Thus, we concluded that although the overall result (MAEs) seems to align with the hypothesis H_2 , i.e., learning a specific discount window size yields more accurate results than the baseline, the significance result doesn't confirm it. We leave the evaluation of H_2 for further investigation.

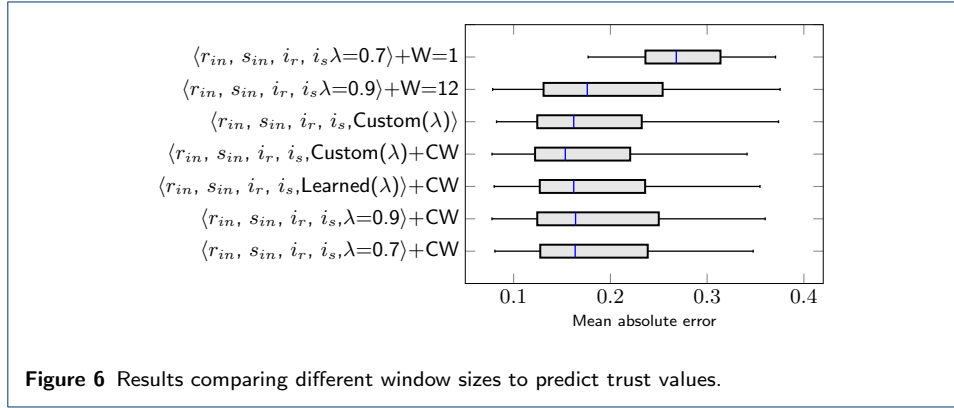


Figure 6 Results comparing different window sizes to predict trust values.

Table 5 Statistical test results for H_2 .

$\mu_{C(\lambda)+CW} < \mu_{others}$	$\mu_{C(\lambda)+CW}$	μ_{others}	p-value
$\mu_{C(\lambda)+CW} < \mu_{\lambda=0.7+CW=1}$	0.17	0.27	0.00
$\mu_{C(\lambda)+CW} < \mu_{\lambda=0.9+CW=12}$	0.17	0.20	0.08
$\mu_{C(\lambda)+CW} < \mu_{C(\lambda)}$	0.17	0.18	0.27
$\mu_{C(\lambda)+CW} < \mu_{L(\lambda)+CW}$	0.17	0.18	0.32
$\mu_{C(\lambda)+CW} < \mu_{\lambda=0.9+CW}$	0.17	0.19	0.17
$\mu_{C(\lambda)+CW} < \mu_{\lambda=0.7+CW}$	0.17	0.18	0.30

[Verifying H_3] From the first phase of our experiment, we obtained subjects' assessments of the weights of commitments. From their assessment we obtained different orderings among the feature values shown in Table 6. The ordering among the values for the creditor type shows that our initial assumption about it was incorrect. For the rest of the feature values, our assumptions correctly aligned with the subjects' assessments.

Table 6 Ordering among the feature values obtained from subjects' assessments.

Features	Ordering
1. Commissive, Directive	Commissive (2) > Directive (1)
2. Debtor's type	"I" (2) > "We" (1)
3. Creditor's type	multiple creditor (2) > single creditor (1)
4. Modal Verbs	"must, will" (4) > "would, should, shall" (3) > "can, could" (2) > "may" (1)
5. Deadlines	messages with "deadline" (2) > messages without "deadline" (1)

From the orderings shown in Table 6 we calculated a weight for each commitment (CWT). Based on the weights, we recalculated trust parameters using different λ s, namely, (0.1, ..., 0.9, learned λ). For H_3 , we calculated a customized λ (Custom(λ)+CWT) and compared the results with $C(\lambda)$ from H_1 . From the results shown in Figure 7, we found that the median of Custom(λ)+CWT (0.161) is slightly less than the median of Custom(λ) (0.162). CWT in the figure means considering

commitment weight. From the one-tailed t-test results shown in Table 7, the mean of Custom(λ)+CWT is not significantly lower than the means of other approaches. Thus, we concluded that although the overall result (MAEs) seems to align with the hypothesis H_3 , i.e., inferring strengths of positive and negative experiences yields more accurate results than the baseline, the significance result doesn't confirm it. We leave the evaluation of H_3 for further investigation.

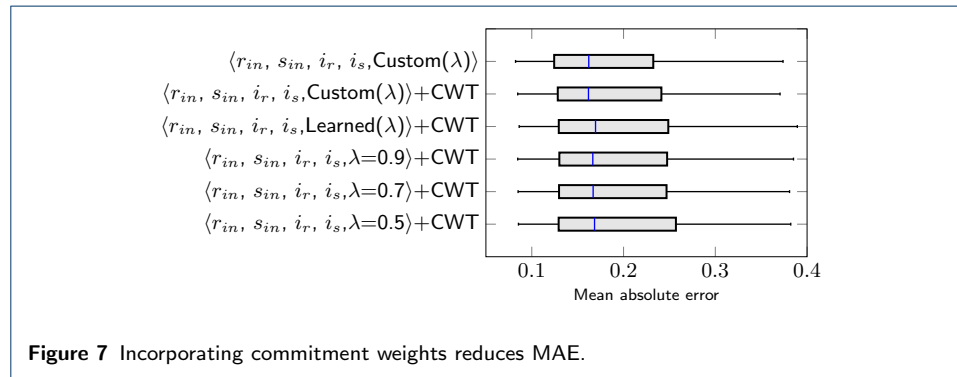


Table 7 Statistical test results for H_3 .

$\mu_{C(\lambda)+CWT} < \mu_{others}$	$\mu_{C(\lambda)+CWT}$	μ_{others}	p-value
$\mu_{C(\lambda)+CW} < \mu_{C(\lambda)}$	0.18	0.18	0.50
$\mu_{C(\lambda)+CW} < \mu_{L(\lambda)+CWT}$	0.18	0.19	0.37
$\mu_{C(\lambda)+CW} < \mu_{\lambda=0.9+CWT}$	0.18	0.19	0.40
$\mu_{C(\lambda)+CW} < \mu_{\lambda=0.7+CWT}$	0.18	0.19	0.34
$\mu_{C(\lambda)+CW} < \mu_{\lambda=0.5+CWT}$	0.18	0.19	0.32

[**Verifying H_4**] We verified the subhypotheses H_{41} and H_{42} via two-tailed tests at the alpha level of 0.05. For both of them, there are no significant differences between the means of the correlation coefficient (R) obtained from the emails and the game, respectively (H_{41} : p-value=0.32, H_{42} : p-value=0.19). Therefore, subjects' trust assessment behaviors in emails and in games are not different, thereby rejecting Hypothesis H_4 .

8 Discussion

Our main contribution is a computational approach for trust that overlays a domain-independent concept describing the social relationships and outcomes of interpersonal interactions. Previous theoretical approaches, both cognitive [4] and architectural [26, 5], have considered rich concepts but they are not easy to be used as bases for computing trust in the field. By contrast, previous computational approaches have largely worked in an ad hoc manner that bind the trust reasoning to a particular domain.

For the email dataset, comparing the means of the MAEs, our approach yields a correlation between subjects' intuitions regarding trust values and those computationally predicted values. Discounting windows customized for each subject yield improved predictions. Considering commitment weights improves predictions further, though not significantly. We additionally evaluate if subjects trust assessment behavior varies across two decision contexts, namely, as bystanders (reading emails

exchanged between agents) and as game players (playing the Colored Trails game with opponents that involves competitive as well as cooperative elements). From the t-test results, we find that both hypotheses (H_{41} and H_{42}) are rejected thereby suggesting that subjects' trust assessment behaviors in emails and in games do not vary.

The limitation of our results may be due to the following reasons: (1) lack of adequate data and (2) a greater fraction of experiences being judged neutral than positive or negative. Also, we lack an existing approach to compare our results. However, we submit that our contribution is valuable for having launched a new research direction on computational techniques unifying trust and commitments. Publishing imperfect results, as in this submission, might serve as an antidote to a systematic bias in academic research to favor "success stories" over accurate reporting of empirical results, a bias that is increasingly decried in the scientific disciplines, e.g., [16, 22].

9 Future Directions

First, our dataset is not large. A challenge we faced was motivating subjects to provide trust values truthfully for a larger dataset.

Second, our work is limited to predicting trust updates and ignores certainty. According to Wang *et al.* [31] certainty is the measure of confidence that a truster places in a trustee based on its experiences with the trustee. A truster's certainty increases with increasing number of consistent experiences of the truster with the trustee. Thus, certainty is crucial to trust. However, it is difficult to elicit certainty from subjects since certainty may be more subjective than even trust. A more careful social science style qualitative investigation, as suggested by a reviewer, may be appropriate. In the future, we plan to address these limitations by adopting an incentive scheme that motivates subjects to provide trust values truthfully.

Third, we plan to extend our model to Hidden Markov Models (HMM) [30, 18] that incorporate the temporal aspect of trust, i.e., an agent's current trust is computed based on its past trust. For the same, we will motivate subjects to provide intermittent trust labels by reading emails.

Fourth, there is no reason to be limited to commitments: indeed, we have begun work on bringing in cognitive aspects such as goals and emotions, suitably elicited from subjects, as a basis for creating commitments and judging commitment outcomes and overall trust.

Fifth, a subtle potential benefit of our approach is that it seeks to understand communications and can thus provide more natural explanations for trust estimates than an approach that is purely heuristic. Evaluating this potential benefit would require additional human study, which we defer to future work.

Competing interests

The authors declare that they have no competing interests.

Author's contributions

Kalia developed the model, methods, and experiments. Zhang evaluated the data collected from experiments. Singh provided the overall direction and helped realizing the model and methods.

Acknowledgment

We thank the anonymous reviewers for their helpful comments. This research is partially supported by a US Army Research Laboratory (ARL) Oak Ridge Institute for Science and Education (ORISE) Fellowship and the US Department of Defense through the Science of Security Label.

Author details

¹Department of Computer Science, North Carolina State University, , 27695 Raleigh, USA. ²IBM Corporation, Research Triangle Park, , 27709 Durham, USA.

References

1. Sibel Adali, Fred Sisenda, and Malik Magdon-Ismael. Actions speak as loud as words: Predicting relationships from social behavior data. In *Proceedings of the 21st International Conference on World Wide Web*, pages 689–698, Lyon, 2012. ACM.
2. Aljoscha Burchardt, Katrin Erk, and Anette Frank. A WordNet detour to FrameNet. In *Proceedings of the GLDV 2005 Workshop GermaNet II*, pages 1–15, Bonn, 2005.
3. Chris Burnett and Nir Oren. Sub-delegation and trust. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, pages 1359–1360, Valencia, 2012. IFAAMAS.
4. Cristiano Castelfranchi and Rino Falcone. *Trust Theory: A Socio-Cognitive and Computational Model*. Agent Technology. John Wiley & Sons, Chichester, United Kingdom, 2010.
5. Amit K. Chopra, Elda Paja, and Paolo Giorgini. Sociotechnical trust: An architectural approach. In *Proceedings of the 30th International Conference on Conceptual Modeling (ER)*, volume 6998 of *Lecture Notes in Computer Science*, pages 104–117, Brussels, 2011. Springer.
6. Thomas F. Colman and Yuying Li. An interior trust region approach for nonlinear minimization subject to bounds. *SIAM Journal on Optimization*, 6(2):418–445, 1996.
7. Marco Comerio. Value-based service contract selection. In *Proceedings of the 2013 IEEE International Conference on Services Computing*, pages 611–618, Santa Clara, 2013.
8. Thomas DuBois, Jennifer Golbeck, and Aravind Srinivasan. Predicting trust and distrust in social networks. In *Proceedings of 3rd International Conference on Social Computing*, pages 418–424, Boston, oct 2011.
9. Andrew Fiore and Jeff Heer. UC Berkeley Enron email analysis. 2004.
10. Ya'akov Gal, Barbara Grosz, Sarit Kraus, Avi Pfeffer, and Stuart Shieber. Agent decision-making in open-mixed networks. *Artificial Intelligence*, 174(18):1460–1480, 2010.
11. Audun Jøsang. A subjective metric of authentication. In *Proceedings of the 5th European Symposium on Research in Computer Security (ESORICS)*, volume 1485 of *Lecture Notes in Computer Science*, pages 329–344, Louvain-la-Neuve, Belgium, September 1998. Springer.
12. Audun Jøsang, Ross Hayward, and Simon Pope. Trust network analysis with subjective logic. In *Proceedings of the 29th Australasian Computer Science Conference, ACSC*, pages 85–94, Hobart, Australia, 2006. Australian Computer Society, Inc.
13. Anup K. Kalia, Hamid R. Motahari Nezhad, Claudio Bartolini, and Munindar P. Singh. Monitoring commitments in people-driven service engagements. In *Proceedings of the 10th IEEE International Conference on Services Computing*, pages 160–167, Santa Clara, 2013. IEEE.
14. Anup K. Kalia and Munindar P. Singh. Muon: Designing multiagent communication protocols from interaction scenarios. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 29(4):621–657, July 2015.
15. Georgia Kastidou, Kate Larson, and Robin Cohen. A trust model for truthful disclosure of anticipated contributions in multiagent systems. In *Proceedings of the 3rd Workshop on Incentives and Trust in E-Communities*, pages 19–24, Québec City, 2014.
16. Michal Kicinski. Publication bias in recent meta-analyses. *PLoS ONE*, 8(11):e81823, November 2013.
17. Bryan Klimt and Yiming Yang. The Enron corpus: A new dataset for email classification research. In *Proceedings of the 15th European Conference on Machine Learning*, volume 3201 of *LNCS*, pages 217–226, Pisa, 2004. Springer.
18. Xin Liu and Anwitaman Datta. Modeling context aware dynamic trust using Hidden Markov Model. In *Proceedings of the 26th National Conference on Artificial Intelligence*, pages 1938–1944, Toronto, 2012.
19. Roger C. Mayer, James H. Davis, and F. David Schoorman. An integrative model of organizational trust. *The Academy of Management Review*, 20(3):709–734, 1995.
20. Mark Nartey and Ferguson Ebo Yankson. A semantic investigation into the use of modal auxiliary verbs in the manifesto of a Ghanaian political party. *International Journal of Humanities and Social Science*, 4(3):21–304, February 2014.
21. Nardine Osman, Carles Sierra, Fiona Mcneill, Juan Pane, and John Debenham. Trust and matching algorithms for selecting suitable agents. *ACM Transactions on Intelligent Systems and Technology*, 5(1):16:1–16:39, January 2014.
22. Jeffrey D. Scargle. Publication bias: The “file-drawer” problem in scientific inference. *Journal of Scientific Exploration*, 14(1):91–106, 2000. http://www.scientificexploration.org/journal/jse_14.1_scargle.pdf.
23. Lauren E. Scissors, Alastair J. Gill, Kathleen Geraghty, and Darren Gergle. In CMC we trust: The role of similarity. In *Proceedings of the 27th International Conference on Human Factors in Computing System*, pages 527–536, Boston, 2009. ACM.
24. Munindar P. Singh. An ontology for commitments in multiagent systems: Toward a unification of normative concepts. *Artificial Intelligence and Law*, 7(1):97–113, March 1999.
25. Munindar P. Singh. Semantical considerations on dialectical and practical commitments. In *Proceedings of the 23rd Conference on Artificial Intelligence (AAAI)*, pages 176–181, Chicago, July 2008. AAAI Press.
26. Munindar P. Singh. Trust as dependence: A logical approach. In *Proceedings of the 10th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 863–870, Taipei, May 2011. IFAAMAS.
27. Munindar P. Singh, Amit K. Chopra, and Nirmal Desai. Commitment-based service-oriented architecture. *IEEE Computer*, 42(11):72–79, November 2009.
28. W.T. Luke Teacy, Jigar Patel, Nicholas R. Jennings, and Michael Luck. TRAVOS: Trust and reputation in the context of inaccurate information sources. *Journal of Autonomous Agents and Multi-Agent Systems*, 12(2):183–198, 2006.
29. Pankaj R. Telang and Munindar P. Singh. Specifying and verifying cross-organizational business models: An

- agent-oriented approach. *IEEE Transactions on Services Computing*, 5(3):305–318, July 2012. Appendix pages 1–5.
30. George Vogiatzis, Ian MacGillivray, and Maria Chli. A probabilistic model for trust and reputation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, pages 225–232, Toronto, 2010. IFAAMAS.
 31. Yonghong Wang, Chung-Wei Hang, and Munindar P. Singh. A probabilistic approach for maintaining trust based on evidence. *Journal of Artificial Intelligence Research*, 40:221–267, January 2011.
 32. Yonghong Wang and Munindar P. Singh. Evidence-based trust: A mathematical model geared for multiagent systems. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 5(4):14:1–14:28, November 2010.
 33. Jie Zhang and Robin Cohen. A framework for trust modeling in multiagent electronic marketplaces with buying advisors to consider varying seller behavior and the limiting of seller bids. *ACM Transactions on Intelligent Systems and Technology*, 4(2):24:1–24:22, April 2013.

10 Appendix

10.1 Data

We provide the data content in the following URL: "http://tinyurl.com/q59joom". In the data folder there are two sub folders: (1) Email.Documents and (2) Processed.Data. In the *Email.Documents* folder, there are 14 files. Each file represents email exchanges between Kimberly and one of her colleague at Enron. In the *Processed.Data* folder, there are 15 files. Each file represents data in the format given in Table 3. Each files contains experiences observed by the truster from the trustee and corresponding trust values of the truster toward the trustee as assigned by subjects. There are four kinds of processed data files.

- *Processed.Data.Emails.Experiences.Trust.Values* file contains experiences obtained from emails between senders and receivers and corresponding trust value assigned by subjects.
- *Processed.Data.Game.Experiences.Trust.Values* file contains experiences obtained from chats exchanged between subjects during their game play and corresponding trust value assigned by subjects
- *Processed.Data.Emails.Windows.Size.1-12.Trust.Values* files contain experiences based on different window sizes (from 1 to 12) and corresponding trust values assigned by subjects. The maximum window size considered is 12 since we mention in Section 7.1 that increasing the window size beyond 12 does not improved the result (reduce the MAE).
- *Processed.Data.Emails.Experiences.Strength.Trust.Values* file contains experiences strength computed from emails based on ranks of features provided in Table 6 and corresponding trust values assigned by subjects.

10.2 Questionnaires

We asked the following questions to subjects to collect their perceived importance about features discussed in Section 5.3.

Commissive over directive. To assess whether a commissive carry a greater weight than a directive, we asked subjects the following question (choice 1 indicates a directive whereas choice 2 indicates a commissive).

- According to you which is more important?
 - 1. Please review the attached agreement for the Big Sandy Interconnect.
 - 2. I will review the attached agreement for the Big Sandy Interconnect.
 - 3. 1 and 2 are equal

Debtor's type. To assess if a single debtor type carry a greater weight than multiple debtors, we asked subjects the following questions.

- According to you which is more important?
 - 1. We will follow up with Mike to make sure he understands how the numbers were derived.
 - 2. I will follow up with Mike to make sure he understands how the numbers were derived.
 - 3. 1 and 2 are equal
- According to you which is more important?
 - 1. SENDER: Kimberly; RECEIVER: Steven, Teb, Mark, Mansoor, Earl, Stephen, Robert, Jan, Mark, Mansoor, Earl, Stephen, Robert, Jan; Please review the attached work order for the tap and side valve for the new Agave interconnect.
 - 2. SENDER: Kimberly; RECEIVER: Steven; Please review the attached work order for the tap and side valve for the new Agave interconnect.
 - 3. 1 and 2 are equal

Creditor's type To assess whether multiple debtors carry a greater weight than a single creditor, we asked subjects the following question (Choice 1

- According to you which is more important?
 - 1. SENDER: Kimberly; RECEIVER: Lorraine, Lohman, Michelle, Mark, Paul; Would you please send me your bullets by the end of today (before we leave for the Cirque show).
 - 2. SENDER: Kimberly; RECEIVER: Lorraine; Would you please send me your bullets by the end of today (before we leave for the Cirque show).
 - 3. 1 and 2 are equal
- According to you which is more important?
 - 1. SENDER: Michelle; RECEIVER: Rich Cc: Earl; Kimberly; I will speak to Mark Kraus (EOG Commercial) to recap, in case he was not aware of the results.
 - 2. SENDER: Michelle; RECEIVER: Rich, Earl, Kimberly; I will speak to Mark Kraus (EOG Commercial) to recap, in case he was not aware of the results.
 - 3. SENDER: Michelle, RECEIVER: Rich; I will speak to Mark Kraus (EOG Commercial) to recap, in case he was not aware of the results
 - 3. 1 and 2 are equal
 - 4. 2 and 3 are equal
 - 5. 1 and 3 are equal

Modal verbs. To assess which modal verb convey high confidence over others, we asked subjects the following question.

- 1 Rank the following sentences in a order that ranges over a scale of 1 to 8 where 1 indicates the lowest and 8
 - 1. I would call you on Monday to discuss so we can give it to Danny quickly [Assign Rank: 1–8].
 - 2. I must call you on Monday to discuss so we can give it to Danny quickly [Assign Rank: 1–8].
 - 3. I can call you on Monday to discuss so we can give it to Danny quickly [Assign Rank: 1–8].
 - 4. I will call you on Monday to discuss so we can give it to Danny quickly [Assign Rank: 1–8].
 - 5. I could call you on Monday to discuss so we can give it to Danny quickly [Assign Rank: 1–8].
 - 6. I should call you on Monday to discuss so we can give it to Danny quickly [Assign Rank: 1–8].
 - 7. I may call you on Monday to discuss so we can give it to Danny quickly [Assign Rank: 1–8].
 - 8. I shall call you on Monday to discuss so we can give it to Danny quickly [Assign Rank: 1–8].

Deadlines. To assess if noun phrases with deadlines may convey more importance than noun phrases without deadlines we asked subjects the following questions

- According to you which is more important?
 - 1. Would you please send me your bullets by the end of today?
 - 2. Would you please send me your bullets?
 - 3. 1 and 2 are equal