Trustworthy Decision Making via Commitments*

Anup K. Kalia, Zhe Zhang, and Munindar P. Singh

North Carolina State University
{akkalia, zzhang13}@ncsu.edu, m.singh@ieee.org

Abstract. Existing approaches to calculate trust between agents rely on the strength of their relationships based on their degrees of friendship, the frequencies of their interactions, the sentiments extracted from their interactions, and so on. These approaches rely heavily on numerical measures and disregard the deep structure underlying trust. By contrast, we establish the idea of trust among agents based on their normative relationships (norms in short). Norms provide the standard of correctness of the interactions between agents and provide a basis for accountability for their actions. We focus on commitments to capture relationships between agents. Based on commitments, we provide a framework using which we can calculate trust probabilistically between agents. We evaluate our approach in two parts. First, we specified interactions between agents and asked subjects to estimate the level of trust between each pair of agents. Second, we selected over 5,000 email sentences from the well-known Enron dataset and asked subjects to estimate the level of trust between each pair of participants. We learned subjects' trust parameters from their estimation. We recalculated the trust values between participants in datasets using these parameters. We found that the mean absolute error (MAE) between trust values calculated from learned parameters and estimates provided by subjects is minimum when compared against trust values calculated using any fixed parameter.

1 Introduction

Trust is important in sociotechnical systems where agents enact different roles and interact with each other. An important research problem is *how to build trustworthy sociotechnical systems*. Existing approaches to build such systems rely on hardcoded policies, adopt a centralized perspective, and do not provide standards to measure the correctness of behavior in such systems. Also, existing methods for calculating trust are based on numerical measures such as number of transitions or connections, relationship strength, and so on [1, 11, 17]. Some approaches provide complex formalizations to capture trust, but are difficult extracting the formal models or patterns from real-world datasets for prediction and analysis [4, 10]. By contrast, we emphasize the normative relationships captured between agents from their interactions with one another in a sociotechnical system. We leverage these normative relationships and determine trust values between agents based on the progression of these relationships. Singh [16]

* This research was partially supported by the US Army Research Office under the Science of Security Lablet grant.

defines five types of norms: commitments, prohibition, authorization, power, and sanction which characterize a sociotechincal system. Here, we consider we consider only commitments.

Singh [14] introduces a commitment as a *conditional* business relationship directed from a debtor (an agent) to a creditor (an agent). A commitment is formalized as C(*debtor, creditor, antecedent, consequent*). Here, the debtor commits to bringing about the consequent for the creditor provided the antecedent holds. When a debtor offers a commitment to a creditor, the commitment is *created* and becomes active. When the antecedent is brought about, the commitment is *detached* or becomes unconditional. When the consequent holds, the commitment is *satisfied*. For a detached commitment if the consequent does not hold before the consequent timeout the commitment is *violated*. Consider C(*buyer, seller, goods, pay*), where the buyer commits to paying a specified amount provided the seller delivers the goods. When the seller delivers, the commitment is detached. When the buyer pays, the commitment is discharged or satisfied. If the seller delivers but the buyer does not pay, the commitment is violated.

The above example provides a basis for correlating commitments with trust. The trust between the buyer and the seller is affected when the commitment transits from one state to another. If the seller detaches the commitment, the trust of the buyer for the seller may increase but not necessarily as the seller brings a positive outcome by delivering the goods. When the commitment discharges, the trust of the seller for the buyer increases as the buyer meets the expectation of the seller by paying for the delivered goods. If the buyer somehow violates the commitment then the trust of the seller for the buyer decreases as the violation of the commitment leads to a negative outcome. We map transactions between agents to different commitment operations: create, detach, discharge, violate, cancel, release, assign, and delegate so as to calculate trust between them. We discuss these operations in detail in Section 2.

We emphasize commitments to calculate trust because commitments capture important business relationships between agents and provide a logical basis to verify whether their expectations are met or not. The success or failure of a task directly affects the trust values of agents. For example, if a customer requests a quote from a merchant without any prior commitment from the merchant to provide it then decreasing the trust value placed by the customer for the merchant will be incorrect if the merchant does not provide a quote. However, if there had been a prior commitment from the merchant to the customer that whenever the customer asks for a quote the merchant will provide it then the merchant's not providing the quote would affect the trust of the customer for the merchant.

Existing approaches [5, 15] for finding trust on the basis of commitments are based on logic. We focus on calculating trust as a probabilistic measurement based on collected evidence. Here, the evidence depends on how the normative relationships between agents progress. We focus on a probabilistic approach so that we can provide a practical framework for agents to interact based on their trust for other agents. To calculate trust, we use Wang and Singh's trust model [19]. In this model, trust is represented as binary evidence $\langle r, s\rangle$ where $r \geq 0$ and $s \geq 0$. Here, r ≥ 0 represents the number of positive experiences a truster has with a trustee and $s \geq 0$ represents number of negative experiences a truster has with a trustee. Both r and s are real numbers. Wang and Singh calculate the expected value of the probability of a positive outcome as $\alpha = \frac{r}{r+s}$, defining α = 0.5 when r+s=0. Consider an example where a buyer makes 10 transactions with a seller. If out of the 10 transactions 8 succeed and 2 fail, then the trust value of the seller for the buyer is 0.8 However, if out of the 10 transactions seven fail and three succeed, then the trust value of the seller for the buyer is 0.3. If five succeed and five fail then the trust value of the seller for the buyer remains unaffected and remains 0.5. Wang and Singh motivate certainty as a function of r and s. The certainty decreases with an increase in the number of conflicts provided the number of transactions is fixed.

The main contribution of this paper is to provide a framework to derive trust values between agents when they interact to create and progress their commitments toward each other. Our framework shows how agents can enable agents to learn their trust values for other agents and guide them to make trustworthy decisions. We use $\mathsf{C}_{x,y}$ to denote a commitment from x (debtor) to y (creditor) and $\mathsf{T}_{a,b}$ to denote the trust of a (truster) for b (trustee). The idea behind using these notations is to highlight the directionality of commitments and trust between the agents.

This paper is organized as follows: Section 2 provides the model for deriving trust values between agents. Section 3 provides evaluation and results. Section 4 provides conclusion and Section 5 discusses related work.

2 Commitment and Trust

We use commitment operations as a basis for agents to update their trust values for each other. Figure 1 represents these operations as part of the commitment lifecycle. A commitment is *created* when a debtor voluntarily offers to do a task or when he is directed to do a task; *detached* if a condition or an antecedent present for a commitment holds true; *discharged* when a debtor executes a committed task. Additionally, a commitment can be delegated, assigned, canceled, and released. 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. The commitment can be *canceled* by the debtor and *released* by the creditor.

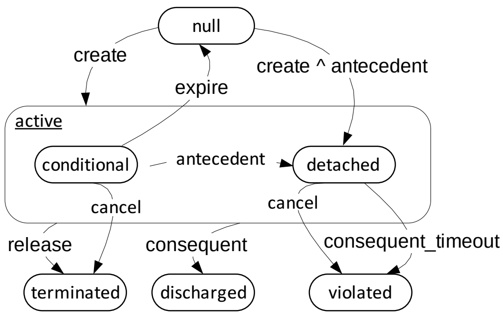

Fig. 1. The life cycle of a commitment [18].

We update the trust $\langle r, s \rangle$ of a trustor in a trustee for an interaction as given in the following equations.

$$r_{new} \leftarrow i_r + (1-\beta)r_{old} \tag{1}$$

$$s_{new} \leftarrow i_s + (1-\beta)s_{old} \tag{2}$$

In the above equations, r_{old} and s_{old} represent the initial trust between agents; β represents a temporal discount factor; i_r and i_s represent the trust update values where r_{new} and s_{new} represent the resulting evidence values when the trust is updated. When a positive event happens, trust is updated by i_r. When a negative event happens, trust is updated by i_s. Once we determine r_{old} and s_{old}, we find α, a function of r_{old} and s_{old}, where $\alpha(\mathrm{r}_{old},\mathrm{s}_{old}) = \frac{r_{old}}{r_{old}+s_{old}}$ and certainty c(r_{old}, s_{old}). Finally the trust value we get is αc(r_{old}, s_{old}). In the following paragraphs, we explain how trust values are updated among agents based on their commitments. For convenience, we represent the trust from a trustor A towards trustee B as $\mathsf{T}_{A,B} = \langle r_{A,B}, s_{A,B} \rangle$.

Create, detach, discharge, cancel, and release Consider Figure 2(a), which represents the states for commitment C_1 = C(Buyer, Seller, deliver, pay). When Buyer (B) commits Seller (S) to paying if S delivers, C_1 goes from the null state (N) to the conditional state (C). At this moment, there is no change to the trust values of B for S ($\mathsf{T}_{B,S}$) and S for B ($\mathsf{T}_{S,B}$). When S delivers, C_1 goes to the detached state (D) and $\mathsf{T}_{B,S}$ increases. If S does not deliver then C_1 expires and $\mathsf{T}_{B,S}$ decreases. When B pays, C_1 is discharged (S) and $\mathsf{T}_{S,B}$ increases. If B does not pay or cancels C_1, C_1 is violated and $\mathsf{T}_{S,B}$ decreases. However, if B cancels C_1 without the C_1 being detached, then C_1 is terminated and $\mathsf{T}_{S,B}$ remains the same. Now, suppose S releases the commitment. Then $\mathsf{T}_{B,S}$ is not affected.

Delegation and discharge Figure 2(b) shows how the trust values between B, S, and Buyer' (B') are affected when B delegates commitment C_1 to B'. When B delegates C_1 in the conditional state to B' and S detaches it, then trust values of B and B' for S increase. Now, if B' discharges the commitment, then the trust values of S for both B and B' increase as well as the trust value of B for B' increases. When B delegates C_1 in the detached state to B' the trust values between B, B', and S remain unaffected until B' discharges it.

Delegation and cancel Figure 2(c) shows how trust values are updated when B delegates C_1 to B' and B' cancels it. When B delegates C_1 to B' without detaching it and S detaches it, then the trust values of B and B' for S increase. Now, when B' cancels the detached commitment, the trust values of S for B and B' decrease as well as the trust values of B for B' decreases. If B cancels the conditional commitment, then the trust values between B, B', and S remains the same. When B' violates the commitment by canceling it in the detached state, B remains committed to S.

Delegation and release If B delegates C_1 to B' and S releases the commitment either before or after detaching the commitment, the trust values between B, B', and S remains unaffected.

Delegation and violation Here, B delegates C_1 to B'. However, B' never discharges the commitment, thereby violating it. Therefore, the trust values of S for B and B' decrease as well as the trust value of B decreases for B'.

Assignment and discharge When S assigns C_1 before detaching it to S' and S' detaches C_1, then the trust values of S and B for S' increase as well as the trust value of

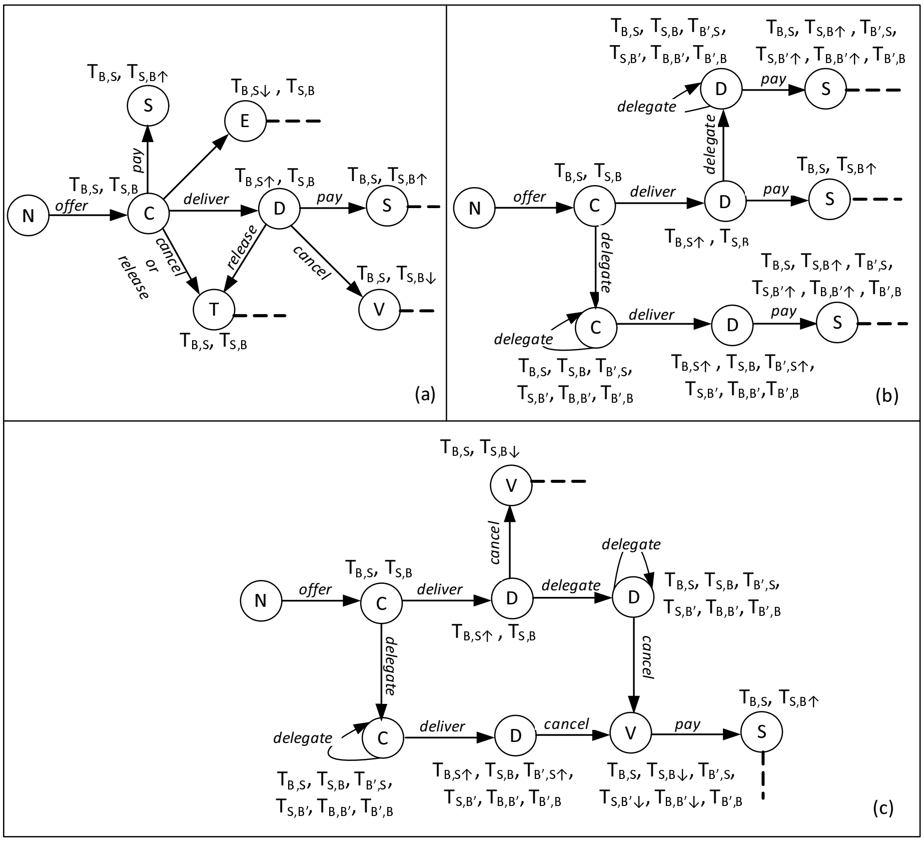

Fig. 2. Different paths of a commitment progression and trust updates along these paths. $\mathsf{T}_{x,y}\uparrow$ and $\mathsf{T}_{x,y}\downarrow$ respectively indicates that the trust from x to y increases and decreases.

B for S. If S' does not detach C_1, the trust values of S and B for S' decrease. When B discharges C_1, the trust values of S and S' for B increase.

Assignment and cancel When S assigns the detached commitment C_1 to S' and B cancels it, then the trust values of S and S' for B decrease. If S assigns C_1 without detaching it to S' and B cancels it even before S' detaches it then the trust values among S, S' and B remain unaffected. However, if S' detaches it and then B cancels it, then the trust values of S and S' for B decrease.

Assignment and release When S assigns C_1 to S' and S' releases it, then the trust values between S, S', and B remain unaffected.

Assignment and violation When S assigns C_1 to S' and B violates it by not executing the consequent or causing consequent_timeout to occur, the trust values of S and S' for B decrease.

3 Evaluation

We evaluated our approach using the process flow as shown in Figure 3.

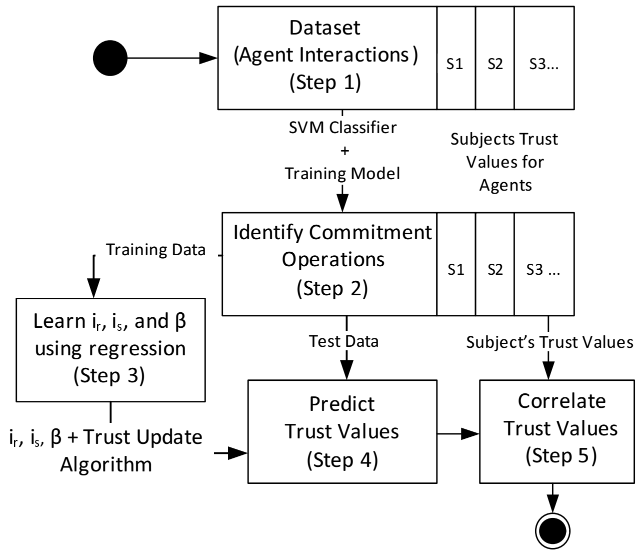

Fig. 3. Process for evaluation.

(Step 1) We build a dataset consisting of agent interactions and trust values for agents assigned by subjects (S1, S2, S3, . . .) based on the agent interactions.

(Step 2) We identify commitment operations from each of the interactions using our support vector machine (SVM) classifier and a training model [8].

(Step 3) We divide the dataset into the training and test datasets. We collate each subject's assigned trust values with the identified commitment operations to learn trust update values i_r and i_s and discount factor β for each subject individually. We do so by performing *regression analysis* [] on the training data using commitment operations and trust values provided by each subject.

(Step 4) Once we learn i_r, i_s, and β for the subject, we predict trust values as estimated by the subjects for agents in the test data using commitment operations and Equations 1 and 2.

(Step 5) We correlate predicted trust values from the equations and the subject's labeled trust values from the test data.

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

3.1 Learning Subjects' Trust Parameters

Trust is subjective and its assessment varies across people. For example, some subjects may reward or penalize agents more than others when a commitment is respectively discharged or canceled. Therefore, we learn trust parameters (i_r, i_s, r_{new}, r_{old}, β) for each subject from the trust values assigned by that subject. For our evaluation, we ignore β to keep the experiment simple and easy to analyze. Thus, Equations 1 and 2 reduce to $r_{old} \leftarrow i_r$ + r_{new} and $s_{old} \leftarrow i_s$ + s_{new} respectively. From these two equations, we derive Equation 3 where $\hat{\alpha_k}$ represents trust values predicted for the kth directed pair of interacting parties using the learned parameters and POS$_k$, NEG$_k$, and NEU$_k$ represent number of positive, negative, and neutral interactions, respectively.

$$\hat{\alpha_k} = \frac{r_{new} + \text{POS}_k i_r + 0.5 \cdot \text{NEU}_k i_r}{r_{new} + s_{new} + \text{POS}_k i_r + \text{NEG}_k i_s + 0.5 \cdot (\text{NEU}_k i_r + \text{NEU}_k i_s)} \quad (3)$$

Now to learn the parameters for each subject, we define our objective function as

$$\Omega = \sum_{k=1}^{n} |\hat{\alpha_k} - \alpha_k| \quad (4)$$

where α_k represents actual trust values assigned by the subject to the left pair of interacting parties. Now the optimization problem is

$$E = \operatorname{argmin} \Omega \quad (5)$$

where we try to achieve a minimum mean absolute error (MAE) of $\hat{\alpha_k}$ with respect to α_k. By transforming this optimization problem into a nonlinear programming problem, the solution can be found efficiently.

3.2 Evaluation 1: The Self-Created Dataset

We created a sample multiagent interaction inspired from a case study in the insurance domain [2] that consists of seven agents and 90 messages. In the case study, an insurance company AGFIL (AG) provides car insurance to a customer John Doe (JD) provided JD pays the premium. AG provides his support to JD by outsourcing his key functionalities such as handling and resolving JD's claims to Europ Assist (EA, a call center) and Lee Consulting Services (LCS, a claim inspector) respectively. When JD files a complaint, EA handles it by sending JD's car to a mechanic (M) and informs AG about it. AG requests LCS to resolve JD's claim. LCS requests M to provide a quote for JD's car repair. If LCS accepts the quote it asks M to repair the car. Once the car is repaired, LCS asks AG to pay M for the estimation and the repair. AG pays M for the repair. Also, AG pays EA and LCS for handling and resolving JD's claim, respectively. LCS asks JD to pick his car from M. In addition to these interactions that represent happy paths, we add some exceptions such as LCS rejects M's estimate and ask M2 to estimate the car repair. Similarly, AG hires LCS2 when LCS denies to resolve JD's car issue. In the following sections, we provide our hypotheses and statistical results based on the correlation of each subjects trust values and the trust values predicted using Equations 2 and 3. We gave this interaction to nine subjects (graduate students in computer science) and asked them to intuitively label trust values between these seven agents after reading all the messages based on their individual intuitions about trust. We provide a few of the messages in Table 1. After reading the messages between two agents LCS and M, the subjects are asked to come up with values from 0 to 1 for each $\mathsf{T}_{LCS,M}$ and $\mathsf{T}_{M,LCS}$.

Table 1. Sample messages from the AGFIL case study.

S	R	Content
LCS	M	Can you please provide the estimate for the repairs? If I accept it, I will pay for it.
M	LCS	Sure.
M	LCS	2500$. Have sent you the details to your email.
LCS	M	Thanks. Will verify and let you know.
LCS	M	The estimation is too high compared to the damages. Can you give me a lower quote?
M	LCS	Can't be lower.
LCS	M	Thanks but I have to reject your quote.

Hypotheses 1 **[H1]** An increase in the number of interactions between any two agents, correlates with the number of commitments created and discharged resulting in an increase of trust values between agents. **[H2]** Trust values calculated using our approach correlate better with trust values estimated by subjects than trust values calculated using selected fixed approaches.

Results from Evaluation 1 **[Verifying H1]** To verify this hypothesis, we correlated the number of interactions between any two agent and the sum of the number of commitments created and discharged between them. We obtained a medium correlation of 0.47. **[Verifying H2]** We collected trust values from nine subjects and learned their individual trust parameters using the objective function in Equation 4. To evaluate our approach we compared the trust values predicted using learned parameters against trust values calculated using fixed parameters for i_s, i_r, r_{new}, r_{old} as 10, 10, 1, 1 (Fixed 1), 12, 8, 1, 1 (Fixed 2), and 8, 12, 1, 1 (Fixed 3). We calculated the mean absolute error between subject-assigned trust values and trust values calculated using learned parameters and fixed parameters. Figure 4(a) shows our results. The X-axis represents mean absolute error and the Y-axis lists the different approaches. In the figure, we can see that the median (0.2644) of the mean absolute error values using the learned parameter approach is lower than the median of the other approaches (0.3494, 0.3356, 0.3728). This suggests that our approach of calculating trust values using learned parameters strongly correlates better with subject's assigned trust values than the selected fixed approaches.

3.3 Evaluation 2: The Enron Dataset

We evaluate our approach using the Enron email corpus [6, 9]. From this corpus, we collected 5,487 email sentences that were exchanged between Kimberly Watson, an employee of Enron, and more than 30 people, including her coworkers at Enron, clients, friends, and family members. In this evaluation, we correlate agent's sentiments with commitment operations and trust values. Therefore, we requested subjects to identify

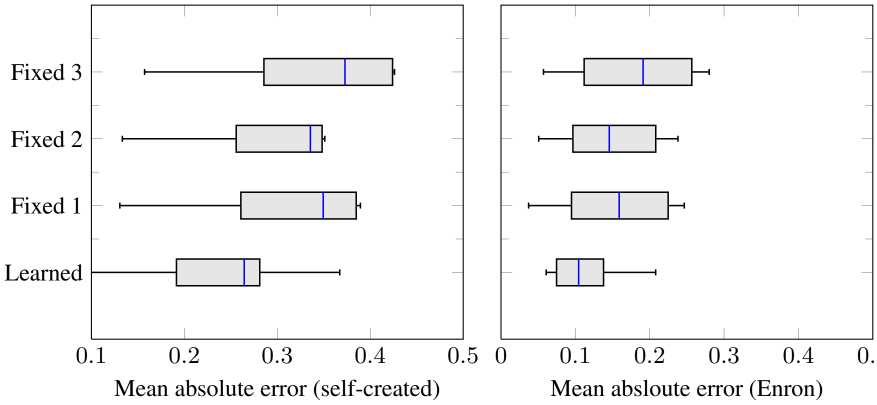

Fig. 4. Results from the self-created and Enron datasets comparing different approaches to predict trust values.

the sentiments expressed in each email (*positive*, *negative*, or *neutral*) along with trust values for agents after all emails. We consider sentiments for our evaluation as they reflect the outcome of agent actions. Suppose a debtor discharges its commitment for a creditor. Then, the creditor may display its positive sentiment by praising the debtor for the successful outcome. Therefore, in the similar direction, trust values may correlate with sentiments. In the above example, the creditor develops positive trust as well as positive sentiment for the debtor. In the following sections, we provide our hypotheses and statistical results.

Hypotheses 2 **[H3]** The number of commitments created and discharged increases with an increase in the number of interactions between Kimberly and her coworkers thereby increasing trust between them. **[H4]** Trust values calculated using our approach strongly correlate with trust values contributed by subjects. **[H5]** There is a strong correlation between positive sentiments and trust labeled by subjects for each agent.

3.4 Results from Evaluation 2

[Verifying H3] We correlated the number of commitments created and discharged between Kimberly and her coworkers with the number of interactions between them to find a high correlation of 0.86. **[Verifying H4]** Similar to our previous evaluation, this time, we asked six external subjects to assign trust values of Kimberly for her coworkers as well as her coworkers for Kimberly. Based on the approach, in Figure 3, we extracted commitments operations between Kimberly and her coworkers. Then we tried to learned trust parameters for each subject based on his or her assigned trust values. Figure 4(b) shows our results. In the figure, we can see that the median of the mean absolute error using the learned parameter approach (0.1044) is lower than for the fixed approaches (0.1589, 0.1456, 0.1911). **[Verifying H5]** To verify the hypothesis we asked the above six external subjects to label sentiments (positive, negative, neutral) for each email and assign trust values to agent pairs ($\mathsf{T}_{debtor,creditor}$). Then, we correlated positive sentiments labeled by subjects with trust values assigned by subjects. Overall, we obtained a correlation (0.1078, 0.5853, 0.2644, 0.6933, 0.3124, 0.5498) between trust values and positive sentiments as labeled by subjects. The result is a medium correlation value because in Enron most of the emails convey neutral sentiments.

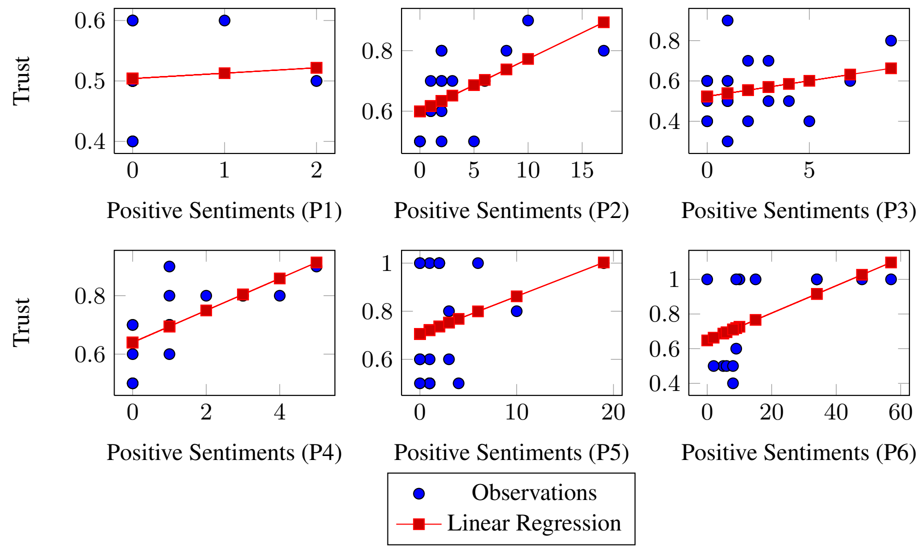

Fig. 5. Correlation between trust and positive emotion.

4 Conclusion

Existing works to calculate trust between agents are based on numerical measures such as attributes, behavior, and relationships. In contrast, we provide a novel approach in which the progression of normative relationships between participants provides means to measure trust values. Using our approach we show strong correlation between the human intuition of trust values and predicted trust values for different datasets. Our approach yields a medium correlation between positive sentiments and trust values labeled by subjects.

Currently, our work is limited to predicting trust updates without using a discount factor and certainty. In the future we plan to address these limitations by learning both these parameters and correlate the trust values calculated using these parameters with the actual trust values assigned by subjects. In the next section, we discuss some relevant existing works.

5 Discussion

Burnett and Oren [3] examine the effects of delegation using a probabilistic trust model [7] and propose a model of weighting trust updates based on shared responsibility. They evaluate their model using different weighting strategies such as All-First/Last Weighting, Increasing/Decreasing Weighting, Full Weighting, and so on. In Burnett and Oren's work, delegation refers to passing the responsibility of performing a task from one agent to another whereas in our work delegation refers to delegating a commitment from one debtor to another. We restrict our trust update to delegation chains of length three agents (the debtor, the new debtor and the creditor). This means if the new debtor delegates the commitment to another debtor (debtor'), then trust between the debtor and the new debtor remains unaffected. In Burnett and Oren's work there is no restriction to the delegation chain length. However, finding longer chains in a text corpus is rare.

Adalı et al. [1] present behavioral features that capture relationship between people without emphasizing the textual content exchanged between them. The features they present are reciprocity, assortativity, attention, and latency. Adalı et al. also introduce a methodology for determining such features. They evaluate their approach using data from Twitter. Unlike Adalı et al. approach, we emphasize commitment created and discharged between people based on the text content exchanged between them. We leverage commitments to find trust values among these people.

Scissors et al. [13] performed an empirical evaluation with 62 pairs of male and female students and found that content (positive emotion words, task-related words), structure (verb tense, phrasal entrainment), and stylistic (emoticons) reflect high trusting pairs while content such as negative emotions reflect low trusting pairs. Similar to their work, we correlate commitments and trust with positive sentiments. Our results show that a correlation exists between positive sentiments and trust values labeled by subjects.

References

1. S. Adalı, F. Sisenda, and M. Magdon-Ismail. Actions speak as loud as words: Predicting relationships from social behavior data. In *Proceedings of the 21st International Conference on World Wide Web*, WWW, pages 689–698. ACM, 2012.
2. S. Browne and M. Kellett. Insurance (motor damage claims) scenario. Document D1.a, CrossFlow Consortium, 1999.
3. C. Burnett and N. Oren. Sub-delegation and trust. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS, pages 1359–1360. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
4. C. Castelfranchi and R. Falcone. Principles of trust for MAS: cognitive anatomy, social importance, and quantification. In *Proceedings of the 3rd International Conference on Multiagent Systems (ICMAS)*, pages 72–79. IEEE Computer Society, 1998.
5. A. K. Chopra, E. Paja, and P. Giorgini. Sociotechnical trust: An architectural approach. In *Proceedings of the 30th International Conference on Conceptual Modeling (ER)*, volume 6998 of *LNCS*, pages 104–117, Brussels, 2011. Springer.
6. A. Fiore and J. Heer. UC Berkeley Enron email analysis. 2004.
7. A. Jøsang, R. Hayward, and S. Pope. Trust network analysis with subjective logic. In *Proceedings of the 29th Australasian Computer Science Conference*, ACSC, pages 85–94. Australian Computer Society, Inc., 2006.

8. A. Kalia, H. R. M. Nezhad, C. Bartolini, and M. P. Singh. Identifying business tasks and commitments from email and chat. Technical report, 2013. http://www.hpl.hp.com/techreports/2013/HPL-2013-4.pdf.
9. B. Klimt and Y. Yang. The Enron corpus: A new dataset for email classification research. In *European Conference on Machine Learning*, volume 3201 of *LNCS*, pages 217–226. Springer Berlin / Heidelberg, 2004.
10. C.-J. Liau. Belief, information acquisition, and trust in multi-agent systems–a modal logic formulation. *Artificial Intelligence*, 149(1):31–60, 2003.
11. C.-Y. Lin, L. Wu, Z. Wen, H. Tong, V. Griffiths-Fisher, L. Shi, and D. Lubensky. Social network analysis in enterprise. *Proceedings of the IEEE*, 100(9):2759 –2776, sep 2012.
12. OOI. Generic service request protocol, 2011. https://confluence.oceanobservatories.org/display/syseng/CIAD+COI+OV+Service+Request+Protocol.
13. L. E. Scissors, A. J. Gill, K. Geraghty, and D. Gergle. In CMC we trust: The role of similarity. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI, pages 527–536. ACM, 2009.
14. M. P. Singh. An ontology for commitments in multiagent systems: Toward a unification of normative concepts. *Artificial Intelligence and Law*, 7(1):97–113, Mar. 1999.
15. M. 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.
16. M. P. Singh. Norms as a basis for governing sociotechnical systems. *ACM Transactions on Intelligent Systems and Technology (TIST)*, pages 1–21, 2013. To appear; available at http://www.csc.ncsu.edu/faculty/mpsingh/papers.
17. J. Tang, H. Gao, H. Liu, and A. D. Sarma. eTrust: Understanding trust evolution in an online world. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD, pages 253–261. ACM, 2012.
18. P. R. Telang and M. P. Singh. Specifying and verifying cross-organizational business models: An agent-oriented approach. *IEEE Transactions on Services Computing*, 5(3):305–318, July 2012.
19. Y. Wang and M. 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, Nov. 2010.