WCCI 2012 IEEE World Congress on Computational Intelligence
June, 10-15, 2012 - Brisbane, Australia FUZZ IEEE

Modelling of Robot Attention Demand in Human-Robot Interaction Using Finite Fuzzy State Automata

Jamil Abou Saleh, Fakhreddine Karray, and Michael Morckos
Center of Pattern Analysis and Machine Intelligence, University of Waterloo, Canada
$\{jaabousa, karray, mmorckos\}$ $@uwaterloo.ca$

***Abstract*— Many systems have been implemented towards achieving effective human-machine interaction, but run the risk of being ignored if appropriate performance metrics are not in place. As a result, our goal becomes that of providing a foundation upon which we can assess how well the human and the robot perform as a team. Toward the efficient modelling of such metrics, we attempt to determine the true amount of time that an operator has to dedicate to the robot. Therefore, we define the robot attention demand (RAD) as a function of both direct interaction time (DIT) and indirect interaction time (IIT), where the IIT is a direct consequence of the human trust in automation. We propose a two-level fuzzy temporal model to evaluate the human trust in automation while collaborating with robots to complete some tasks. The model combines the advantages of fuzzy logic and finite state machines to best model this phenomenon, and reduces the system complexity and the size of the knowledge base by grouping perceptions into first- and second-order perceptions. The fuzzy knowledge base is further updated by implementing an application robotic platform where robots and users interact via natural language to complete tasks with varying levels of complexity. User feedback is noted and used to tune the knowledge base where needed.**
Keywords: performance metrics, human-robot interaction, fuzzy logic, finite state automata

I. Introduction

Robots have always been touted as powerful tools that could be used effectively in a number of applications ranging from automation to human-robot interaction. In order for such systems to operate adequately and safely in the real world, they must be able to perceive, and must have abilities of reasoning up to a certain level. Toward this end, performance evaluation metrics are used as important measures. The idea of developing a common toolkit of performance metrics has been discussed by many researchers. Olsen and Goodrich discuss six interrelated performance metrics that can lead the design of human-robot interaction systems [15], [3]. The most instrumental one is the robot attention demand (RAD). RAD is a measure of the fraction of total task time that a user must spend interacting with a robot. RAD is defined as a relationship between neglect tolerance (NT) and interaction effort (IE) as shown in equation 1.

$$RAD = \frac{IE}{IE + NT} \quad (1)$$

Doing so, Olsen suggests to approximate the value of the neglect tolerance (NT) by measuring the time between the human instruction and either a drop of the robot performance below the effectiveness threshold, or the intervention of the human with another instruction. Olsen states, however, that in this scenario, this metric is no more independent from the user, and hence the operator's trust in the robot's autonomous abilities becomes a critical issue.

Trust in automation has been extensively discussed in the literature by many researchers, especially in the fields of human-machine interaction [9], [10], [2], [4]. Lee et al. [6] examined the relationship between trust in automatic controllers and user's self-confidence in manually operating a simulated semi-automatic pasteurization plant. Muir et al. [11] present experimental studies of trust and human intervention in a process control simulation, which demonstrates that operators' subjective evaluation of trust in the machines are based mainly on their perception of the robot's competence. There was a high correlation between operators' trust and the use of automation, where operators chose to rely on manual operation when the trust in automation was low [5].

Oleson et al. [14] identify and describe factors that have significant impact on a human's level of trust in a robotic teammate. Such key antecedents of trust can be categorized as human-related (e.g. expectations about the system), robot-related (e.g. the presence or absence of system errors), and environment-related (e.g. more complex and demanding tasks may have different effects on trust than less demanding ones). Li et al. [7] suggest that a more human-like robot, in both appearance and behaviour, can create an emotional bond with the robot which leads to an increasing trust in the robotic teammate. Humans, for instance, trust a polite and friendly automated system. A human user also tends to trust a robot more when they feel that the robot is highly predictable [1] [13].

Other researchers also tried to formulate a model for trust between humans and machines. Lee [6] fitted a time series model and found relationships for trust in a feedstock pump as shown in equation 2, where T refers to trust, P to productivity, F to fault size, C to some weight coefficient, and v to residual error.

$$T_n = C_1 T_{n-1} + C_2 P_n + C_3 P_{n-1} + C_4 F_n + C_4 F_{n-1} + v \quad (2)$$

In our related work [17] [18] [19], we propose a framework that attempts to determine the true time that an operator has to dedicate to the robot. Therefore, an alternative definition of the RAD as a function of both direct interaction time

(DIT) and indirect interaction time (IIT) is presented, where the IIT is a direct consequence of trust, and can represent the time being spent when the robot is independent in its work, but still with much of the user's attention drawn to it as a result of the operator's distrust in the machine. This relationship is shown in equation 3, where NT represents the neglect tolerance, and Tr is the human operator's trust in the robot. A two-level fuzzy temporal model to estimate the trust value is proposed [17]. The framework is then further augmented to propose a generic performance metric for multi-robot human interaction systems [18]. Sequential and parallel robot cooperation schemes with varying levels of task dependency are also addressed [19].

$$RAD = DIT + IIT = DIT + NT \times (1 - Tr) \tag{3}$$

In this work, we further extend our previous findings by implementing an application robotic platform where robots and users interact via natural language to complete tasks with varying levels of complexity and success. User feedback is noted and used to tune some rules in the proposed fuzzy knowledge base. The remainder of the paper is organized as follows: section II presents a brief overview of our proposed two-level trust fuzzy temporal model, along with some preliminary simulation results. Experimental setup along with the detailed implementation of the application robotic platform is described in section III. Further discussions and experimental results are presented in Section IV. Finally, section V concludes this paper.

II. Proposed Fuzzy Temporal Model

Previous discussions on trust and its evolution, along with the attempts to quantify this instrumental phenomenon in human-robot interaction teams, makes it clear that coming with a unique mathematical formula that governs the temporal behaviour of trust is going to be unrealistic, and application specific. Trust phenomenon is fuzzy by nature, and determining it depends on many factors, including time and previous state. Modelling this phenomenon using only fuzzy logic might not give the anticipated results, and therefore, a state machine is required, in which case a finite fuzzy state machine (FFSM) that combines both advantages of fuzzy logic and state machine is proposed. The size of the knowledge base, however, increases exponentially with the number of inputs. Consequently, identifying the crucial factors which contribute largely to building up trust becomes crucial. This led to the idea of building a two-level framework, shown in figure 1. Level II is the direct relation between trust and the most influential second-order perceptions that contribute to its temporal evolution. A finite fuzzy temporal inferencing model is used to model such a relationship since trust highly depends on its previous state. Second-order perceptions inputted to the FFSM are explained with lower-order perceptions, and hence should be properly modeled. In the following, we briefly present the system's interconnection and its main building blocks. Further details can be found in [17].

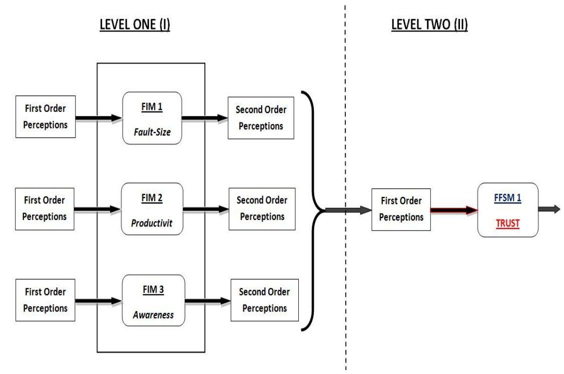

Fig. 1. Overall Architecture of the Proposed Trust Model

A. Level Two

Level II is implemented using a finite fuzzy state machine that takes three inputs (the second-order perceptions): fault size, productivity, and awareness, and outputs the corresponding level of trust based on the previous state of trust and the current perceived inputs (inferred from level I). Fault size, productivity, and awareness are modelled using three membership functions: low, medium, and high. Five membership functions are used to model trust, therefore, 5 states are required: very low, low, medium, high, and very high; and the size of the knowledge base would be: 3*3*3*5 = 135 rules.

Initially at time zero, when the human starts interacting with the robot, the trust is thought to be based on both the intrinsic human trust and reputation. Every human being has a different level of initial intrinsic trust. This value varies from one human being to another, depending on many psychological and sociological factors. People also tend to be biased with what they hear about machines. Therefore, a robot's reputation is very important in determining the basic initial level of trust it will be granted. Robots with good reputation will be more trusted when first activated, and vice versa.

The output function maps the trust states (very low, low, medium, high, very high) to the following zero-order consequences $(0.1, 0.3, 0.5, 0.7, 0.9)$. The crisp overall consequence is then generated by aggregating the qualified crisp output of each rule using the weighted average method, as described in equation 4, where w_i is the i^{th} rule firing strength, and c_i is the rule consequent.

$$\hat{c} = \sum_{i=1}^{n} w_i c_i / \sum_{i=1}^{n} w_i \tag{4}$$

B. Level One

Each of the second-order perceptions is fuzzy in nature. Those factors, however, are not temporal, thus no FFSM is needed. As a result, a fuzzy Mamdani inferencing model is dedicated for this purpose. Three membership functions are chosen to represent each of these factors: low, medium, and high. More membership functions could be used but at the expense of a more complex fuzzy inferencing system. Further details can be found in [17]. Fault size is thought to be highly determined by three lower-order factors: fault frequency (FF), fault cruciality (FC), and the ability of the robot to recover from its faults (FR). Each is represented by three membership functions, therefore, 27 rules are needed. The same applies to awareness. Three inputs mainly determine the value of this factor: machine awareness of its capabilities (MA), context awareness of the task (CA), and machine awareness of the human operator's availability and cognitive and physical abilities and limitations (HA). Another 27 rules are needed to represent this model. Finally, as for productivity, two inputs mainly determine the value of this factor: task/goal successfulness and completion (TC), and task complexity and sophistication (TS); therefore, 9 rules are required.

The two-level architecture to estimate the trust factor is crucial as it dramatically decreases the complexity of the system and the size of the knowledge base. If only a FFSM were used to model this trust factor, $5 * 3^8 = 32,805$ rules would have been needed, which is impractical and almost impossible to accurately implement.

C. Preliminary Simulations Results

In the following, we briefly present some simulation results that intuitively support our proposed trust model. Figure 2 shows the trust evolution in accordance with the

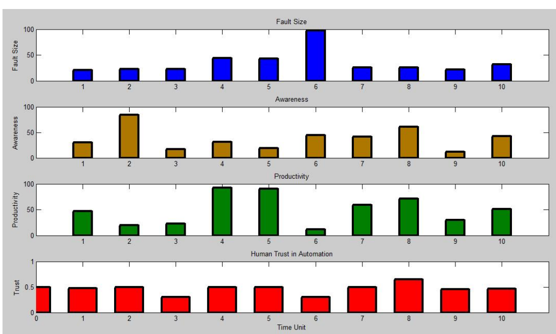

Fig. 2. Level II Trust Simulation Results -1-

temporal change of the three inferred second-order perceptions discussed earlier. Results show that at $t = 0$ (time at which the human-robot interaction starts taking place), the human trust in automation is assumed to be neutral, and the trust factor varies according to new perceptions. Results show the smoothness of change in the trust factor value, which is observed by the incremental decrease and increase in its value, rather than an abrupt change. Human trust is a step-by-step process, and strongly depends on its previous state. Results also show that when the fault size is low, and both awareness and productivity are high, the human trust in automation increases, and vice versa. Another scenario is shown in figure 3, where most extreme conditions are taken into account. Initially at $t = 0$, trust is assumed to be medium. From t = 1 to t = 3, fault size is assumed to be low, and both productivity and awareness are assumed to be high. Results show that trust only increased smoothly

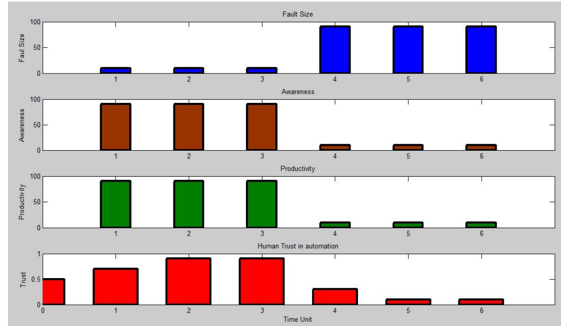

Fig. 3. Level II Trust Simulation Results -2-

in a step-by-step process, which reflects the most intuitive trust characteristics. Confirming this phenomenon, Sisodia et al. [20], and Notter et al. [12], state that building trust is a slow process, and sustaining it is always a challenge. Then, starting from t = 4 to t = 7, fault size is assumed to be high, where productivity and awareness are assumed to be low. Trust evolution shows a relatively faster decrease in trust compared to when trust was building up. The reason for this is intuitive and rooted in human psychology. People tend to lose trust much faster than they can build it. Seeing something that conflicts with our faith and beliefs makes us suspicious and cautious. Trust is hard to earn, yet so easy to lose [8].

Figure 4 shows the implications of the human trust in automation factor on both interaction time (IT) and free time (FT). DIT is assumed to be 25% of the overall task time, during which time the human user instructs and informs the robot about the task to be completed. Figure 4 shows that when the human trust in automation increases, the IIT spent monitoring the robot and interfering when needed, decreases. Results show that at time t = 2, when the human trust in automation is very high, the IIT is almost negligible. Therefore, the practical free time (during which the human operator can neglect the robot and conduct another task or instruct another robot) becomes closer to the ideal free time (obtained using only the DIT). Results also show that at time t = 4, when the human trust in automation becomes really low, the IIT becomes significantly high, and the human user is assumed to have too little time to spend on other tasks.

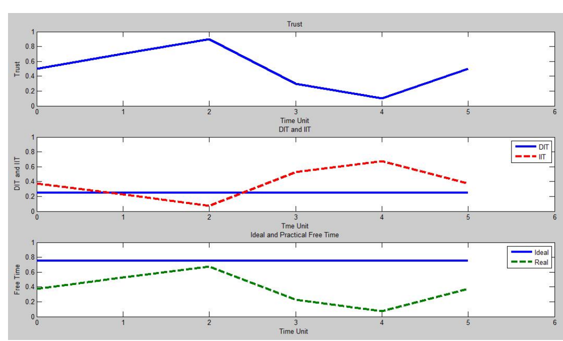

Fig. 4. True Interaction Time and Practical Free Time

III. Experimental Setup

The knowledge bases presented in levels I and II are based on a human expert's knowledge, and the most recent work in the area of cognitive human-machine interaction and performance evaluation metrics. However, and in order to enrich our proposed system, an application robotic platform that enables human-machine interaction is implemented, and users' feedback while interacting with the system was noted. The knowledge base was then fine-tuned to better reflect the user's knowledge.

A. Experimental System

In the following, we discuss the main software and hardware components of the robotic platform. The proposed platform design and implementation will follow in the subsequent subsection.

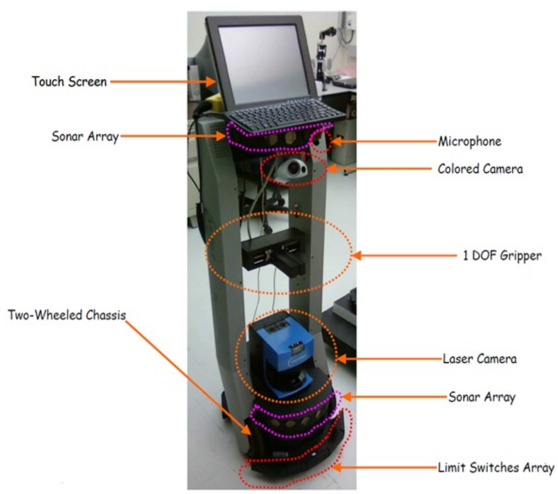

Fig. 5. PeopleBot Robotic System

1) PeopleBot: PeopleBot [16], shown in figure 5 is a differential-drive robot that is well known for human-robot interaction projects. PeopleBot is equipped with infrared table sensors and a gripper with sensors which allow the robot to pick up an object from one location and place it at another. PeopleBot features come with a laser navigation package with an autonomous robotic navigation and localization (ARNL) software that uses Monte Carlo/Markov based techniques for localization and navigation, which allows the PeopleBot to safely navigate autonomously while avoiding obstacles with great precision. It also comes with pan/tilt/zoom camera that can be used for object and people recognition, color tracking, or other robot vision tasks. This can be accomplished using the advanced color tracking (ACT) software that comes with the system. PeopleBot comes with an advanced robotics interface for applications (ARIA), which offers an API to communicate with all the robot components, control the robot's parameters, and also provide tools to integrate input/output with other custom hardware. PeopleBot SDK package also provides some tools for creating maps of a robot's operating environment, for autonomous localization and navigation.

2) Verbal Interaction: Vestec's automatic speech recognition engine (VASRE) was used in this work [21]. It is a speaker independent speech recognition engine that supports a distributed architecture of servers and clients. VASRE supports multiple languages, large vocabulary, and continuous speech recognition. Its acoustic models were trained based on continuous hidden Markov modelling. Equipped with noise reduction techniques and voice detection algorithms, it ensures smooth data input and more accurate speech recognition. The output of the engine contains such information as the raw recognized text, confidence scores, and logical parsing for generating semantic results.

B. Platform Design and Implementation

The framework supports a distributed architecture for reliable and scalable operation of clients and robot servers. The proposed distributed architecture comprises three components: the robot server, the client, and the resource manager (RM). The RM is the control tower of the distributed architecture. It manages one or more robot servers and coordinates communication sessions between servers and clients. The robot server has two states: busy or idle. A server is busy if it is executing an action upon receiving a client command. A server is idle if it is not busy. Under the idle state, the robot server periodically communicates with the RM to report its status. The RM balances server loads over different machines. For example, if a robot server A is loaded with several queued commands, the RM will guide the next client request to robot server B that has similar capabilities, if available.

The resource manager is the central entity of the framework that connects the various components to one another. As each of the robotics entities turn on, it contacts the RM registering itself along with its capabilities. In this way, the RM keeps a record of all active robots and their states, and this information can then become available to other clients and interfaces by request. The RM is also the main entity that communicates with the client. The recognition client initiates a communication session with the RM. It contacts the RM registering itself and receiving a client ID. In this way the RM keeps a record of all active clients. Then after a command is ready on the client side, the client contacts the RM asking for a command dispatch. The client might specifically ask for a specific robot, or might send a generic command that is to be executed by the first available server robot with such ability to perform the task. The RM keeps a queue of instructions for each robot. This allows multiple robots to be controlled simultaneously. Figure 6 illustrates a simplified data flow view of the RM. The robot store contains records of all currently connected robots, their actions, and locations. The task queue contains sequences of tasks that need to be executed. It consists of a list of parallel tasks, where each parallel task is a collection of tasks that need to be executed in series. The scheduler, on the other hand, is responsible for receiving task requests from clients and loading them into the task queue in order. The dispatcher loads tasks out of the task queue and sends them

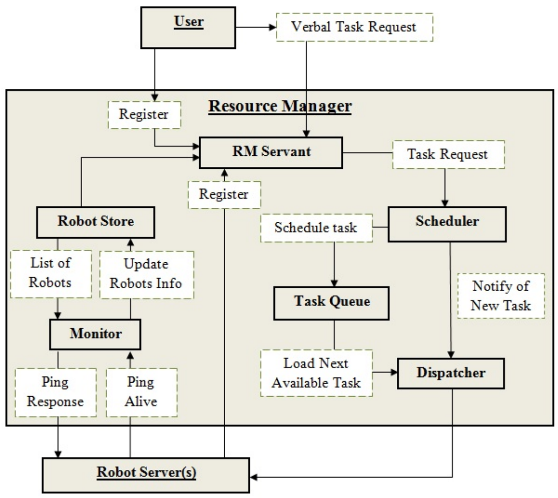

Fig. 6. Generic Resource Manager Design

to individual robots as they become available. The dispatcher can handle numerous parallel tasks simultaneously. Finally, the monitor is a separate thread that runs in the background and periodically checks with every robot to make sure it is still connected and if any information has been updated. If a change has happened, the robot store will be updated on the RM.

IV. Experimental Results

The set of experiments conducted in this work involves two Peoplebot robots, working singularly or together toward independent tasks, and a human operator. The purpose of these experiments is to support the correctness and validity of the proposed fuzzy knowledge base, and tune rules where needed to best represent the human expert's knowledge. Nine users were chosen for this purpose, and each was exposed to a set of five to six scenarios where the robot attempts to complete a set of different tasks, with varying levels of complexity, under the command and operation of the human user. Human trust in automation, along with other first- and second-order perceptions, are marked at different time units and compared to those obtained/inferred using our proposed framework. The scenarios emphasize how the human trust in automation varies with time, according to the system's success fulfilling the required tasks. The tasks vary from simple to more complex. In some scenarios, the robot is instructed to perform a series of simple tasks of moving a certain distance forward or backward, turning left or right at a certain angle, and/or controlling its gripper. More complex tasks require the robot to pick an object from a certain location and place it at a goal location. Other scenarios require the robot to locate, grab, or follow a predefined coloured object in a room, etc.

Users' perceptions were helpful to enrich the expert's knowledge base. Several rules were tuned after receiving feedback from users. For instance, some rules belonging to the productivity knowledge base were tuned when feedback showed that users tend to give more weight toward task completion than task complexity and sophistication. Figure 7 shows a sample simulation for one user, and the implication of this rule tuning on the overall percent estimation error between the practical feedback and those inferred using the Mamdani fuzzy inference model for the productivity factor. Table I shows the implication of such tuning on the remaining users. Results show some overall significant reduction of error when tuned rules are put in place. Similar findings

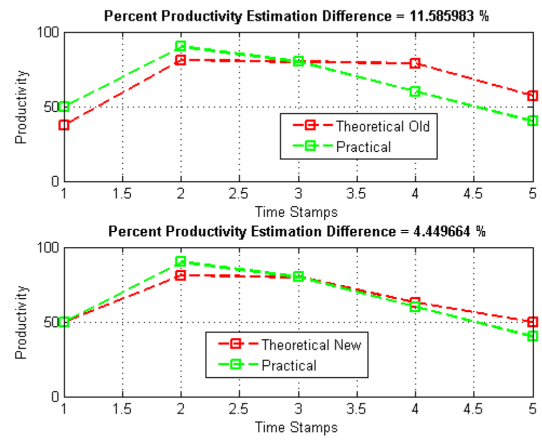

Fig. 7. Percent Productivity Estimation Error: Before and After Rules Tuning (Sample User)

TABLE I

Productivity % Error Reduction

Subject	Old Rules	New Rules	% Error Reduction
Subject#1	2.28	2.28	0.00
Subject#2	11.40	9.27	2.13
Subject#3	5.05	3.90	1.15
Subject#4	6.03	6.03	0.00
Subject#5	9.92	5.36	4.56
Subject#6	11.59	4.45	7.14
Subject#7	5.02	5.02	0.00
Subject#8	8.14	8.14	0.00
Subject#9	9.40	9.67	-0.27

were also reported for trust inference at level II of the proposed framework. Users feedback showed that users tend to generally build trust rather slower than earning it, but when the trust is already at a *very low* state, this build up process becomes even a bit slower. The implication of such observation is reported in table II. Results also show some overall significant gain in approximation accuracy. Figure 8 shows a comparison between the theoretical results obtained using our proposed framework, and the practical ones obtained from one sample user. Two independent sets of three scenarios each took place. The sets are independent and separate which explains the discontinuity at time stamp $t = 3$, which represents time stamp $t = 0$ for the second set. The first set focused on good robot performance, and successful task completion. The user's trust evolution was noted. In the second set, the user is asked to start interacting (starting with the same initial human trust in automation at $t = 0$) with the robot with a different set of scenarios, which focused

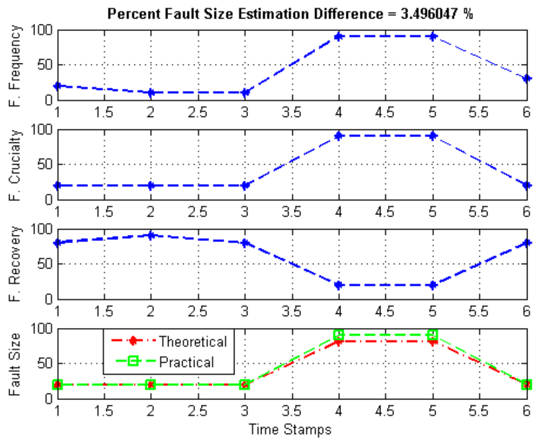

(a) Fault Size Inference

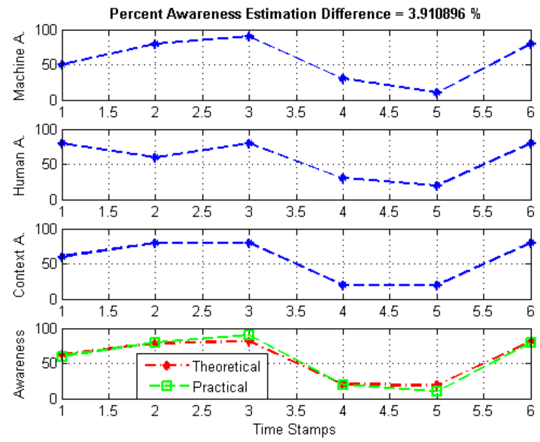

(b) Awareness Inference

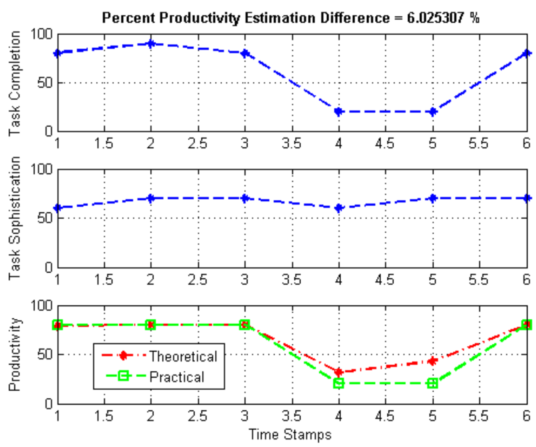

(c) Productivity Inference

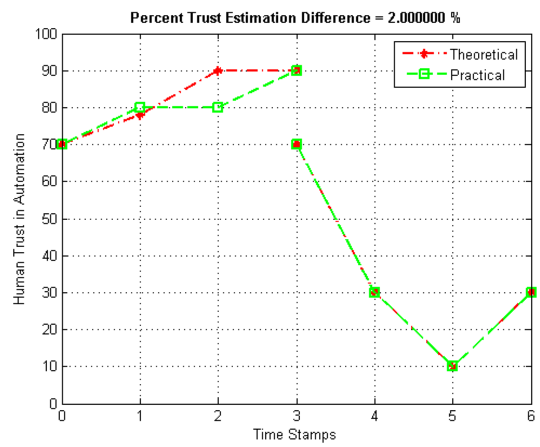

(d) Human Trust in Automation

Fig. 8. Subject #4 - Levels I and II Inferences

TABLE II
TRUST % ERROR REDUCTION

% Subject	Old Rules	New Rules	% Error Reduction
Subject #1	6.43	8.49	-2.06
Subject #2	6.38	3.04	3.34
Subject #3	7.83	4.16	3.67
Subject #4	5.33	2.00	3.33
Subject #5	6.40	6.40	0.00
Subject #6	10.60	10.60	0.00
Subject #7	8.86	8.86	0.00
Subject #8	5.29	5.29	0.00
Subject #9	15.14	9.14	6.00

on poor robot performance. The user's trust in the system automation was also noted. Figure 8(a) shows the user's first-order perceptions of fault frequency, fault cruciality, and fault recovery, along with the overall fault size. The latter value is compared to that obtained using our fault size fuzzy inference model. Figures 8(b) and 8(c) address the same manner for both the awareness and the productivity factors. Finally figure 8(d) compares the trust value as noted from the user and generated using our proposed fuzzy level II. Results show accurate trust approximation and good inferences in both levels I and II, which reflects proper and representative knowledge bases design. Similar results are also shown in figure 9, where five continuous scenarios took place with varying levels of success and completion. User feedback was noted and compared to the inferred values. Results also show that the proposed system, with its set of modified rules, is representative and within estimated accuracy. Table III shows the results for all the nine users selected in this work. The table shows the approximation errors for all inferred values, starting from fault size, to awareness, productivity, and finally human trust in automation. The results are very encouraging for the correctness of the knowledge base. Future work will further include more users to take part in this work, interacting with different types of other robots.

A. *Experienced vs Inexperienced Users*

In this part, we address another observation made when users with different expertise interacted with the robotic system. It was noted that inexperienced users with no experience working with sophisticated machines or robots provided feedback that tends to show some slight differences when compared to those obtained from more experienced users in the same experimental scenarios. Inexperienced users tend to

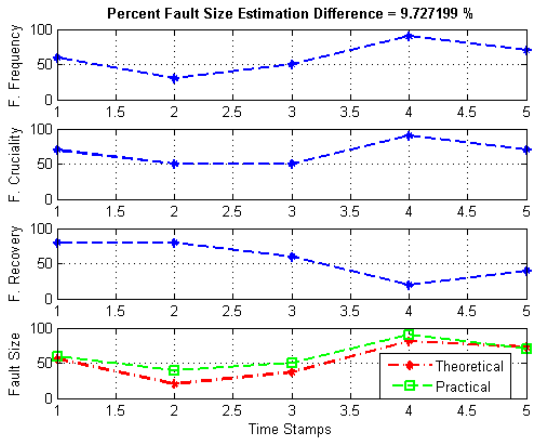

(a) Fault Size Inference

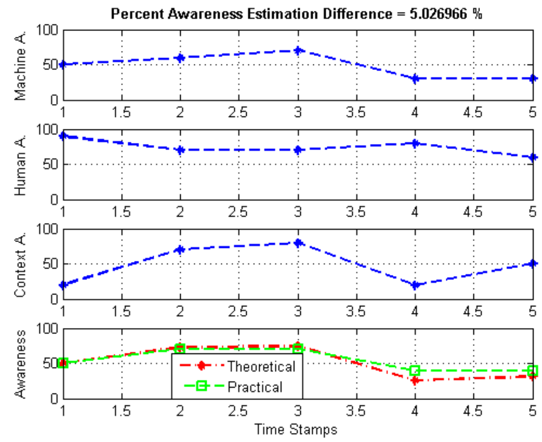

(b) Awareness Inference

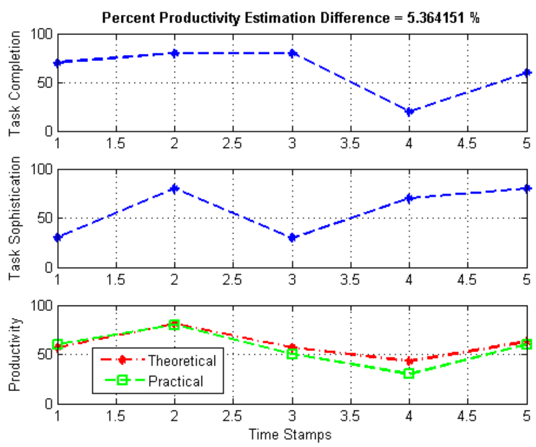

(c) Productivity Inference

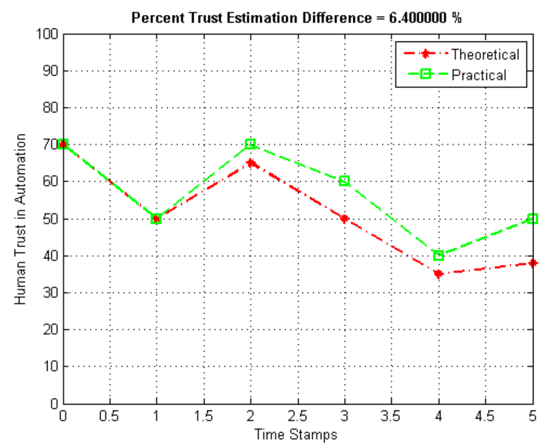

(d) Human Trust in Automation

Fig. 9. Subject# 5 - Levels I and II Inferences

TABLE III
INFERENCE % APPROXIMATION ERRORS

% Error	Fault Size	Awareness	Productivity	Trust
Subject#1	8.13	8.12	2.28	8.49
Subject#2	9.32	5.27	9.27	3.04
Subject#3	6.50	4.83	3.90	4.16
Subject#4	3.50	3.91	6.03	2.00
Subject#5	9.73	5.03	5.36	6.40
Subject#6	4.61	8.90	4.45	10.60
Subject#7	10.72	9.84	5.02	8.86
Subject#8	7.48	7.62	8.14	5.29
Subject#9	9.18	7.31	9.67	9.14
Avg Error	7.69	6.76	6.01	6.44
Std Dev	2.28	1.94	2.37	2.84

show signs of being *over impressed* with the system when it shows successful task completion, without paying attention to minor mistakes that did not affect the overall system task completion. They also get *more frustrated* with the system when it shows strong signs of incompetence. Toward this end, special considerations had to be taken into account to further accommodate this category of users to preserve the generic aspect of the proposed trust evaluation metric that is fed on its lowest level with first-order perceptions from the user. One solution would be to build another knowledge base to accommodate such an audience. This, however, adds further complexity to the system. This problem could be easily avoided with the use of some variation of the original membership functions for those inexperienced users, as shown in figure 10(b). In doing so, a total of 70 feedbacks obtained from both experienced and inexperienced users for the same scenarios are recorded and used to optimize the support set and the height of the membership functions. Four parameters are used in the optimization process: a, b, h1, and h2, as shown in figure 10(a). The optimization process is to search for an optimal combination as to minimize the total error between both experienced and inexperienced user perceptions. This is achieved by reducing the overall fuzzification error between the two set of users. Value of 10, 14, 1, and 1 were found to achieve such minimal error. This is shown in figure 10(b).

Although this preliminary suggested solution helps reducing the fuzzification error between the two set of users, this, however, does not solve the problem that users in general tend to provide subjective perceptions based on personal

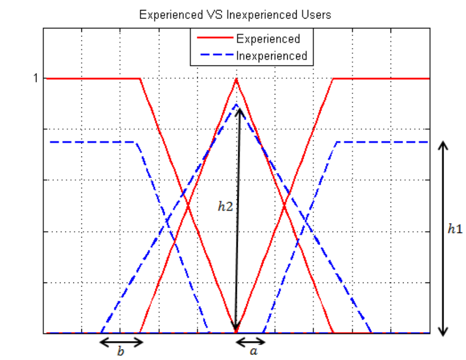

(a) Parameter to Learn

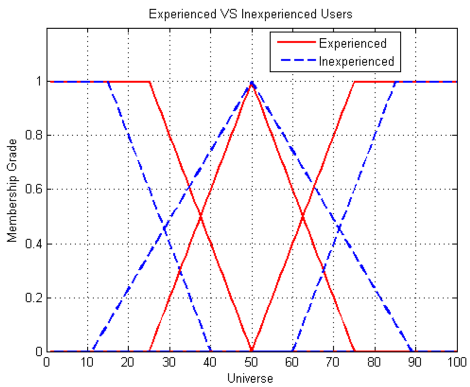

(b) Optimized Sets

Fig. 10. Experienced Vs Inexperienced MFs

and relative judgements that are highly related to factors such as experience, confidence, and level of expertise. Such perceptions vary from one person to another, thus introducing further fuzziness into the system. This observation further motivated our work to address **interval fuzzy type II sets** in our future related work to accommodate for further fuzziness in the proposed fuzzy sets. Future work shall address such extension.

V. Conclusion

Human-robot performance evaluation metrics have been receiving a large deal of the researchers' attention, especially with the fast growth in the fields of robotics and human-robot interaction systems, and the emergence of higher order functions. Robots are becoming more involved in increasingly more complex and less structured tasks and activities that require indispensable interaction with people to complete the required tasks. Therefore, designing a performance metric that can assess this effectiveness of such performance is crucial. Toward the efficient modelling of such a metric, we attempt to determine the true time that an operator has to dedicate to the robot. Therefore, we define the robot attention demand (RAD) as a function of both direct interaction time (DIT) and indirect interaction time (IIT), where the IIT is a direct consequence of the human trust in automation. Then, we propose a two-level fuzzy temporal model to evaluate the human trust in automation while collaborating with robots to complete some tasks. The fuzzy knowledge base is further updated by implementing an application robotic platform where robots and users interact naturally to complete tasks with varying levels of complexity. User feedback is recorded and used to tune the knowledge base where needed.

References

[1] M. Desai, K. Stubbs, A. Steinfeld, and H. Yanco. Creating trustworthy robots: Lessons and inspirations from automated systems. In *Proceedings of Artificial Intelligence and Simulation of Behaviour Convention: New Frontiers in Human-Robot Interaction*, 2009.

[2] A. Freedy, E. DeVisser, G. Weltman, and N. Coeyman. Measurement of trust in human-robot collaboration. In *International Symposium on Collaborative Technologies and Systems*, pages 106–114, 2007.

[3] M. Goodrich and D. Olsen. Seven principles of efficient human robot interaction. In *Proceedings IEEE International Conference on Systems, Man and Cybernetics*, pages 3943–3948, 2003.

[4] J. Y. Jian, A. M. Bisantz, and C. G. Drury. Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics*, 1:53–71, 2000.

[5] M. T. Khasawneh, S. R. Bowling, X. Jiang, A. K. Gramopadhye, and B. J. Melloy. A model for predicting human trust in automated systems. In *Proceedings of the 8th Annual International Conference on Industrial Engineering Theory, Applications and Practice*, 2003.

[6] J. D. Lee and N. Moray. Trust, self-confidence, and operators adaptation to automation. *International Journal of Human-Computer Studies*, 40:153–184, 1994.

[7] D. Li, P. L. P. Rau, and Y. Li. A cross-cultural study: Effect of robot appearance and task. *International Journal of Social Robotics*, 2:175–186, 2010.

[8] J. Love. Trust is so hard to earn, yet so easy to lose. eTapestry.com, 2009.

[9] B. M. Muir. Trust between humans and machines, and the design of decision aids. *International Journal of Man-Machine Studies*, 27:527–539, 1987.

[10] B. M. Muir. Trust in automation: Part i. theoretical issues in the study of trust and human intervention in automated systems. *Ergonomics*, 37:1905–1922, 1994.

[11] B. M Muir and N. Moray. Trust in automation: Part ii. experimental studies of trust and human intervention in a process control simulation. *Ergonomics*, 39:429–460, 1996.

[12] J. Notter and M. Blair. Conflict and trust: expressions of culture and keys to organizational success. *Journal of Association Leadership, Center for Association Leadership*, 2:535–545, 2004.

[13] S. Ogreten, S. Lackey, and D. Nicholson. Recommended roles for uninhabited team members within mixed-initiative combat teams. In *Proceedings of International Symposium in Collaborative Technologies and Systems*, 2010.

[14] K. E. Oleson, D. R. Billings, V. Kocsis, J. Y. C. Chen, and P. A. Hancock. Antecedents of trust in human-robot collaborations. In *IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support*, 2011.

[15] D. R. Olsen and M. A. Goodrich. Metrics for evaluating human-robot interactions. In *Performance Metrics for Intelligent Systems Workshop*, 2003.

[16] Mobile robots. Peoplebot. mobilerobots.com.

[17] J. Abou Saleh and F. Karray. Temporal fuzzy based modeling as applied to the class of man-machine interaction. In *IEEE International Conference on Fuzzy Systems*, 2010.

[18] J. Abou Saleh and F. Karray. Towards generalized performance metrics for human-robot interaction. In *International Conference on Autonomous and Intelligent Systems*, 2010.

[19] J. Abou Saleh and F. Karray. Towards unified performance metrics for multi-robot human interaction systems. In *International Conference on Autonomous and Intelligent Systems*, 2011.

[20] R. Sisodia, D. B. Wolfe, and J. N. Sheth. *Firms of endearment*. Business Economics, 2007.

[21] Vestec. Voice enabling systems technology. evestech.com.