

**A COMPUTATIONAL MODEL OF HUMAN TRUST
IN SUPERVISORY CONTROL OF ROBOTIC
SWARMS**

by

Huao Li

Submitted to the Graduate Faculty of
the School of Computing and Information in partial fulfillment
of the requirements for the degree of
Master of Sciences

University of Pittsburgh

2019

UNIVERSITY OF PITTSBURGH
SCHOOL OF COMPUTING AND INFORMATION

This thesis was presented

by

Huao Li

It was defended on

August 21st 2019

and approved by

Michael Lewis, School of Computing and Information

Stephen Hirtle, School of Computing and Information

Hassan Karimi, School of Computing and Information

Thesis Advisor: Michael Lewis, School of Computing and Information

A COMPUTATIONAL MODEL OF HUMAN TRUST IN SUPERVISORY CONTROL OF ROBOTIC SWARMS

Huao Li, M.S.

University of Pittsburgh, 2019

Trust is an important factor in the interaction between humans and automation to mediate the reliance action of human operators. In this work, we study human factors in supervisory control of robotic swarms and develop a computational model of human trust on swarm systems with varied levels of autonomy (LOA). We extend the classic trust theory by adding an intermediate feedback loop to the trust model, which formulates the human trust evolution as a combination of both open-loop trust anticipation and closed-loop trust feedback. A Kalman filter model is implemented to apply the above structure. We conducted a human experiment to collect user data of supervisory control of robotic swarms. Participants were requested to direct the swarm in a simulated environment to finish a foraging task using control systems with varied LOA. We implement three LOAs: manual, mixed-initiative (MI), and fully autonomous LOA. In the manual and autonomous LOA, swarms are controlled by a human or a search algorithm exclusively, while in the MI LOA, the human operator and algorithm collaboratively control the swarm. We train a personalized model for each participant and evaluate the model performance on a separate data set. Evaluation results show that our Kalman model outperforms existing models including inverse reinforcement learning and dynamic Bayesian network methods.

In summary, the proposed work is novel in the following aspects: 1) This Kalman estimator is the first to model the complete trust evolution process with both closed-loop feedback and open-loop trust anticipation. 2) The proposed model analyzes time-series data to reveal the influence of events that occur during the course of an interaction; namely, a users

intervention and report of levels of trust. 3) The proposed model considers the operators cognitive time lag between perceiving and processing the system display. 4) The proposed model uses the Kalman filter structure to fuse information from different sources to estimate a human operator's mental states. 5) The proposed model provides a personalized model for each individual.

TABLE OF CONTENTS

1.0 INTRODUCTION	1
1.1 Human-automation trust	1
1.2 Supervisory control and level of automation	2
1.3 Robotic swarm system	3
1.4 Thesis structure	4
2.0 RELATED WORK	5
2.1 Theoretical Trust Models	5
2.2 Computational Trust Models	10
3.0 HUMAN EXPERIMENT	12
3.1 Supervisory Control of Robotic Swarms	12
3.2 Human-Swarm Interface	13
3.2.1 Level of autonomy and Search mode	14
3.3 Human Experiments	15
3.3.1 Experimental Design	15
3.3.2 Participants	16
3.4 Experimental Results	17
3.4.1 Survey results	17
3.4.2 In-progress trust feedback	17
3.4.3 Performance	18
3.4.4 User intervention commands	19
3.4.5 Mode switch in the MI LOA	21
4.0 HUMAN-SWARM TRUST MODEL	22

4.1	Kalman Estimator	22
4.2	Model Parameterization	25
4.2.1	Process and measurement noise	26
4.2.2	Control parameters	27
4.2.3	Perception parameters	28
4.2.4	Time delay	28
4.2.5	Grid search	29
5.0	TRUST MODEL EVALUATION	30
5.1	Data Process	30
5.2	Model training	31
5.3	Comparison with existing model	31
5.4	Discussion	32
6.0	CONCLUSIONS	35
7.0	Bibliography	38

LIST OF TABLES

- 1 The result of trust predictions (RMSE). The Kalman filter model is compared with the IRL model proposed in [20] and DBN model in [38]). Fig. 10 shows the example prediction results in graphs of three individual participants. . . . 32

LIST OF FIGURES

1	Lee and See’s dynamic model, including the closed-loop trust evolution and factors that influence this process.	6
2	Kalman estimation trust model.	8
3	An illustration of the swarm simulator. The swarm navigates in the unknown area to find targets. The participants adjust the trust slider on the right panel by using the mouse wheel to give <i>trust feedback</i> as their trust changes. The left panel shows task-related information.	15
4	Some swarm parameters: (a) the variance of the heading angles of all swarm members; (b) the convex hull area that the swarm makes.	16
5	Results from surveys. Error bars are 1 Standard Error from means (SEM). (a) The participants had a significantly low trust towards the autonomous LOA than the MI LOA. (b) The workload of the autonomous LOA is much less than the other two LOAs. (c) The participants had the smallest negative trust change in the MI LOA.	18
6	Results from in-process measurements (error bars are SEM). (a) The participants had a significantly lower average trust feedback values in the autonomous LOA while the manual LOA had the highest trust feedback. (b) The task performance of the manual LOA was significantly higher than the autonomous LOA.	19
7	The average trust feedback when users issued intervention or nonintervention commands. Error bars are 1 standard error from the mean.	20

8	The average trust feedback when the mode switches occurred (user-initiated vs. system-initiated switches). Error bars are 1 standard error from the mean.	21
9	A block diagram of the implementation of a Kalman filter.	26
10	Example predicted results on the test data set. Each graph shows the prediction result of one participant and the x and y axis represent the time step and the trust level, respectively. The blue line represent model prediction while green is user trust feedback values.	33

1.0 INTRODUCTION

1.1 HUMAN-AUTOMATION TRUST

Trust is an important factor when human operators interact with automation. The term refers to a human's belief in automation's capability and the willingness to rely on automation in uncertain situations [14]. Human operators need to trust automation enough to rely on it in conditions where this reliance would lead to improved performance [15, 13]. The situations in which people fail to do so are recognized as *distrust* or *undertrust* because the human lacks confidence in automation. As a result, the operator may undertake additional tasks that could have been delegated to automation or issues unnecessary intervention, both of which can harm joint performance. A running example could be a pilot who distrusts the auto-pilot system, and tends to monitor all the state parameters, double-check the output of automated systems, or even chooses to manually control the plane. Each additional task increases the cognitive workload and occupies the mental resources of the pilot, which could have been used instead for high-level activities, like communication and planning.

The other maladaptive attitude is called *overtrust*, where human operators hold unrealistic expectations of the capabilities of automation and thus over-rely on automation. Under such circumstances, humans may fail to properly monitor automation or blindly accept the recommendations of automation because they are overconfident about the reliability of automation systems [23]. Overtrust in automation prevents human operators from intervening at necessary moments and can lead to severe consequences [16]. In the pilot case, the operator might overestimate the capacity of autonomous navigation system such that they accept the advice without serious consideration or fail to respond to emergencies in a timely fashion.

Therefore, in human-automation interactions, a well-calibrated level of trust in automa-

tion is needed to optimize collaboration and overall joint performance. As autonomous systems become more intelligent and self-governed, their behaviors are no longer restricted to the specific actions for which they have been designed and the situations to which they need to respond also go beyond pre-programmed or anticipated cases. More flexibility comes with more complexity in analyzing and modeling human trust in progress. A computational model of human trust is necessary for autonomous systems to adapt to human operators and calibrate their levels of trust. For example, when an auto-pilot system detects the dropping of human trust, it may present a more defensive strategy or display more explanations for its behavior to regain trust from the operator.

1.2 SUPERVISORY CONTROL AND LEVEL OF AUTOMATION

This thesis focuses on the supervisory control of robotic systems, which signify a supervisor-worker relationship between human operators and robots [37]. In our task scenario, both the human and robot work collaboratively toward a shared task goal. The autonomous agent is capable of controlling robots to accomplish given tasks, while the human supervises the whole progress and intervenes when necessary. Depending on different levels of automation (LOA), the human operator and the autonomous system collaborate in different ways. For example, in a relative low LOA, a human needs to issue commands to change the heading direction of unmanned vehicles while the system automatically moves them to that direction while avoiding obstacles along the way. When it comes to a higher LOA, the agent is able to self-navigate vehicles around different areas to finish given tasks, but the operator can intervene by taking direct control.

When supervising autonomous systems with diverse LOA, a human operator employs different mental models and trust evolution processes. For instance, previous research has indicated that participants give a more "performance-centric" trust report when they are only required to passively monitor the system, as compared to those in the active control LOA [22]. The author proposed that in higher LOAs, participants were saved from constant motor control and therefore had more cognitive resources to better perceive the overall status

and performance of automation. On the other hand, in systems that can operate either automatically or manually, control takeover is often considered a signal of trust dropping [15, 13]. For supervised automation systems that cannot be manually controlled, an operator’s interventions in the ongoing tasks can be interrupted in a similar way. Therefore, it is important to continue exploring the relationship between system LOA, control takeover, and user intervention and human trust, and especially how those events that occur during the course of interactions influence the time series of trust evolution.

1.3 ROBOTIC SWARM SYSTEM

A supervisor-worker relationship exists in diverse real-world contexts [31] e.g. remote control of unmanned ground vehicles, human drivers operating a self-driving car, and human-robot teaming. In this thesis, we concentrate on the control of swarm robots.

A robotic swarm is a group of simple, typically homogeneous robots that is capable of accomplishing complex tasks. Individual swarm members are coordinated via local control laws to form global behaviors (e.g. flocking, deployment, and rendezvous) which enable the swarm to coherently interact with each other and the environment. Swarms benefit from their decentralized nature, in that the system is robust to individual failure and does not require extraordinary individual capability. Hence, swarms are theorized to be important for large-scale applications in unknown and dynamic environments, including environmental monitoring [6], structure inspection [11], search and rescue [25], and even space exploration [32].

From a human point of view, interacting with swarm robots is tremendously different from controlling a single robot because of the unique characteristic of swarms. For instance, the nonlinear dynamics of swarm systems [3] has been shown to prevent a human operator from correctly perceiving the swarm’s state or performance [35], and further, issuing interventions in time (a.k.a. neglect benevolence [19]). Moreover, unlike automation systems that are directly controlled by operator commands, swarms can only be indirectly influenced by changing their control laws. In other words, human interventions are not directed

to each group member to tell them the heading direction and velocity; rather, operators must issue commands to the whole swarm by changing their emergent behaviors or parameters [17]. Such a feature inevitably extends the lag time of system feedback and complicates the decision-making process of human operators. When considering the perception difficulties and indirect control present in human-swarm interaction, there are significant time delays in multiple steps within the operator’s cognitive processes. Most importantly, when modeling human trust during this process, it is necessary to take additional feedback paths and their own unique time lag into account.

1.4 THESIS STRUCTURE

This thesis will focus on the problem introduced in the previous section: modeling dynamic trust evolution during supervisory control of robotic swarms. The following section begins with an introduction of related work in Chapter 2. In addition to previous research that only focuses on human control takeovers and interventions as the signal of a loss of trust, we track the dynamic evolution of human trust in both directions. With the help of such a computational model, an adaptive swarm system is envisioned to calibrate human trust to optimize the overall joint performance. After the literature review, Chapter 3 shows the experiment design and the observational results of supervisory swarm control. During the experiment, we measure the operator’s in-progress trust during the course of their interaction, which is believed to have higher reliability and validity. A novel trust model based on a Kalman filter is introduced in Chapter 4. When considering the different cognitive paths in processing system feedback and anticipated consequences of intervention in human swarm interaction, two distinct modules with time lags are constructed in the model to represent both processes. Based on the data collected from human experiments, the performance of our Kalman-based model is verified and compared with other existing trust models in Chapter 5. Finally, Chapter 6 discusses the contributions of this work and possible directions for future research.

2.0 RELATED WORK

2.1 THEORETICAL TRUST MODELS

Lee and See [14] have provided a thorough review of early research on the topic of *trust in automation*. A formal definition of trust is given as, “the attitude that an agent will help achieve an individual’s goals in a situation characterized by uncertainty and vulnerability” considering its major characteristics. In addition, Lee and See also summarized the basis of trust in three dimensions: purpose (why the automation was developed), process (how the automation operates), and performance (what the automation does). Those three elements provide important insights that human trust is not only influenced by whether the robot could finish a given task, but also by the degree of the operator’s understanding of what the robot is designed to do and how it functions.

Fig. 1 shows a fundamental dynamic model of trust and reliance on automation proposed in Lee and See’s work. This theoretical model consists of a closed-loop evolution of trust, the influence of context, and the role of information display in calibrating trust. In the core closed-loop system, human operators perceive the physical state of the system from a display and form their own beliefs about the system’s state. The trust level is then established based on the belief in the automation’s capability and the current state. Based on the trust level, the operator may form the intention to either use or intervene with the automation and finally take the selected action. The action that operators take will affect the state of automation.

The bottom line of this closed-loop structure is that the dynamic interaction with the automation influences trust and trust influences the dynamic interaction. When the operator uses or relies on the automation, they can observe information about it and therefore have a better chance to establish a robust level of trust in the automation [18]. Having a good

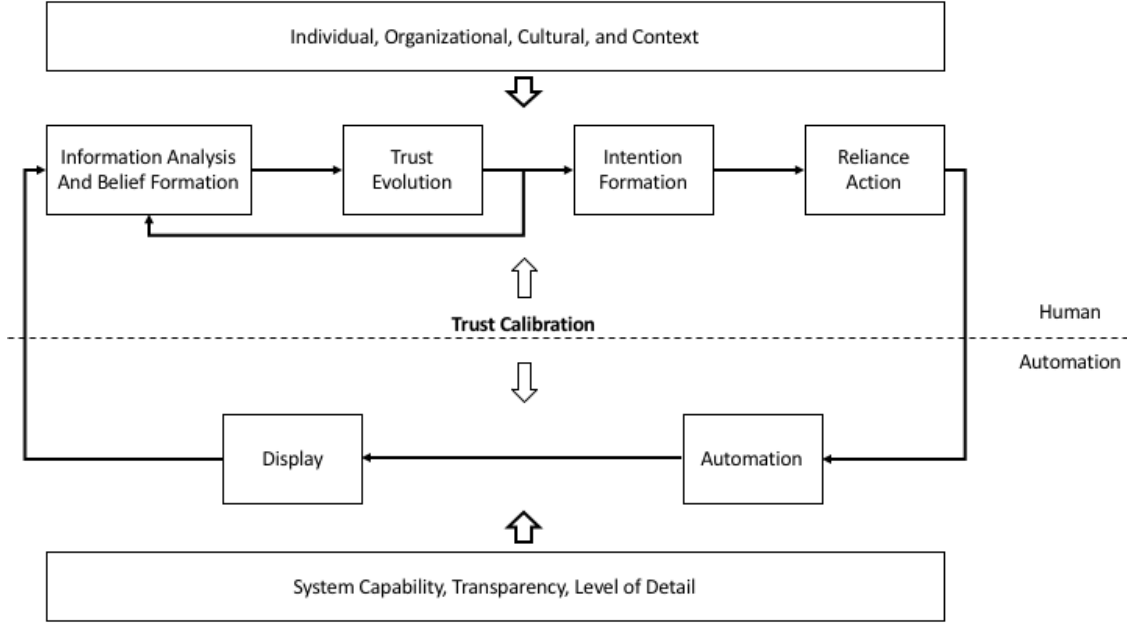


Figure 1: Lee and See’s dynamic model, including the closed-loop trust evolution and factors that influence this process.

understanding of the system’s capability and working conditions is vital to develop greater trust in the automation. Unlike most theories that analyze the decision-making process, Lee and See’s dynamic trust model is more suitable for analyzing trust in automation because it better reflects the factors that influence reliance and their effect over time. Static approaches that have been used to identify the mis-calibration of self-confidence in decision-making address only cumulative experience, rather than the evolving experience and continuous re-calibration that are critical for an appropriate level of reliance on automation.

In addition, the type of automation also has a huge effect on the evolution of human trust. Depending on the stages of information processing that it involves, the automation system can assist humans in information acquisition, information analysis, decision selection, and action implementation [24]. This difference has been argued as the level of automation (LOA) that combines the degree of automation in different types and stages [33]. Lee and See proposed that different LOAs will significantly change the trust dynamics, especially the observation process. For example, it is possible for operators to observe the behavior of

information acquisition automation, even when they are not using it directly, because both the raw and processed data are available. However, it is difficult for operators to observe action implementation automation unless they are relying on it directly. This difficulty in perceiving system performance may lead operators to fail to recover trust after adapting manual control mode, even when the reliability of the automation improves [15]. Therefore, it is important to compare the trust dynamic evolution in different LOAs.

An important consideration of trust in automation is the effect of system failures. The occurrence of faults usually leads to the development and erosion of trust as a dynamic process. Depending on the type and severity of the failure, trust may decline and recover shortly (when the fault is mild and temporary) or decline until operators subjectively accommodate it (when the fault permanently harms the capability of the automation) [10, 15]. The dynamic changes in trust that are brought about by failures do not happen immediately, but occur over a period of time [15]. As noted in an early work of Lee and Moray [13], a time-series analysis has shown that the influence of automation failures on human trust can be modeled by a first-order differential equation, in which the largest effect will be seen immediately, with a residual effect distributed over time. This time-series analysis identifies the time constant of trust and determines how quickly the trust changes to reflect changes in capabilities of the automation.

Recently, Sheridan summarized three different types of approach analyzing trust in automation, including signal detection, statistical parameter estimation and model-based control [29]. As proposed in the paper, Lee and See's trust model frames human trust as a closed-loop model with six elements, as shown in Fig 1. The model structure can be easily mapped with a classical control process with a minimum modification, as proposed in [30]. The trust evolution of the human operator refers to the internal model of reality that needs to be estimated. Based on the trust level, the operator may form an intention to intervene in the operation and finally take an action to influence the state. Those two elements represent the state-based policy of action and physical action that modify the state in control theory. The action that operators take will affect the system state of the automation. The physical changing of reality will be feedback to the operator via the system display, which corresponds to the measurement process in the language of control. Finally, the human operator takes

of trust anticipation and verification. First, the degree of trust is equivalent to the size of the open-loop commitment and the anticipation made after the action of decision-making. The closed-loop trust verification is equivalent to the difference between the actual displayed system state and the internal desired state. As commented in [29]: ‘In this trust model, the control law determines whether or how far to commit to open-loop action, based on the internal model of the current state of the automation’s trustworthiness and in consideration of ones vulnerability.’

Another concern in the evolution of trust is the time delays between perceiving and processing information in the human cognitive system. Unlike classical control systems, in which this value tends to be quite short, it takes much longer for human operators to correctly perceive the displayed information and form their own belief of trust in automation. For a longer time delay, a longer open-loop decision making process is taken, based on the anticipation of automation. When the feedback from the actual system eventually arrives after T time steps, the operator can use this feedback information to verify their own trust belief. The longer the time lag, the lower the expectation that the internal belief and feedback will match, and the more vital it becomes to involve both the open and closed loops to better estimate the overall trust state.

This trust-action-verify structure is commonly used in modeling motor operating with a response time lag. For example, factory operators making machine settings may not able to see results for minutes or even hours, due to the slow response of the machine. The same is true for teleportation robots in space: human supervisors give open-loop trust commands and waits for closed-loop verification after several seconds of delay. As we can see, when the time lag grows larger, there is more room for the operator to consider their trust criteria and involve more cognitive consideration. In a human-swarm interaction, the unique physical characteristics of swarms bring additional difficulties for operators to correctly perceive their current behaviors and operational states. In addition, the control algorithm and communication constriction makes the swarm group slower in response command compared to single robots. Therefore, a Kalman system that considers both open-loop trust anticipation and closed-loop trust feedback is the optimal choice for modeling human trust in supervisory control of robotic swarms.

2.2 COMPUTATIONAL TRUST MODELS

In the work of Lee, a computational model based on the extended decision field theory was developed to predict operator trust in and reliance on supervisory control [8]. The application scenario had operators monitoring the operation of an orange juice pasteurization plant, where a human can choose to either monitor automatic control or intervene manually. The model tracked both the operator's trust toward the automation and self-confidence of their manual control capacity. The decision of whether or not to rely on a specific control mode was made upon the accumulated difference between those two values over a period of time. Basically, each of the two processes was represented by a closed-loop dynamic, as in the classical model previously mentioned. If the operator holds a great estimation of automation's capacity, they would have a high trust towards it and would be more willing to rely on the automatic control. While the operator is using automatic control, they gain information from the system display and update their belief about the system's capability. On the other hand, the higher self-confidence that an operator has on their manual control skills, the more likely they would be to choose to manually control the system. Thus, we could have a preference indicator by calculating the difference between trust in automation and self-confidence at a given time step. According to the decision field theory, while this valence difference accumulates, an operator's preference towards a certain control model becomes stronger. When it finally goes beyond a certain threshold, the operator would make the decision to either relay or intervene.

This computational model applied the closed-loop dynamic structure in Lee and See's theoretical trust model and provided decent results in predicting human reliance on automation. More importantly, it captured the accumulation of trust over time and considered the influence of trust levels on decision making. However, this model has its own limitations when applying to general task settings. First, this model focuses on the interaction between a single operator and a single automation. When the human operator interacts with a multi-agent system, the cognitive workload may increase and the mental model may become more complex. Second, this model considers the use of automatic and manual control as purely complementary, while this may not be true for many systems. Third, the parameters in this

model are predefined without learning or adaptation to either the automation system or to a human operator. However, many parameters such as the decision making threshold and manual capabilities can vary widely from operator to operator.

Xu proposed a dynamic Bayesian inference trust model for a single robot [38]. The model uses a robot’s performance to predict the operator’s latent trust state and constructs a Bayesian network to consider influence from past time steps. The participants’ in-progress trust report was used as the measurement of the latent trust state. The above model structure is based on the assumption that the changes in trust rise and fall, based on the task performance of the robot. In Xu’s paper, the human operator was asked to monitor a UAV taking a given boundary and to take over the control whenever they noticed a failure. Since the task setting was fairly straightforward, the performance of robot can be easily recognized by humans with no significant delay. The takeover behaviors of human operator also directly indicate a loss of trust in automation. However, in the swarm control tasks, the automation performance was not intelligent to humans, due to the characteristics of swarms. Humans need more time to process complicated swarm behaviors, which brings a longer time delay in the decision-making process. As a result, the effect of open-loop trust anticipation is even more dominant and should be taken into consideration.

In [21], a computational trust model for human swarm interaction was proposed on the basis of inverse-reinforcement learning. The task scenario in this research is a foraging task, where the human operator was asked to control the swarm to search for hidden targets in an unknown environment. In this case, because human interventions may not be necessarily related to a drop in trust but may instead indicate a changing of intentional searching areas, a classifier was developed to distinguish human interventions caused by trust loss and intention shifting. The human trust evolution process is modeled as a Markov decision process (MDP), in which the state space consists of the task performance, trust-related human intervention, and swarm state. The model learns an individualized reward function from each operator’s event log and uses it predict human trust at given states. In a later work [20], the author extended this approach to human swarm systems with different LOAs. Although this model provides an accurate and direct trust prediction, it does not consider the temporal sequence of events that occur during human-swarm interaction.

3.0 HUMAN EXPERIMENT

3.1 SUPERVISORY CONTROL OF ROBOTIC SWARMS

We considered a target search task where a swarm of robots is controlled by a human operator to explore an unknown environment with static obstacles. Each robot is assumed to have the capability of sensing surrounding environments, but with a limited field of view. The human operator controls the swarm remotely via command inputs, as they select the heading direction towards which the swarm should move. The swarm has leader robots who receive command inputs from the operator and communicate heading directions towards which the swarm should flock to swarm members via peer-to-peer communication. The swarm is equipped with basic obstacle avoidance function, but is not able to change its own heading direction; for example, escaping from a corner. Observable swarm parameters include velocity, centroid, connectivity, mean heading angle, heading variance, and convex hull area (Fig. 4).

We have three different levels of autonomy (LOAs) of the swarm: (1) the manual LOA, (2) the autonomous LOA, and (3) the mixed-initiative (MI) LOA. In the manual LOA, the human operator needs to give heading directions for the swarm to navigate (i.e., *manual search mode*). In the autonomous LOA, the swarm searches the map by itself using a search algorithm [4], while the operator is out of the control loop (i.e., *autonomous search mode*). The search algorithm enables the swarm to search the entire space of an unknown area if the time is sufficient. In the MI LOA, the operator is allowed to use both the manual and autonomous search modes mentioned above to optimize the searching performance of swarm robots. In addition, a recommendation system is used to automatically switch the control mode, based on the performance within the current task. The system may recommend a

mode switch if the performance remains low in the current mode; however, the operator may or may not follow the recommendation and can switch the mode at any time, even if there is no recommendation to switch. Here, we differentiate *search mode* and *LOA* as different concepts. In the manual or the autonomous LOA, only the manual search mode or autonomous search mode is accessible, respectively. In the MI LOA, the operator is able to access a mix of the manual and autonomous search modes, which is determined by interaction with the swarm and the recommendation system.

3.2 HUMAN-SWARM INTERFACE

A swarm simulator [34] was used for testing human-swarm interactions. Figure 3 illustrates the simulation interface. The center panel gives the operator a top view of the undiscovered map. Swarm robots, targets and static obstacles are represented by red dots, green stars, and gray rectangles on the map respectively. The line between swarms refers to the peer-to-peer communication link. The map is covered at the beginning of each trial, and will be explored as long as the swarm can move and sense the environment. The left panel shows additional task-related information, such as the time remaining for the current trial, the cumulative number of targets found, a low-performance alert, a mode switch countdown given by the recommendation system, and the current search mode (manual or autonomous). There is a trust slider on the right panel of interface that the participants can adjust, using the mouse wheel, to indicate their current subjective trust ratings. The participants were required to report their level of trust toward swarms (*trust feedback*) every 30 seconds on a scale from -10 (strongly distrust) to $+10$ (strongly trust). They were also encouraged to adjust this value when they felt that their level of trust changed.

The swarm consisted of 32 homogeneous robots that were set at random positions at the center of a $200\text{ m} \times 200\text{ m}$ environment. In each trial, random environmental configurations were generated for robot poses, obstacles, and targets. The operator could give commands to navigate the swarm during flocking by dragging a line (the purple line shown in Fig. 3) to indicate the heading direction. In different search modes, the swarm received the heading

for flocking [26] either from the operator or the autonomous search algorithm. The data from user command inputs (the angle and length of command vectors), swarm parameters (the mean and variance of heading angles of the robots shown in Figure 4a, convex hull area defined by the robots shown in Fig. 4b, connectivity), user or system initiated mode switch, and the number of targets found were recorded for each time step (60 Hz).

3.2.1 Level of autonomy and Search mode

During the supervisory control task, the participants could use two different search modes to finish the task, depending on the LOA of the system. In the manual LOA, only the manual search mode was available to the operator, in which participants could give the swarm a heading direction by dragging a line on the screen using the mouse. In the autonomous LOA, the swarm could only be controlled by the algorithm in the autonomous search mode.

The MI LOA offered a flexible interaction between the human and the swarm by appropriately adjusting the search mode. The simulator initially started with the autonomous search mode, where the participants had no control over the swarm. However, the operator could switch to the manual search mode by giving a mouse input or pressing a toggle key. In the manual search mode, the participants could also use the toggle key to switch back to the autonomous search mode.

In any mode, current task performance (i.e., number of targets found) both in total and during the last 15 seconds (the red text on the left panel in Fig. 3) were shown on the interface. If the swarm found less than three targets for the last 15 seconds (the predefined threshold was determined by a pilot study), the recommendation system would prompt an alert for low performance (in red) on the left panel. If the low performance continued for 10 seconds, the system recommended that the participants switch the mode in another 10 seconds (a countdown appeared on the left panel along with the recommendation). If the operator did not initiate a mode switch before the countdown ended, the system switched the mode automatically. However, the participants could reverse the forced switch by pressing the toggle key or by giving a heading direction. The current search mode is always shown on the interface in text with a color coding that corresponds to the map boundary to increase

its visibility (e.g., the green text in the left panel and the green bounding box in Fig. 3).




Figure 3: An illustration of the swarm simulator. The swarm navigates in the unknown area to find targets. The participants adjust the trust slider on the right panel by using the mouse wheel to give *trust feedback* as their trust changes. The left panel shows task-related information.

3.3 HUMAN EXPERIMENTS

3.3.1 Experimental Design

The experiment employed a 3-level within-subject design, in which each participant ran three different LOAs in a counterbalanced sequence. In the beginning of experiment, participants were asked to finish a survey in order to measure their general trust towards autonomy (*trust pre-test*). The questionnaire was adopted from [5] considering three trust components (performance, process, and purpose) in 5-Likert scale. For each of the three LOAs, a 2-min training session and three 5-min identical trials were given to the participants. The main task of participants was to navigate the swarm through the environment to discover 100 initially hidden targets. Participants were told to find as many targets as possible within



Figure 4: Some swarm parameters: (a) the variance of the heading angles of all swarm members; (b) the convex hull area that the swarm makes.

the given time (5 minutes). After finishing three trials in each LOA, the participants were asked to fill out a survey to collect their subjective trust towards the swarm that they just interacted with (*trust post-test*) and a NASA-TLX survey [9] to measure their workload. The participants were told to consider the swarm (e.g., individual robots and the search algorithm) and the system (e.g., the interface and the alert/recommendation) as a whole when they rate trust. The experimental procedure lasted for 75 minutes.

3.3.2 Participants

20 participants were recruited from the University of Pittsburgh and Carnegie Mellon University communities with an average age of 24.1 ($\sigma = 2.89$). Each of them were paid \$10 to finish a 75-minute experiment. The experiment’s protocol was approved by the University of Pittsburgh Institutional Review Board. Participants had no prior experience with controlling a swarm of robots.

3.4 EXPERIMENTAL RESULTS

The results from the experiment have been published in [22, 20], which focuses on user’s different trust feedback and behavior among LOAs. In the present work, the analysis of the results concentrates more on identifying trust-related factors in order to build the computational model.

3.4.1 Survey results

An one-way repeated measures ANOVA was run to analyze the post-test trust survey data. Results showed that there was a significant difference in the post trust report between LOAs ($F(1.37, 19.20) = 7.80, p = 0.007$, see Fig. 5a). Pairwise comparison showed that the participants had a significantly lower trust towards the autonomous LOA than the MI LOA ($p = 0.001$). Workload is also shown to be significantly different among LOAs ($F(1.38, 19.37) = 13.52, p = 0.001$, see Fig. 5b), in which the workload of the autonomous LOA is much less than the other two LOAs ($p = 0.023, p = 0.001$). This result is consistent with our hypotheses and confirms previous findings on trust [1, 27] and workload [7] with varied LOAs in a new domain of swarm supervisory control.

To take participants’ preexisted trust levels towards the autonomy into consideration, we compared the difference between the pre- and post-trust survey (*trust change*) (Fig. 5c). In all LOAs, trust decreased after the participants experienced swarm control which may due to the low controllability and intelligibility of swarm behaviors. Among all LOAs, the MI LOA had the smallest negative change in trust. The survey results of trust and workload show that the mode switching and recommendation in the MI LOA neither damaged trust towards the swarms nor increased operator’s workload.

3.4.2 In-progress trust feedback

The average trust feedback values (i.e., the mean of the in-process ratings of trust) had a significant difference among the three LOAs (one-way ANOVA, $F(2, 57) = 3.35, p = 0.0423$, shown as Fig. 6a). The participants had a significantly lower trust feedback values




Figure 5: Results from surveys. Error bars are 1 Standard Error from means (SEM). (a) The participants had a significantly low trust towards the autonomous LOA than the MI LOA. (b) The workload of the autonomous LOA is much less than the other two LOAs. (c) The participants had the smallest negative trust change in the MI LOA.

in the autonomous LOA whose mean was 2.571 while the manual LOA had the highest trust feedback (the mean was 5.086). The MI LOA’s mean trust feedback was 4.014. A likely explanation for the low trust in the auto LOA is that the searching algorithm is lack of transparency and the participants did not like being excluded from the decision-making loop.

3.4.3 Performance

The most common thought about human trust in automation is goal-oriented, which establishes trust in a robot’s performance in finishing given tasks.

The average task performance in different LOAs have a significant difference ($F(2, 57) = 55.18, p \ll 0.01$). The means of the targets found in the three LOAs are 74.3 (manual), 66.4 (MI), and 55.6 (autonomous), respectively (Fig. 6b). The result indicates that the autonomous search algorithm did not outperform humans in the given environment, perhaps owing to the presence of obstacles.

To identify if the current task performance influences the user’s reported level of trust, we run a correlation analysis. The result shows that the correlation coefficients between task performance and trust were low ($-0.0768, -0.0140, 0.1353$ in the manual, MI, autonomous LOA respectively). This finding reveals an important fact: in human-swarm interaction, the

performance of swarm robots is less intelligible to the operator in terms of evaluating trust. As noted in [22], the different relationship between task performance and trust among LOAs may be due to the operator’s diverse workload. For instance, the decreased workload in the autonomous LOA would enable the participants to correctly perceive task performance. An evidence is that, within the MI LOA, there was no statistically significant correlation between the performance and trust feedback ($R = 0.1230$, $p = 0.3620$) in the manual search mode. However, for the autonomous search mode, a correlation ($R = 0.2715$, $p = 0.041$) indicates that participants were able to align their in-process trust feedback to task performance when they were not actively engaged in controlling the swarm.

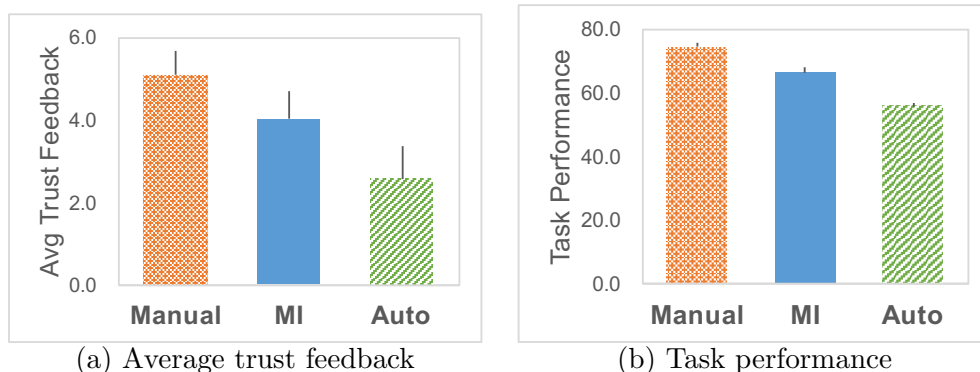


Figure 6: Results from in-process measurements (error bars are SEM). (a) The participants had a significantly lower average trust feedback values in the autonomous LOA while the manual LOA had the highest trust feedback. (b) The task performance of the manual LOA was significantly higher than the autonomous LOA.

3.4.4 User intervention commands

In [21], a classifier was developed to distinguish the human operator’s command input into *intervention* and *nonintervention* categories. It was shown that only the parts of the human commands that correct the heading direction of the swarm occur due to low levels of trust, which refers to *intervention* commands. Other *nonintervention* commands may be due to non-trust-related factors, such as swarm deployment or a switch in intentions. The

experimental data indicated a linear classifier, which uses the length of the vector drawn by the participants to distinguish a command input. Shorter lines were associated with interventions that indicate levels of dissatisfaction with swarm behavior, while longer lines were used to redirect the swarm to different search regions could indicate a change of intention instead of a loss of trust. In this work, the same classifier is used to identify user commands, because the experimental setting is identical to that of [21].

With 317 px as the threshold of the linear classifier, we compare the trust feedback when users issued different kinds of commands. In both the manual and the MI LOA (the autonomous LOA is not applicable since it does not have a command input), the two groups showed a significant difference in the trust feedback (manual: 2-tailed $t = 25.98$, $p \ll 0.001$, $df = 1080058$, MI: 2-tailed $t = 32.76$, $p \ll 0.001$, $df = 1074381$) with participants tending to give low trust feedback when they issued interventions. The average trust feedback values of the intervention and non-intervention groups were 4.789 and 5.114 in the manual LOA (3.445 and 4.040 for the MI LOA), respectively (shown as Fig. 7). This finding indicates that the participants issued intervention commands when their level of trust was lower.




Figure 7: The average trust feedback when users issued intervention or nonintervention commands. Error bars are 1 standard error from the mean.

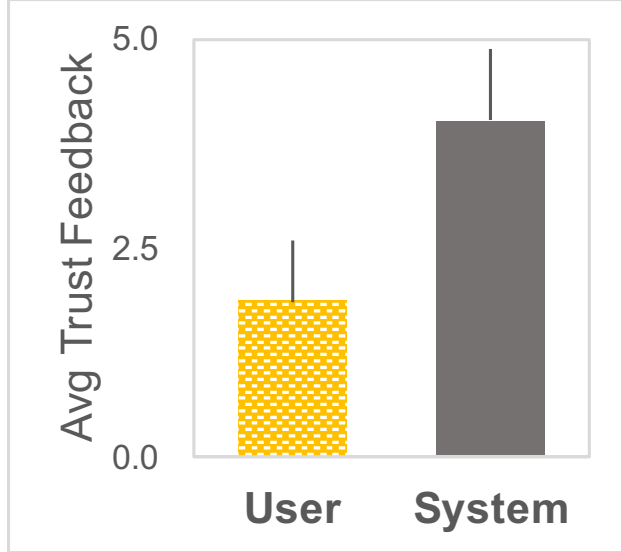


Figure 8: The average trust feedback when the mode switches occurred (user-initiated vs. system-initiated switches). Error bars are 1 standard error from the mean.

3.4.5 Mode switch in the MI LOA

Because participants were able to access both manual and autonomous search modes in the MI LOA, another important indicator of operator’s trust is the control switch between two modes. On one hand, human operators may take the control over autonomous searching algorithm when they lose trust in it, or leave the control to the algorithm when they feel confident. On the other hand, a control switch initiated by the recommendation system may also influence human trust, since it is related to the performance of the swarm. The average numbers of user-initiated and system-initiated switches in each trial were 6.500 ($\sigma = 3.138$) and 5.050 ($\sigma = 3.916$), respectively. Trust feedback values when users or the system initiated the mode switch were significantly different (2-tailed $t = -8.988$, $df = 1045$, $p \ll 0.001$). The means of trust feedback were 1.858 and 3.954 (Fig. 8), which suggest that the participants had significantly lower levels of trust when they switched the search mode themselves.

4.0 HUMAN-SWARM TRUST MODEL

In this chapter, a novel computation model of human trust in supervisory control of robotic swarms will be introduced. This model is based on a Kalman filter that estimates human trust states from temporal events that occur during the course of interaction. Both the open-loop trust anticipation and closed-loop trust feedback models are updated with processing time lags to model the overall trust dynamics. A personalized model is obtained from each individual's interaction experience to provide a customized trust prediction that adapts to an individual's behavior and attitude.

This chapter starts with a brief introduction to the implementation of the Kalman estimator and the extended assumptions made for human trust prediction settings. Next, the model's structure and chosen procedure of parameters are described in detail.

4.1 KALMAN ESTIMATOR

As previously mentioned, the modified theoretical trust model can be easily implemented as a Kalman estimator. The internal model updates the trust estimation via two different loops, as the process update and measurement update in a Kalman estimator.

It is assumed that the overall human trust state can be represented by the following linear time-variant equation:

$$x_k = Ax_{k-1} + Bu_{k-1} + w_{k-1} \quad (4.1)$$

where x_k represents the human trust state, u_{k-1} is the human's control input to the system,

and w_{k-1} refers to random disturbances. This equation reveals the open-loop trust evolution over time. The first term Ax_{k-1} is the natural evolution of the trust state, which remains the same over time if parameter A is set as a vector consisting of ones. This setting is based on the assumption that human trust will remain at the same level if no feedback from external events or changing of internal belief is received. The second term Bu_{k-1} refers to the influence of trust anticipation when operators issued commands or controlled switches. While control parameter B may be changing during each time step, we assume it as a constant for each operator and learn it from individual interaction history. The last term w_{k-1} is a vector of independent zero-mean normalized Gaussian white noises, which represents the unexpected noise that occurs in this process.

The following equation represents the human operator's perception process:

$$z_{k-\tau} = Cx_k + v_k \quad (4.2)$$

where $z_{k-\tau}$ is the observation that the human operator gets from the system display, the perception parameter C is also assumed as a constant vector and learned from each operator's interaction history, v_k is the Gaussian white noises during measurement. For simplification, we assume that the displayed swarm parameters can be easily perceived by human operators and that they have linear relationships with human trust evolution. The parameter C reveals the relationship between the trust state x and the measurement z . In addition, as mentioned in an earlier section, there is a time lag for human operators to perceive and analyze the information from the system display. Therefore, τ seconds are subtracted from the time step to indicate the later arrival of system feedback. This measurement update equation corresponds to the closed-loop feedback in trust evolution.

The random variables w_{k-1} and v_k represent the process and measurement noise, respectively. They are assumed to be independent of each other, white, and with Gaussian probability distributions of:

$$p(w) \sim N(0, Q) \quad (4.3)$$

$$p(v) \sim N(0, R) \quad (4.4)$$

We define $\hat{x}_k^- \in \mathfrak{R}^n$ to be our a priori state estimate at step k , given knowledge of the

process prior to step k , and $\hat{x}_k \in \mathfrak{R}^n$ to be our a posteriori state estimate at step k , given measurement $z_{k-\tau}$. A priori and a posteriori estimate errors can be defined as:

$$e_k^- = x_k - \hat{x}_k^-$$

$$e_k = x_k - \hat{x}_k$$

The a priori estimate error covariance is then:

$$\hat{P}_k^- = E[\hat{e}_k^- \hat{e}_k^{-T}] \quad (4.5)$$

and the a posteriori estimate error covariance is:

$$\hat{P}_k = E[\hat{e}_k \hat{e}_k^T] \quad (4.6)$$

The goal of a Kalman filter is to find an equation that estimates an a posteriori state \hat{x}_k as a linear combination of an a priori estimate \hat{x}_k^- and a weighted difference between a measurement $z_{k-\tau}$ and a measurement estimation $C\hat{x}_k^-$. This difference is called residual, which refers to any discrepancy between the predicted measurement and the actual measurement.

$$\hat{x}_k = \hat{x}_k^- + K(z_{k-\tau} - C\hat{x}_k^-) \quad (4.7)$$

The matrix K in (4.7) is chosen to be the gain that minimizes the a posteriori error covariance (4.6). One form of the resulting K that minimizes (4.6) is given as follows:

$$K_k = P_k^- C^T (C P_k^- C^T + R)^{-1} \quad (4.8)$$

To put these equations together to estimate the human trust state, a discrete Kalman filter algorithm is implemented, based on [2]. The Kalman filter estimates human trust at some time steps and then obtains feedback in the form of measurements with both noise and time lag. Thus, the equations for the Kalman filter are divided into two groups: process update equations and measurement update equations. The process update equations are responsible for projecting the current state and error covariance estimates forward to obtain the a priori estimates for the next time step, while the measurement update equations are responsible for the feedback; for example, incorporating a time-delayed measurement into

the a priori estimate (when it is available) to obtain an improved a posteriori estimate.

Kalman filter process update equations:

$$x_k = x_{k-1} + Bu_{k-1} \quad (4.9)$$

$$P_k^- = P_{k-1} + Q \quad (4.10)$$

Kalman filter measurement update equations:

$$K_k = P_k^- C^T (HP_k^- C^T + R)^{-1} \quad (4.11)$$

$$\hat{x}_k = \hat{x}_k^- + K(z_{k-\tau} - C\hat{x}_k^-) \quad (4.12)$$

$$P_k = (I - K_k C)P_k^- \quad (4.13)$$

The process update equations are a time-series function that projects the state and covariance forward from time $k - 1$ to time k . The algorithm first computes the Kalman gain K_k during the measurement update and then measures the state to obtain $z_{k-\tau}$. Notice that the measurement has a time delay τ , so the measurement at time step at $k - \tau$ will be used to generate an a posteriori state estimate at time k in (4.12). The final step is to obtain an a posteriori error covariance estimate via (4.13).

With the process and measurement updates, this process repeats at each time step by putting the a posteriori estimate at the last time step to predict the a priori estimates at the current time step. The recursive fashion enables the Kalman filter to condition the current estimate on all of the past measurements, which makes the Kalman filter's implementation much more practical and feasible than with previous methods.

4.2 MODEL PARAMETERIZATION

In the implementation of the Kalman filter, five sets of parameters need to be predefined to run the simulation. They are process and measurement noise covariance Q and R , control and perception parameters B and C , and processing time lag τ . The specific value of each parameter is based on rules, as well as a combination of heuristic and grid search using a




Figure 9: A block diagram of the implementation of a Kalman filter.

mean square error as the criterion; thus, the resulting fit might not be optimal.

4.2.1 Process and measurement noise

In the common practice of implementing Kalman filters, the measurement noise covariance R is usually measured before operating the filter, because an off-line sample measurement can help determine the variance of the measurement noise. The determination of the process noise covariance Q is typically more difficult, because we have no access to the human trust dynamic that we are estimating. A simplified but effective solution provided by previous research is that providing enough uncertainty to the process via Q .

Literature on the Kalman filter has pointed out that tuning the filter parameters Q and R can bring superior filter performance, even without a rational basis for choosing the parameters [2]. Considering the lack of quantitative work in measuring human trust dynamics and the influence of swarm physical state on human trust, the measurement noise R and the process noise Q are tuned off-line when using the training data set. The tuning

results are $R = 1$ and $Q = 1 * 10^{-3}$, and remain static throughout the following process.

4.2.2 Control parameters

Parameters B refer to the influence of temporal user control inputs that occur during the interaction with a human operator’s level of trust. For the control inputs, we track the control mode switch initiated by both the user and system and the command input issued user by the user. Therefore, B becomes a parameter vector with three values that correspond to three types of input.

Control takeovers between the automation and human indicate the changing trust state of human operator. When humans decide to switch the control from automation to manual search mode, their trust towards swarm robots’ autonomous searching algorithms are more likely to drop because they have a higher self-confidence in their manual control skills. On the other hand, the reverse control switch from manual search to autonomous search indicates an improvement in levels of human trust in automation. The above assumptions are supported by the observational results presented in Section 3.4, which showed that participants hold a significantly lower trust level when they initiated a control switch. Therefore, the contribution of a user-initiated switch should be negative when the control switch direction is from manual to auto, and positive in the opposite direction. When considering that the trust rating ranges between 0 and 1, we set the searching range of control parameter b_1 from 0 to 0.1. The process of setting the parameters of a system-initiated control switch is similar to setting the parameters of a user-initiated switch. However, the influence degree may be of different significance for different individuals, so a separate searching process is run for b_2 .

Based on the findings of the supervisory control of robotic swarms in [20, 21], users’ command inputs generally fall into one of two categories: *intervention commands* that occur due to a loss of trust and *nonintervention commands* that occur due to other reasons, like changing an intention. Here, only the *intervention commands* issued by operators are used to predict human trust. Because intervention commands indicate a loss of trust, the range of the corresponding parameter b_3 is then set between -0.1 and 0.

4.2.3 Perception parameters

Parameters C represent the influences of perceived swarm states on levels of human trust. According to previous studies [20, 38], both robots' task performance and physical states have significant influence on human trust formation and updates to its levels. However, when applying the model to human-swarm interaction scenarios, the overall performance is shown to be less informative (see Section 3.4). Therefore, the performance increment in each second is then adopted to predict human trust levels. The corresponding parameter c_1 is set as positive with a range between 0 and 0.1, because better performance usually leads to greater levels of trust. When considering the unique physical characteristics of swarms [20, 21], the heading variance and convex hull area are chosen as the indexes that are most influential to human trust levels. The heading variance is the deviation of a given swarm member's heading direction, which reveals how coherently the swarm is moving. The convex hull area equals to the area covered by connected swarm members, which refers to how concentrated the swarm team is distributed in a given area. As shown in Section 3.4, the impact of swarm physical parameters on human trust are different between LOAs, so the parameter ranges of c_2 and c_3 are set from -0.1 to 0.1.

4.2.4 Time delay

Time delay τ represents the processing time for a human operator to correctly perceive the swarm's state and performance. Since similar Kalman filter structure has only been used manual control task, published data of time delay only considers the neuromotor time, which was reported as 0.15-0.25 seconds [12]. However, a previous study on neglect benevolence in human swarm interaction [19] shows that the time for a closed-loop feedback from user input to swarm behavior convergence ranges around several seconds. Considering that the total length of a experiment trial is 5 minutes, we searched τ from 0 s to 30 s.

4.2.5 Grid search

Because each individual has unique trust criteria and anticipation due to their previous experiences, a personalized model is needed to reach a better overall level of fitness. Therefore, a grid search process is employed to find the best parameter combination for each participant that uses root mean square error as the metric. Grid search is a commonly used method of hyperparameter selection for a given model with clearly defined criteria. When considering the searching space and parameter number in our experiment settings, we adopted a modified version of the searching strategy. We employed this method to search the best control parameter B , perception parameters C , and time lag τ for each participant. Based on the estimated ranges introduced above, we equally sampled several values for each parameter and tried out all possibilities for one single parameter once, with other parameters fixed at the median value. The sampling step length is 0.01 for B and C and 1 for τ . The assumption of applying a simplified grid search is based on the following reasons: 1) B and C are each a vector with three elements, and the searching space of τ ranges from 0 - 30s, so exclusively trying out all the combinations is computationally inefficient. 2) The Kalman filter is a nearly linear system and the influence of each term and corresponding parameter is addable. 3) During the course of interaction, significant events (e.g. control switch or intervention command) occur in a temporal sequence with considerable intervals. Therefore, we assume that parameters do not have interaction effects with each other and only change one of them during the search task.

5.0 TRUST MODEL EVALUATION

Data collected from the human experiment was used to evaluate the proposed Kalman filter model. In the experiment, each participant interacted with swarms in three different LOAs, and in each LOA, they finished three 5-minute trials. During the interaction, human inputs (user-initiated control switch, system-initiated control switch, and control commands) and swarm states (heading variant, convex hull area, and task performance) were recorded with a sampling rate of 60 Hz. To reveal the unique cognitive process of different individuals in varied LOAs, separate models were trained for each participant in each LOA.

5.1 DATA PROCESS

First, the sampling rate of original data is 60 Hz, which is relatively high for a human cognitive model. In order to better capture the temporal relationship of interaction events and human trust, we compute the average value of recorded data during each second to decrease the data sampling rate to one data point per second. Second, since data inputs are at different scales and are unable to be normalized into a shared space, only the count number or sign of incremental value is kept for further model computation. For control inputs, we calculate the total number of events occurring during each second and use the count number in our model (e.g. the number of user-initiated switches or the number of intervention commands). For swarm states, we calculate the difference of physical parameters or performance in each second and take the sign of it for the model. For example, 1 represents a increase of heading variant during a certain time window (1 second in our case) and -1 represents a decrease of the value, while 0 represents the value remaining the same. Third,

in order to compare the model performance with an existing model, we adopted the same scale of depended variable, the reported level of trust. The original $[-10,10]$ scale is linearly transformed to $[0,1]$.

Finally, the 5-minute experiment log of each trial was converted into three input matrices: x , with a dimension of $(300,1)$, consists of the average trust level during each second; u , with a dimension of $(300,3)$, consists of the count number of three types of user input; z , with a dimension of $(300,3)$, consists of the changing directions of swarm physical parameters and overall performance.

5.2 MODEL TRAINING

To train the personalized model for each participant in each LOA, the data of first two trials was used as training data, and the model was validated with the data of the last trial. A modified grid search was used to learn the parameters $B, C,$ and τ . For each simulation, a combination of parameters was given and the root mean square error of model prediction on the training data set was reported. The searching algorithm compares all possibilities of parameters within the given search space and store the combination with least RMSE. The learning results of two trials were combined by computing the arithmetic mean of each parameter. In addition, the mean value of reported trust level in the first two trials was calculated to be used as the a priori knowledge about each participants' preference.

5.3 COMPARISON WITH EXISTING MODEL

In this section, we evaluate our Kalman filter model on the experimental data of 20 participants in three different LOAs. The inverse reinforcement learning model [20] and dynamic Bayesian network model [38] were chosen as the baseline, to compare our model with state of the art principles. The model implementation was conducted on a system with an AMD Ryzen 1600X 3.80 Ghz chip with 16G RAM and Python 3.7. Fig. 10 shows three instances

of the predicted result on the test data set. Each graph shows the predicted result of one participant and the x and y axis represent the time step and the trust level, respectively. The blue line represents model prediction outputs, while the green line shows user trust feedback values. The average root mean square error between the predicted values and user feedback values are reported in Table 1. The mean RMSE of Kalman Filter model in MI, Manual and Auto LOAs are 0.107, 0.101, 0.116, respectively. Because our human experiment has the exactly same setting as in [20], the performance metrics (RMSE) can be directly compared to one another. Based on the model prediction performance data reported in [20], the RMSE of our Kalman filter model is significantly lower than that of the IRL and DBN models in all three LOAs ($ps < 0.001$).

Table 1: The result of trust predictions (RMSE). The Kalman filter model is compared with the IRL model proposed in [20] and DBN model in [38]). Fig. 10 shows the example prediction results in graphs of three individual participants.

Model	MI	Manual	Auto
Kalman Filter	0.107(0.047)	0.101(0.075)	0.116(0.059)
IRL	0.148(0.075)	0.159(0.108)	0.174(0.071)
DBN	0.233 (0.107)	0.264 (0.132)	0.245 (0.083)

5.4 DISCUSSION

When evaluating the data, our Kalman filter model had a better performance, as compared to existing methods such as IRL and DBN models. There are several reasons for the improvement of trust prediction.

Our model is the first to consider human perception time lag in trust prediction. The trained models are shown to have a time delay of around 10 seconds, which indicates a huge gap between a user issuing commands and finally perceiving any feedback. The time delay




Figure 10: Example predicted results on the test data set. Each graph shows the prediction result of one participant and the x and y axis represent the time step and the trust level, respectively. The blue line represent model prediction while green is user trust feedback values.

is caused by both the non-linear dynamics of swarm systems and the characteristics of the evolution of human trust levels. Since the trust we are tracking is a user-reported value, it may take time for the internal trust increment to accumulate to exceed a certain threshold that triggers a trust report.

The IRL model reduces the state space of MDP by encoding the swarm physical parameters into several categories. In the best-performing model (model 2) in [20], only two states of heading variance and five states of convex hull area were considered. The reduced state space may oversimplify the problem and miss important information in the dynamic features of swarm parameters. On the other hand, the assumption of MDP is that the decision at a given time is based on both the current and next transition states. This mechanism decreases the ability of the IRL method to track the long-term temporal sequences of events that occur during the interaction; for example, the accumulation of evidence in trust evolution. In the Kalman filter, user intervention, control switches, and a swarm’s physical parameters are included in the model, which provides a relatively larger state space. The iterative features of the Kalman filter enables us to better track the features of a time-series event log.

The reason for the DBN model’s poor performance can be explained by different experimental designs and data frameworks. The DBN model in [38] uses both the absolute and relative trust feedback as the input. However, our data does not have relative feedback,

so the DBN model cannot use this important piece of information. Also, we did not allow the DBN model to use absolute trust feedback in testing for a fair comparison, because the other two models directly predict those values. Last, the occurrence of interventions, which is important evidence of trust in the DBN model, is less frequent in our swarm experiments when compared with Xu’s experimental setting.

The Kalman Filter model does not show significantly different performance between LOAs. Three LOAs have different input channels that user interventions are not valid in the autonomous LOA and control switches are only valid in the mixed initiative LOA. Thus, the only shared channel across three LOAs is the physical parameters of swarm. The above fact leads to a conclusion that the swarms’ appearance is the most influential factor towards human trust, which aligns with our previous findings [20].

6.0 CONCLUSIONS

This work has developed a novel method of modeling human trust in a swarm control task where humans are not readily able to perceive swarm states and overall task performance. The proposed implementation method is based on a modified theoretical trust model, in which the trust evolution process is framed as "trust-action-verify." In this model, the human operator will first make an open-loop anticipation about the change in their trust state based on the intervention action that has been made, then wait until they receive the displayed feedback from the system to verify the initial anticipation. Thus, a trust update consists of both an open-loop trust anticipation and a closed-loop trust feedback. A Kalman filter is used to implement this process, which considers the trust anticipation as process update equations and the perceived system feedback as the measurement update equations with a time delay. A personalized model was created for individual operators to reveal different preferences and previous experience.

In summary, the proposed model is novel in the following aspects: 1) This Kalman estimator is the first to model the complete trust evolution process with both closed-loop feedback and open-loop trust anticipation. 2) The proposed model analyzes time-series data to reveal the influence of events that occur during the course of an interaction; namely, a users intervention and report of levels of trust. 3) The proposed model considers the operators cognitive time lag between perceiving and processing the system display. 4) The proposed model uses the Kalman filter structure to fuse information from different sources to estimate a human operators mental states. 5) The proposed model provides a personalized model for each individual.

There are also several limitations in the current work, which can be addressed in future studies. The Kalman filter assumes that there is a linear relationship between control in-

put, measurement, and human trust; but when considering the limited knowledge in such processes, more work needs to be done to provide a convincing basis for this method. An extended Kalman filter may need to be employed to represent a non-linear relationship.

The parameter searching is merely a random grid search without gradient learning; therefore, the search result is not guaranteed to be optimal, even within the searching range. To address this issue, advanced learning techniques may be employed to reach a better fit.

A Kalman filter is directly used to predict the user's reported trust level. However, the reported trust level is only a measurement of the internal human trust state. As a result, it may take time for trust state increments to accumulate until it exceeds a certain threshold that triggers the trust report. A more reasonable way to construct the model should enable it to track a latent state of human trust and to consider the influence of external factors (e.g. user intervention) and the observation method (e.g. trust report).

As for future research directions, multiple options are available based on the findings and model structure proposed in this work.

First, with this online computational trust model, human trust levels can be predicted given the system states and user inputs. We envision the implementation of an adaptive robotic system that is capable of sensing human trust levels and adapting to the human accordingly. For example, when the system detects a drop in human trust levels, it may increase the system's transparency by providing explanations of its own behaviors, or hand over control to the operator if necessary. Real-time detection of human cognitive states is the foundation of such adaptive systems that can optimize joint performance across dynamic scenarios. On the other hand, the adaptive behavior of a swarm could help human operators to calibrate their trust to an appropriate level. For example, when the operator shows an unrealistic expectation for existing levels of automation, the system may alert the human operator to be prepared for automation failures and control takeovers.

Second, besides trust, inferring intent is another promising direction in building adaptive systems. The command classifier used in this thesis is a simplified method to distinguish users' intent from intervention and nonintervention (e.g. redirection). More work is needed to consider other metrics during the interaction (e.g. physiological measurements) to identify human intent. Human intent modeling is especially important in human-agent teaming

scenarios. Robots could have an optimized schedule that considers the changing intents of human during the collaboration.

Third, we found in our experiment that human trust was more influenced by swarm's physical parameters (appearance) than by the number of targets found (performance) Those findings could be attributed to the unintelligibility of swarm behaviors, due to the unique physical characteristics of swarms and the complex interactions that occur among robots. Previous research in human-robot interaction has shown that transparency can lead to an improved calibration of trust levels [36, 28]. Thus, a better design for the interfaces of human-swarm interaction is needed to to better communicate the system state and smooth the control input.

Finally, the trained models are shown to have a huge time delay (10 seconds) between a user issuing commands and finally perceiving the feedback. The composition of this delay is quite complicated and requires further research. For example, the delay contains both the time that the swarm team needs to respond to the user's input and the time that humans need to perceive the changing of swarm states. Since swarm robots have special control laws and communication constraints, their response time is relatively long, and the unique physical characteristics of swarms make their behaviors and states less intelligible to human operators, as compared to single-robot systems. As a result, operators need more time to perceive and process the displayed information from swarms. Finally, the trust we are tracking is a user-reported value of trust, which is a measurement of the internal trust state. It may take time for the trust state increment to accumulate to a level high enough to exceed a certain threshold that triggers the trust report. It would be interesting to conduct an in-depth investigation of this whole process in-depth and separate those compositions.

7.0 BIBLIOGRAPHY

- [1] F Amato et al. “Trust Observations in Validation Exercises”. In: *Proc. of Int. Conf. on Secure Software Integration and Reliability Improvement*. 2011, pp. 216–223.
- [2] Gary Bishop, Greg Welch, et al. “An introduction to the kalman filter”. In: *Proc of SIGGRAPH, Course 8.27599-23175* (2001), p. 41.
- [3] Francesco Bullo, Jorge Cortes, and Sonia Martinez. *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*. Vol. 27. Princeton University Press, 2009.
- [4] Zack J Butler, Alfred A Rizzi, and Ralph L Hollis. “Contact sensor-based coverage of rectilinear environments”. In: *Proceedings of the IEEE International Symposium on Intelligent Control/Intelligent Systems and Semiotics*. IEEE. 1999, pp. 266–271.
- [5] Shih-Yi Chien et al. “An Empirical Model of Cultural Factors on Trust in Automation”. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. Vol. 58. 2014, pp. 859–863.
- [6] Miguel Duarte et al. “Application of swarm robotics systems to marine environmental monitoring”. In: *OCEANS 2016-Shanghai*. IEEE. 2016, pp. 1–8.
- [7] Mica R Endsley and Esin O Kiris. “The out-of-the-loop performance problem and level of control in automation”. In: *Human Factors* 37.2 (1995), pp. 381–394.
- [8] Ji Gao and John D Lee. “Extending the decision field theory to model operators’ reliance on automation in supervisory control situations”. In: *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 36.5 (2006), pp. 943–959.
- [9] SG Hart and LE Staveland. “Development of a multi-dimensional workload rating scale: Results of empirical and theoretical research. In, PA Hancock and N. Meshkati”. In: *Human Mental Workload* (1988).

- [10] Makoto Itoh, Genya Abe, and Kenji Tanaka. “Trust in and use of automation: their dependence on occurrence patterns of malfunctions”. In: *IEEE SMC’99 Conference Proceedings. 1999 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No. 99CH37028)*. Vol. 3. IEEE. 1999, pp. 715–720.
- [11] Mohammad R Jahanshahi et al. “Reconfigurable swarm robots for structural health monitoring: a brief review”. In: *International Journal of Intelligent Robotics and Applications* 1.3 (2017), pp. 287–305.
- [12] D Kleinman, Sheldon Baron, and W Levison. “A control theoretic approach to manned-vehicle systems analysis”. In: *IEEE Transactions on Automatic Control* 16.6 (1971), pp. 824–832.
- [13] John D Lee and Neville Moray. “Trust, self-confidence, and operators’ adaptation to automation”. In: *International journal of human-computer studies* 40.1 (1994), pp. 153–184.
- [14] John D Lee and Katrina A See. “Trust in automation: Designing for appropriate reliance”. In: *Human factors* 46.1 (2004), pp. 50–80.
- [15] John Lee and Neville Moray. “Trust, control strategies and allocation of function in human-machine systems”. In: *Ergonomics* 35.10 (1992), pp. 1243–1270.
- [16] Michael Lewis, Katia Sycara, and Phillip Walker. “The role of trust in human-robot interaction”. In: *Foundations of Trusted Autonomy*. Springer, Cham, 2018, pp. 135–159.
- [17] Huao Li et al. “Human Interaction Through an Optimal Sequencer to Control Robotic Swarms”. In: *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE. 2018, pp. 3807–3812.
- [18] Bonnie M Muir and Neville Moray. “Trust in automation. Part II. Experimental studies of trust and human intervention in a process control simulation”. In: *Ergonomics* 39.3 (1996), pp. 429–460.
- [19] Sasanka Nagavalli et al. “Bounds of neglect benevolence in input timing for human interaction with robotic swarms”. In: *Proceedings of the tenth annual acm/ieee international conference on human-robot interaction*. ACM. 2015, pp. 197–204.
- [20] Changjoo Nam et al. “Models of Trust in Human Control of Swarms With Varied Levels of Autonomy”. In: *IEEE Transactions on Human-Machine Systems* (2019).

- [21] C. Nam et al. “Predicting trust in human control of swarms via inverse reinforcement learning”. In: *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. Aug. 2017, pp. 528–533. DOI: [10.1109/ROMAN.2017.8172353](https://doi.org/10.1109/ROMAN.2017.8172353).
- [22] C. Nam et al. “Trust of Humans in Supervisory Control of Swarm Robots with Varied Levels of Autonomy”. In: *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. Oct. 2018, pp. 825–830. DOI: [10.1109/SMC.2018.00148](https://doi.org/10.1109/SMC.2018.00148).
- [23] Raja Parasuraman and Victor Riley. “Humans and automation: Use, misuse, disuse, abuse”. In: *Human factors* 39.2 (1997), pp. 230–253.
- [24] Raja Parasuraman, Thomas B Sheridan, and Christopher D Wickens. “A model for types and levels of human interaction with automation”. In: *IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans* 30.3 (2000), pp. 286–297.
- [25] Jacques Penders et al. “A robot swarm assisting a human fire-fighter”. In: *Advanced Robotics* 25.1-2 (2011), pp. 93–117.
- [26] C.W. Reynolds. “Flocks, herds and schools: A distributed behavioral model”. In: *ACM SIGGRAPH Computer Graphics*. Vol. 21. ACM. 1987, pp. 25–34.
- [27] Heath A Ruff, Sundaram Narayanan, and Mark H Draper. “Human interaction with levels of automation and decision-aid fidelity in the supervisory control of multiple simulated unmanned air vehicles”. In: *Presence: Teleoperators and Virtual Environments* 11.4 (2002), pp. 335–351.
- [28] Tracy L Sanders et al. “The influence of modality and transparency on trust in human-robot interaction”. In: *2014 IEEE International Inter-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*. IEEE. 2014, pp. 156–159.
- [29] Thomas B Sheridan. “Extending Three Existing Models to Analysis of Trust in Automation: Signal Detection, Statistical Parameter Estimation, and Model-Based Control”. In: *Human factors* (2019), p. 0018720819829951.
- [30] Thomas B Sheridan. *Modeling Human–System Interaction: Philosophical and Methodological Considerations, with Examples*. John Wiley & Sons, 2016.
- [31] Thomas B Sheridan. *Telerobotics, automation, and human supervisory control*. MIT press, 1992.

- [32] Walt Truszkowski et al. “NASA’s swarm missions: The challenge of building autonomous software”. In: *IT professional* 6.5 (2004), pp. 47–52.
- [33] Marialena Vagia, Aksel A Transeth, and Sigurd A Fjerdings. “A literature review on the levels of automation during the years. What are the different taxonomies that have been proposed?” In: *Applied ergonomics* 53 (2016), pp. 190–202.
- [34] Phillip Walker. *CUDA Swarm Simulator*. <https://github.com/pmwalk/cuSwarm>. 2015–2017.
- [35] Phillip Walker, Michael Lewis, and Katia Sycara. “Characterizing human perception of emergent swarm behaviors”. In: *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE. 2016, pp. 002436–002441.
- [36] Robert H Wortham and Andreas Theodorou. “Robot transparency, trust and utility”. In: *Connection Science* 29.3 (2017), pp. 242–248.
- [37] Anqi Xu. “Efficient Collaboration with Trust-Seeking Robots”. An optional note. PhD thesis. Montreal, Quebec, Canada: McGill University, Oct. 2016.
- [38] Anqi Xu and Gregory Dudek. “OPTIMo: Online Probabilistic Trust Inference Model for Asymmetric Human-Robot Collaborations”. In: *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. HRI ’15. Portland, Oregon, USA: ACM, 2015, pp. 221–228. ISBN: 978-1-4503-2883-8. DOI: [10.1145/2696454.2696492](https://doi.org/10.1145/2696454.2696492). URL: <http://doi.acm.org/10.1145/2696454.2696492>.