

## Steps towards Satisficing Distributed Dynamic Team Trust

Edmund R. Hunt<sup>1</sup>, Chris Baber<sup>2</sup>, Mehdi Sobhani<sup>1</sup>, Sanja Milivojevic<sup>3</sup>, Sagir Yusuf<sup>2</sup>,  
Mirco Musolesi<sup>4</sup>, Patrick Waterson<sup>5</sup>, and Sally Maynard<sup>5</sup>

<sup>1</sup>School of Engineering Mathematics and Technology, University of Bristol

<sup>2</sup>School of Computer Science, University of Birmingham

<sup>3</sup>Bristol Digital Futures Institute, University of Bristol

<sup>4</sup>Department of Computer Science, University College London

<sup>5</sup>Human Factors and Complex Systems Group, Loughborough University

edmund.hunt@bristol.ac.uk, c.baber@bham.ac.uk, mehdi.sobhani@bristol.ac.uk, sanja.milivojevic@bristol.ac.uk,  
s.m.yusuf@bham.ac.uk, m.musolesi@ucl.ac.uk, p.waterson@lboro.ac.uk, s.e.maynard@lboro.ac.uk

### Abstract

Defining and measuring trust in dynamic, multiagent teams is important in a range of contexts, particularly in defense and security domains. Team members should be trusted to work towards agreed goals and in accordance with shared values. In this paper, our concern is with the definition of goals and values such that it is possible to define ‘trust’ in a way that is interpretable, and hence usable, by both humans and robots. We argue that the outcome of team activity can be considered in terms of ‘goal’, ‘individual/team values’, and ‘legal principles’. We question whether alignment is possible at the level of ‘individual/team values’, or only at the ‘goal’ and ‘legal principles’ levels. We argue for a set of metrics to define trust in human-robot teams that are interpretable by human or robot team members, and consider an experiment that could demonstrate the notion of ‘satisficing trust’ over the course of a simulated mission.

### Introduction

Defining the interpersonal and technical factors that relate to trust in human-AI/robot teaming is an open problem in the research community (Huang et al. 2021). In particular a key problem is the definition of ‘trust’ in these scenarios. We expect trust to vary according to the developing situation faced by each teammate. Thus, obtaining a trust level sufficient for a given situation will always involve satisficing, i.e. obtaining a minimally acceptable level to progress a team’s mission. We use the metaphor of a *Ladder of Trust*, whereby teammates may ‘climb’ down or up the ladder to satisfy situation-dependent requirements (Baber et al. 2023).

To study trust in human-robot teams, it is necessary to define the concept of trust in a manner which is meaningful for both humans and robots. Previously, trust has been seen as a multi-dimensional concept that focuses on human perceptions (McAllister 1995; Mayer, Davis, and Schoorman 1995), e.g., through self-report questionnaires for the human team members (Jian, Bisantz, and Drury 2000; Schaefer 2016; Malle and Ullman 2021), but this assumes cognitive and cultural capabilities beyond those of robots. While humanlike robots can be perceived to be capable of making ‘moral’ decisions (Malle and Ullman 2021), in general,

people do not accept the idea of machines making moral decisions (Bigman and Gray 2018). We believe that there has been less attention given to trust held by robots of their human team mates. In our project, we seek metrics appropriate for both humans and robots to quantify, and hence regulate, their trust in their teammates.

### Principles, Values and Goals

We see a hierarchy of *principles*, *values* and *goals* (Fig. 1). For human-robot teams, trust requires team members to behave consistently with the team’s goals, team values and legal principles to a standard of performance agreed by the team (Baber et al. 2023). Goals are pre-defined ends, aims or objectives for humans or robots in a team. A team goal, for example, could be to defuse a bomb. Values, on the other hand, have a sense of purpose. They are ‘*those ends deemed worth pursuing*’ (Williams 2009, p.559). Values define desirable outputs: what motivates us, what is worth striving for. As such, they can be individual, team, and social. Social values, such as liberty, freedom, justice, democracy, and respect for fundamental rights, are accepted by society and are aspirational. Legal principles are ‘norms laying down essential elements of a legal order’ (van Bogdandy, cited in Williams 2009, p.559). Following Habermas, we consider legal principles to be general propositions from which norms arise, ‘*certain standards that might be based in law or practice, which contribute to forming a framework for decision-making and action*’ (Williams 2009, p.559). Acting according to legal principles, in our example, could be to defuse the bomb without a loss of human life. Legal principles have a sense of obligation attached to them; they set the bounds or constraints within which activities are permitted. Values could fill the gap where legal principles fail to provide guidance. For example, acting according to social values would be to defuse a bomb even if that means a loss of human life, where such action saves the lives of many others. Acting according to *team* values would be defusing a bomb with a loss of human life, but with preservation of the capability of the human-robot team to continue its mission. Values are moral in character. They are what is the best for the individual, the team, or society. Principles command; values recommend; goals direct. Principles are binary (valid or invalid), values are not. As such, values can be considered in

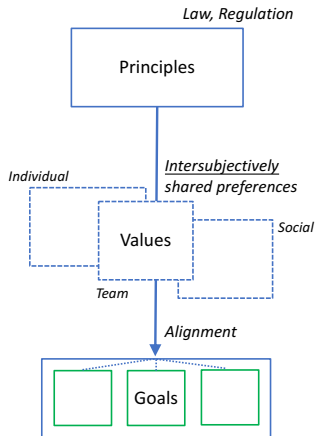


Figure 1: Alignment of goals and principles between humans and robots is possible; expressing ‘values’ in contextualized goal priorities remains the province of human actors.

terms of a trade-off in which individual, team and societal values influence the choice of goal and definition of acceptable outcome in a specific situation.

### Values in Human-Robot Teams

To speak of ‘values’ in a human-robot team might imply that these can be defined with sufficient clarity that they can be quantified. Indeed, the Artificial Intelligence (AI) literature has discussed ‘value alignment’ (Gabriel and Ghazavi 2021), through which the ‘values’ that an AI system pursues match those of the human stakeholders affected by the system (Moor 2006). An early example of explicit ethics in AI systems can be seen in the ‘ethical governor’ (Arkin, Ulam, and Duncan 2009). This was proposed for a weapons targeting system in which selection of target could be overruled if there were rules which the decision violated. In this instance, the ‘rules’ could be defined in a similar manner to our notion of legal principles. We might be able to codify legal principles, as legal norms setting the essential elements of the legal order. But ‘values’ are, as Habermas (cited in Williams 2009) suggests, *intersubjectively shared preferences*. The concepts of ‘value alignment’ and ‘explicit ethics’ in AI seem to rest on an assumption that it is possible to codify ‘social values’ with sufficient clarity that an action might be evaluated. We are skeptical that this might be possible. We believe that alignment of goals and principles (between humans and robots in a system) is possible, but that ‘social values’ must remain the province of the human actors who are either directly participating in the system or acting as overseers, regulators, or other stakeholders. Our emphasis on context could make such an assumption difficult to articulate (because the ‘values’ might vary with context). We note that approaches that use, for example, an ‘ethical consequence engine’ (Winfield, Blum, and Liu 2014), might deal with context through the simulation of the outcome of an action. However, this still implies the codification of ‘values’ within a normative moral framework. Nonetheless, while our position is that ‘values’ remain in the human domain,

we accept that it might be possible for engineered systems to seek ‘goal alignment’ (i.e. between humans and robots).

The interpretation of a goal (in terms of the desirability of its outcome) will depend on the context in which it is performed. From this perspective, we might define the problem as one of inferring goals from the activity of other teammates as a problem of Inverse Reinforcement Learning (Ng and Russell 2000). In a recent study, the tension between selfish and social choices (as a basic version of moral problems in social dilemma games) was explored using Reinforcement Learning (Tennant, Hailes, and Musolesi 2023). In this study, Reinforcement Learning agents are provided with intrinsic rewards that reflect different views of ethics (i.e., utilitarian, deontological, consequentialist) and play a variety of iterated social dilemma games (i.e., Prisoner’s Dilemma, Volunteer’s Dilemma, Stag Hunt). Within a game, an agent seeks to respond to the state of the game by performing an action that will maximize both the game reward and an intrinsic *moral* reward. Pitching agents with different moral stances against each other revealed systematic differences in strategy, particularly in terms of cooperative or exploitative activity. Alternative strategies or stances can be considered in terms of counterfactual analysis, which is a fundamental element of causality analysis (Pearl 2009; Pearl and Mackenzie 2018), and can be used to understand which changes to a particular data model would change the model decision. Counterfactual analysis in a multi-agent reinforcement learning environment (Forney, Pearl, and Bareinboim 2017) could support performance in mission-critical environments through the ability to review alternative courses of action. We regard goals as explicit statements of intent that team members can choose to pursue, and also recognize their pursuit by others. Thus, if an actor (human or robot) sees a teammate performing a sequence of tasks, it might be reasonable for it to assume that this sequence is directed toward achieving a specific goal, and that an alternative goal might be more desirable (to achieve particular team or social values) in that context. This area has been of great interest for the community in the recent years (Jaques et al. 2019; Conitzer and Oesterheld 2023).

### Trust Specification

To define ‘trust’, we follow Lewis and Marsh (2022) in claiming that a minimal specification of trust involves:

1. *Capability*, i.e., is the teammate most appropriate for a given task in that situation?
2. *Predictability*, i.e., is the teammate acting in a way that fits the team goals and set principles, and is appropriate to its situation?
3. *Integrity*, i.e., is the teammate acting to support the team and the set principles?

In order to define metrics for the three-element view of trust outlined above, we consider what could be sensed or perceived when humans and robots interact in a collaborative task. The literature on trust generally refers to dyadic relationships between a ‘trustor’ and ‘trustee’ (e.g. Hurley 2006; Kim, Dirks, and Cooper 2009) and this can be represented as a network of directed edges, where each edge

represents one of the three elements (Figure 2). We do not suppose that the three elements can be measured with equal certainty (some aspects might need to be inferred rather than perceived), but each agent will track the relative increase or decrease of these elements over time to adjust their ‘trust’ in a teammate.

Recent work has introduced the terms system-wide trust and component-level trust (Walliser, de Visser, and Shaw 2023). These map well to our view of distributed, dynamic team trust, where we would clarify component-level trust as comprising the three elements of estimated capability, predictability and integrity, as contextualized by the agents’ situation. As Huang et al. (2021) note, it is important to understand the context in which human-AI-robot teaming occurs “including the tasks, environment, the stakeholders, and artificial agents involved... [as well as]... the kinds of interactions that are available between the entities involved in a situation context, where transitive properties of trust take place.” (Huang et al., p.307). From this, Huang et al. (2021) propose a 4-step process that involves identification of context and stakeholders and defining and measuring trust relationships in the trust network. In our approach, we define a mission using Cognitive Work Analysis (CWA) (Rasmussen, Pejtersen, and Goodstein (1994); Vicente (1999); Jenkins, Stanton, and Walker (2009)).

From the Work Domain Analysis phase of CWA, a mission can be decomposed into an Abstraction Hierarchy that (reading from top-to-bottom) describes the *purpose* of the system and (reading from bottom-to-top) describes the *activity* of the system. The claim is that any system is intended to achieve an outcome (or set of outcomes) that can be evaluated in terms of desirable consequences. Such consequences reflect the values of the stakeholders working with and affected/impacted by the system and can serve as constraints on the goals that the system is seeking to achieve. Goals could be defined by more than one ‘value’, and the values might conflict. Where there is conflict, this either requires negotiation between teammates or intervention by a higher authority. Legal principles could, to some extent, define the constraints within which a mission is performed, but these will need to be filtered through values appropriate to the situation.

With our metaphor of a ‘ladder of trust’ (Baber et al. 2023), it is these sub-component-level trust metrics (e.g.  $C_{12}$ ,  $P_{12}$ ,  $I_{12}$ , indicating Agent 1’s estimation of capability, predictability, integrity of Agent 2 in Fig. 2) that will change as more information is gathered from their interactions. To satisfice trust, then, means for the necessary combination of component (agent) level trust estimates, each comprising these three elements, to reach a certain threshold before a mission can proceed effectively. In practical terms, this might mean Agent 1 improving its estimate of Agent 2’s integrity ( $I_{12}$ , an element of  $T_{12}$ ) before it believes that Agent 2 will carry out a certain task within the mission in an acceptable manner. As a working hypothesis, we assume that system-wide trust is limited by its weakest component. In cases where an action only requires a subset of team members to be carried out, deficiencies in system-wide trust need not hold back progress, but instead the demand is to satisfice



Figure 2: Agent 1’s trust in Agent 2,  $T_{12}$ , depends on its running estimates of Agent 2’s capability, predictability and integrity; these are contextualized by the local situation.



Figure 3: A ‘Leo Rover’ robot equipped with 2D LiDAR and built-in front facing camera, which can scan ARTags.

the relevant combination of component-level trust between the situated team members.

### Motivating Use-Case

We have designed an environment in which the concept developed in this paper can be experimented upon. Human participants work with wheeled rover robots (Figure 3) to collaborate on search tasks in an environment. In the environment, Augmented Reality (AR) Tags are positioned at approximately 1 m height, for the human to scan, and also at approximately 20 cm height, for the robots to scan (Figure 3). The tags provide information about the environment (whether an area is safe or hazardous to the human) and status of objects in the environment (e.g., whether these are operational, risky, or require repair). For instance, the tags could indicate whether a location has low or high levels of radiation or could indicate whether there is a suspect package that needs to be investigated. The team might comprise two robots and one human in the field and one human working as a coordinator (remotely). Together, the robots and humans need to coordinate a search mission in an experimental environment. As the mission is performed, data can be collected from the human (using a computer tablet interface) or from sensors on the robot. The data from human and robot actors follow the same structure, e.g., {agent ID, time, location, object, action...}. All the data and communication between team members (robot(s) and human(s)) are relayed through a Django server and recorded in a database. For instance, the mobile application (‘app’) used as an interface for communicating to the robot is connected to the server so that any commands from a human to a robot will be recorded

along with other relevant information like time. The robot messages to human teammates also goes through the server before appearing on the app. The location of all team members are also being published on the server to record the tracking data before being passed to the interface.

## Defining Metrics for Trust

In order to interpret a sequence of tasks and their relation to a goal, it is necessary to define metrics through which observation of tasks and inference of goals can be performed by team members. Table 1 indicates the variables that we seek to capture and whether these can be directly measured or inferred. In Table 1, entries that are in plain text can be directly obtained from the agents, their activity, or the database recorded during the experiment; entries in italic text can be inferred from the mission plan, i.e., there will be goals that are achieved by performing combinations of tasks. Finally, the ‘why’ will be derived from the competing values of individuals, the team and society.

Having outlined a set of metrics (Table 1), we can relate these to the three-element model as follows:

- Capability (of an Agent) is a function of WHAT (task, object) and WHO (goals). We assume that a given agent will perceive the affordances of the environment in terms of their own and others’ ability to perform a task on an object in pursuit of a goal.
- Predictability (of an Action) is a function of Capability and Context (i.e., WHERE and WHEN). We assume that an agent will perceive a teammate in a context and infer the likelihood of success of an agent (with a perceived capability) achieving an outcome.
- Integrity (of an Action) is a function of Predictability and WHY. For *robot actors*, an action is interpreted by a robot and any onlooking robot teammates in terms of the likelihood of success relative to the individual/team *goal priorities* and compliance with codified principles. For *human actors*, an action is also interpreted by the human and any onlooking human teammates in relation to fulfilling relevant goal priorities in compliance with legal principles; *and also* performance criteria, constraints, and determinants of acceptable outcomes, *as influenced by individual, team and social values*.

Agents will begin a mission with initial teammate-specific estimates for each of these elements, which could be based on factors such as prior experiences with the teammates over the longer-term (e.g. de Visser et al., 2020), or proxies for trustworthiness (e.g. Lewis and Marsh, 2022).

Notice that we are claiming that integrity dynamics are linked to each distinct, situated action rather than an agent’s identity. This is a very different perspective to conventional definitions of integrity; in our view, this allows us to attribute changes in integrity without requiring a theory of mind (as in e.g. de Visser et al.; Mou et al., 2020). Our argument is two-fold. First, integrity arises from the action performed in a context (which we capture in the predictability function) and the individual and team values, as well as set principles by which the action can be judged. For example, assume that

some actions can benefit the agent (‘selfish’) and others benefit the team. In a context in which no other teammate will be affected, a ‘selfish’ action might be judged neutrally, but in a context where the action might be chosen in preference to one which could aid a teammate, the judgment might be negative. Or, in a context where some actions benefit the agent and the team, but are clashing with set principles (‘do not break the law by X’), the judgment will be negative. Second, each agent will have a *reputation* which reflects the history of these instances of integrity. As other teammates learn the history of the actions performed by an agent, so the reputation of that agent will be formed. Given an understanding of a teammate’s reputation, one can define an expectation of the action that they might perform. This can be modeled as a reinforcement learning problem (Anastassacos, Hailes, and Musolesi 2020; Anastassacos et al. 2021).

## Integrity, Reputation and Values

In our definition of trust, each member of a team will form estimates of their teammates’ capability, predictability and integrity as task performance is observed. Observation might be of the actual performance of the task or an outcome of this performance (either the result of the task or a report of from another agent). As information is accumulated, an agent will infer the *reputation* of its teammates. Reputation is a quantifiable factor, and (as we noted above) is based on the history of performing tasks in pursuit of goals relative to the individual and team values. However, rather than the expectation of prosocial behavior being fixed solely by reputation, we assume that this will be moderated by context. For example, a teammate might have a reputation for performing ‘selfish’ actions, i.e., seeking to gain rewards for themselves at the expense of their teammate. However, in a given situation, there might not be an option to assist a teammate or the outcome of the action might not be detrimental to the team. In this case, the severity of the outcome of a selfish action might be minimal. When there is sufficient information about a task being performed, the integrity of this task can be defined in terms of selfish or team values. When there is insufficient information, knowing the relationship between task and goal, and the context in which the task is performed, can enable the prediction of the most likely goal being pursued. We describe this by updating a Bayesian Belief Network that is held by each agent (Figure 4). Note, here we tend to think of reputation as being a short-term, mission-bounded metric, but we also recognize that this metric could interact with other factors, e.g. long-term trust weightings such as relationship equity (de Visser et al. 2020), or short-term weightings such as robot appearance (which may have a proxy trust effect; Lewis and Marsh, 2022); such weightings could be added into the BBN.

## Trust and Distributed Situation Awareness

We assume agents have a partial view of the context. This view consists of their perception of the environment, their inference of which goal is appropriate to perform in the context, their belief in the reputation of their teammates, and the goal that they expect their teammates to be pursuing. Situ-

Category	Variable	Derivation
WHEN	Time at which an action is performed	Clock reading for logged data
	Prediction of future action to be performed	<i>Inferred from mission plan</i>
	History of previous actions	Store of logged data
WHERE	Fixed location	Location of, e.g., AR Tag
	Path	Waypoints recorded from, e.g., ultrasonic tracking
WHAT	Object	Identification of object used
	Task	Action performed on an object
WHO	Agent	Agent ID
	Goal	<i>Inferred from mission plan</i>
	Reputation	Store of integrity from prior actions
WHY	Robots: Compatibility of own and other’s goal priorities and compliance with codified principles	<i>Inferred individual/team goal priorities</i>
	Humans: As above, plus judged performance criteria, constraints, and determinants of acceptable outcomes	<i>Inferred from individual, team and social values</i>

Table 1: Initial set of metrics to be measured or inferred. Similarly applicable to humans and robots, except for ‘Why’.

ation awareness in a team is likely to be distributed (Stanton et al. 2006) and we have previously demonstrated that Distributed Situation Awareness can be formally described using a Bayesian Belief Network (BBN) model (Yusuf and Baber 2022).

In a BBN, the system can be modeled using a graph  $G(N, E)$  where  $N$  is a set of nodes connected by a causal directional edges  $E$ . Each node represents an element of component-level trust with a defined number of states (i.e., selected situations as illustrated in Figure 4). States are defined as probabilities, i.e., between 0%-100% to reflect the assumption that these are estimates formed by an agent. These probabilities can be updated based on mission information (e.g., Equation 1), or learned over time (Yusuf and Baber 2022). As such, the trust elements (e.g., goals, tasks, reputation etc.) can be modeled using BBN nodes with assigned probabilities and causal relationships based on the operating mission context. For example, assume a scenario where an agent has two goals ( $G_A$  and  $G_B$ ), e.g.,  $G_A$ : to mark the locations of hazards in an environment by scanning AR Tags, and  $G_B$ : to construct the map of the environment using Simultaneous Localization and Mapping algorithms (SLAM). Each of these goals can be achieved by completing a number of tasks  $\alpha_i, \forall i \in [1, N]$ ; and these tasks can be spread across goals (e.g., the SLAM task to search for a hazard is the same as the one for a mapping task) i.e.,  $G_A \rightarrow \alpha_i \times \alpha_j$ , such that,  $\alpha_i \cap \alpha_j \neq \{\}$ ,  $\forall \alpha_i, \alpha_j \in G_A \cap G_B, i, j \in [1, N]$  or mapped singly to a goal  $G_i \rightarrow \alpha_i, \forall i \in [1, N]$  as illustrated in Figure 4. From Figure 4, an agent is capable or incapable of achieving a goal, and the goal achievement depends on the assigned task(s) and context (defined by time, location, and opportunity). Thus, each task achieved will increase the goal’s probability of success, i.e., a task with a 90% ‘achieved’ state has a higher contribution to the goal achievement than the one with 50% (though this depends on the criticality of the task towards the goal achievement Yusuf and Baber (2022)). Equation 1 is an example of a protocol-based mode of updating a probability of each state of the BBN after every mission event (e.g., sensor sampling by the agent):

$$P(R_i) = P(R_{i-1}) + f(w_c) \quad (1)$$

where  $P(R_i)$  is the updated prior of the state (e.g., after the

event occurrence),  $P(R_{i-1})$  is the previous prior of the state,  $f(w_c)$  is the probability decrement/increment weight function (to be assigning values based on subject matter expert (SME) judgments of the context  $c$ , e.g., information from a reliable sensor weights a complement of 100%),  $c$  is the mission context (e.g., as defined by time, location, and opportunity in Figure 4), and  $i$  is the event number, such that,  $i \in [1, N]$ . For example, if a goal has two tasks with a weight ratio of 9:1 (e.g., as assigned by  $f(w_c)$  of Equation 1) towards goal achievement, accomplishing task A contributes to 90%, i.e., 0% (as the current probability value of the ‘achieved’ goal state, i.e., assuming no prior progress on task B) +90% = 90% for the designated goal achieved state. Note that,  $f(w_c)$  reduces the prior  $P(R_i)$  for the non-occurring states and sum up the states probabilities to 100%.

The probability of a parent node can be defined by the child(ren) node(s) states using the conditional probability table (CPT), i.e., a table mapping parent node states probabilities with the joint child(ren) states. One of the advantages of modeling the system concepts using a BBN is the ability to predict states using conditional probabilities (Equation 2):

$$P(A) = P(A \cap B) / P(A|B) \quad (2)$$

where  $P(A)$  is the expected probability of the querying state  $A$ ,  $P(A|B)$  is the conditional probability of state  $A$  given  $B$ , and  $P(A \cap B)$  is the joint probability of  $A$  and  $B$ . As such, based on the previous reputation of an agent, its capability on a particular task can be predicted. Expectation maximization algorithms can be used to improve the prediction accuracy, i.e., by checking the agent’s performance history (Yusuf and Baber 2022).

## Illustrating a Ladder of Trust in the Experimental Environment

Having defined trust and proposed how it might be measured and captured in a BBN, we now explain how this framework can be tested in our experimental environment. From the motivating use-case, scanning the AR tags to define the safety of an area could be defined as a ‘team’ goal, in that it will benefit other team members. This goal could be further emphasized by giving the robot a competing ‘selfish’ goal: each scan could involve a cost to the robot (e.g., each scan

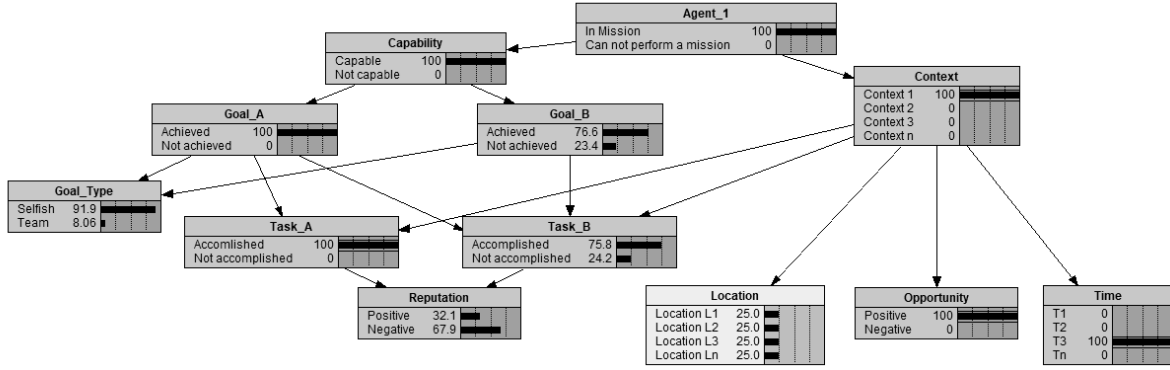


Figure 4: A teammate’s Bayesian Belief Network (BBN) based on their observation of Agent 1’s choice of action in a context.

requires energy and the robot will need to leave the environment to recharge when energy level are below a threshold) and the robot could be rewarded for minimizing cost (e.g., by staying as long as possible in the environment). The human could be rewarded for scanning all AR Tags within a time limit, but there could be a cost to entering unsafe areas. Scanning an AR Tag in an unsafe area could also represent a cost to robots, who would need to enter that area to help the human (e.g., by guiding the human to safety). Defining the potential for conflict between team members allows us to manipulate selfish and team goals, and hence to manipulate the ‘values’ of the team performance. Each team member will be asked, at set intervals, to rate what their teammates are doing (in terms of expectation of the actions between selfish or team goals). It is possible that the task is at the limits of the agent’s ability or demands significant resources. In this case, the expectation of a trustor is that the teammate (trustee) is exerting themselves to achieve the outcome. For a ‘selfish’ goal this could be interpreted as the agent recklessly pursuing a reward at personal cost (leading to an increase in distrust); for a ‘team’ goal this could be interpreted as the agent risking themselves for their teammates (leading to an increase in trust across trustors). As well as expectations about capability, the task will be interpreted in terms of predictability of outcome, i.e., is the outcome likely to be completely successful? As with the previous example, one might expect the reputation assigned to a teammate to be affected by the success of the outcome. Successful outcome(s) will increase perceptions of that teammate’s capability (Grillo et al. 2022); though if not calibrated could lead to overtrust (Ullrich, Butz, and Diefenbach 2021). The final element, integrity, relates to the interpretation of the goal against the goal priorities (robots and humans) and values (humans only) that the individual or team applies.

## Discussion

Distributed, Dynamic Team Trust (Huang et al. 2021) requires metrics that reflect the activity and interactions of members of a team. In this paper we share our definition of such metrics and illustrate how these can be applied to the conduct of a mission. We contribute to the debate on trust in

human-robot teams in the following ways. First, we propose that trust arises from a hierarchy of principles, values and goals. Second, we argue that ‘integrity’ (as a component of trust) should be judged in terms of the observed choice of task in a given context; this can lead to an inference regarding the choice of goal in that context. Where the goal might be considered selfish and where this might have negative consequences for teammates, this will lead to a lower perception of the integrity of the action. Over a history of observations of such choices made by a specific actor, the ‘reputation’ of that actor will be formed. It is likely that such a reputation will (a) reflect the observations of specific agents and, thus, might differ between agents, but (b) could be shared between agents. In contexts where there is no history to draw upon and hence, no evidence on which to define a ‘reputation’, this will either have to be assumed by the observing agent (e.g. ‘proxy trust’), or communicated by the other agent. For example, a robot might be programmed to assume that its human teammates will behave in a prosocial manner. This would lead it to ascribe a positive reputation to its human teammate – until it collected sufficient evidence (from observation or from other reports) to the contrary. Third, we argue that ‘trust’ is dynamic and context-dependent. In this, we are in agreement with Huang et al. (2021). Our approach has been to define the metrics for which each member of a team is able to acquire information. We aim to define these metrics in a way that allow humans or robots to perceive sufficient data to update similar Bayesian Belief Networks (BBN). This is not to assume that human reasoning is reducible to a BBN but allows the human to infer the robots’ choices (from a BBN that represents the robots’ behavior) and the robot to infer human choices. Fourth, the concept of a ladder of trust (on which perceptions of teammates can move up and down) provides a metaphor for the ways in which trust changes during a mission.

We intend to develop these concepts in a concrete way by carrying out a form of the experiment described in this paper, to observe whether our operationalization of trust metrics obtains plausible dynamics and supports our notion of ‘satisficing trust’ on a ‘ladder of trust’ (Baber et al. 2023). We hope to obtain insights that can guide both human and robot actors in dynamic trust building.

## Acknowledgements

The research reported in this paper is supported by grant EP/X028569/1 ‘Satisficing Trust in Human Robot Teams’ from the UK Engineering and Physical Sciences Research Council. This project runs from 2023 to 2026 and involves the Universities of Birmingham, Bristol, Loughborough and UCL.

## References

- Anastassacos, N.; Garcia, J.; Hailes, S.; and Musolesi, M. 2021. Cooperation and Reputation Dynamics with Reinforcement Learning . In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, 115–123. London, United Kingdom.
- Anastassacos, N.; Hailes, S.; and Musolesi, M. 2020. Partner Selection for the Emergence of Cooperation in Multi-Agent Systems using Reinforcement Learning. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*. New York City, NY, USA.
- Arkin, R. C.; Ulam, P.; and Duncan, B. 2009. An ethical governor for constraining lethal action in an autonomous system. Technical Report GIT-GVU-09-02, Georgia Tech.
- Baber, C.; Waterson, P.; Milivojevic, S.; Maynard, S.; Hunt, E. R.; and Yusuf, S. 2023. Incorporating a ‘ladder of trust’ into dynamic Allocation of Function in Human-Autonomous Agent Collectives. In *14th Organizational Design and Management Conference (ODAM)*. Bordeaux, France.
- Bigman, Y. E.; and Gray, K. 2018. People are averse to machines making moral decisions. *Cognition*, 181: 21–34.
- Conitzer, V.; and Oesterheld, C. 2023. Foundations of cooperative AI. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 15359–15367.
- de Visser, E. J.; Peeters, M. M. M.; Jung, M. F.; Kohn, S.; Shaw, T. H.; Pak, R.; and Neerinx, M. A. 2020. Towards a Theory of Longitudinal Trust Calibration in Human–Robot Teams. *International Journal of Social Robotics*, 12: 459–478.
- Forney, A.; Pearl, J.; and Bareinboim, E. 2017. Counterfactual data-fusion for online reinforcement learners. In *Proceedings of the 34th International Conference on Machine Learning (ICML 2017)*, 1156–1164. PMLR.
- Gabriel, I.; and Ghazavi, V. 2021. The Challenge of Value Alignment: From Fairer Algorithms to AI Safety. In Véliz, C., ed., *The Oxford Handbook of Digital Ethics*. Oxford University Press.
- Grillo, A.; Carpin, S.; Recchiuto, C. T.; and Sgorbissa, A. 2022. Trust as a metric for auction-based task assignment in a cooperative team of robots with heterogeneous capabilities. *Robotics and Autonomous Systems*, 157: 104266.
- Huang, L.; Cooke, N. J.; Gutzwiller, R. S.; Berman, S.; Chiou, E. K.; Demir, M.; and Zhang, W. 2021. Distributed dynamic team trust in human, artificial intelligence, and robot teaming. In Nam, C. S.; and Lyons, J. B., eds., *Trust in Human-Robot Interaction*, 301–319. Academic Press.
- Hurley, R. F. 2006. The decision to trust. *Harvard Business Review*, 84: 55–62.
- Jaques, N.; Lazaridou, A.; Hughes, E.; Gulcehre, C.; Ortega, P.; Strouse, D.; Leibo, J. Z.; and De Freitas, N. 2019. Social Influence as Intrinsic Motivation for Multi-Agent Deep Reinforcement Learning. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning (ICML 2019)*, volume 97 of *Proceedings of Machine Learning Research*, 3040–3049. PMLR.
- Jenkins, D. P.; Stanton, N.; and Walker, G. 2009. *Cognitive work analysis: coping with complexity*. Boca Raton, FL: CRC Press.
- Jian, J.-Y.; Bisantz, A. M.; and Drury, C. G. 2000. Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics*, 4: 53–71.
- Kim, P. H.; Dirks, K. T.; and Cooper, C. D. 2009. The repair of trust: A dynamic bilateral perspective and multilevel conceptualization. *Academy of Management Review*, 34: 401–422.
- Lewis, P. R.; and Marsh, S. 2022. What is it like to trust a rock? A functionalist perspective on trust and trustworthiness in artificial intelligence. *Cognitive Systems Research*, 72: 33–49.
- Malle, B. F.; and Ullman, D. 2021. A multidimensional conception and measure of human-robot trust. In Nam, C. S.; and Lyons, J. B., eds., *Trust in Human-Robot Interaction*, 3–25. Academic Press.
- Mayer, R. C.; Davis, J. H.; and Schoorman, F. D. 1995. An integrative model of organizational trust. *Academy of Management Review*, 20: 709–734.
- McAllister, D. J. 1995. Affect-and cognition-based trust as foundations for interpersonal cooperation in organizations. *Academy of Management Journal*, 38: 24–59.
- Moor, J. H. 2006. The Nature, Importance, and Difficulty of Machine Ethics. *IEEE Intelligent Systems*, 21: 18–21.
- Mou, W.; Ruocco, M.; Zanatto, D.; and Cangelosi, A. 2020. When Would You Trust a Robot? A Study on Trust and Theory of Mind in Human-Robot Interactions. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 956–962.
- Ng, A. Y.; and Russell, S. 2000. Algorithms for inverse reinforcement learning. In *ICML*, volume 1, 2.
- Pearl, J. 2009. *Causality*. Cambridge University Press.
- Pearl, J.; and Mackenzie, D. 2018. *The Book of Why: the New Science of Cause and Effect*. Basic Books.
- Rasmussen, J.; Pejtersen, A. M.; and Goodstein, L. P. 1994. *Cognitive Systems Engineering*. New York: John Wiley and Sons.
- Schaefer, K. E. 2016. Measuring trust in human robot interactions: Development of the “trust perception scale-HRI”. In *Robust Intelligence and Trust in Autonomous Systems*, 191–218. Springer.
- Stanton, N. A.; Stewart, R.; Harris, D.; Houghton, R. J.; Baber, C.; McMaster, R.; Salmon, P.; Hoyle, G.; Walker,

- G.; Young, M. S.; Linsell, M.; Dymott, R.; and Green, D. 2006. Distributed situation awareness in dynamic systems: theoretical development and application of an ergonomics methodology. *Ergonomics*, 49: 1288–1311.
- Tennant, E.; Hailes, S.; and Musolesi, M. 2023. Modeling Moral Choices in Social Dilemmas with Multi-Agent Reinforcement Learning. In *Proceedings of the 32nd Joint Conference on Artificial Intelligence (IJCAI 2023)*.
- Ullrich, D.; Butz, A.; and Diefenbach, S. 2021. The development of overtrust: An empirical simulation and psychological analysis in the context of human–robot interaction. *Frontiers in Robotics and AI*, 8: 554–578.
- Vicente, K. J. 1999. *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Walliser, J. C.; de Visser, E. J.; and Shaw, T. H. 2023. Exploring system wide trust prevalence and mitigation strategies with multiple autonomous agents. *Computers in Human Behavior*, 143: 107671.
- Williams, A. T. 2009. Taking Values Seriously: Towards a Philosophy of EU Law. *Oxford Journal of Legal Studies*, 29: 549–577.
- Winfield, A. F. T.; Blum, C.; and Liu, W. 2014. Towards an Ethical Robot: Internal Models, Consequences and Ethical Action Selection. In Mistry, M.; Leonardis, A.; Witkowski, M.; and Melhuish, C., eds., *Advances in Autonomous Robotics Systems*, 85–96. Springer International Publishing.
- Yusuf, S. M.; and Baber, C. 2022. Formalizing Distributed Situation Awareness in Multi-Agent Networks. *IEEE Transactions on Human-Machine Systems*, 52: 1166–1175.