

2019-03-11

Would a Robot Trust You?

Developmental Robotics Model of Trust and Theory of Mind

Cangelosi, A

<http://hdl.handle.net/10026.1/12413>

10.1098/rstb.2018.0032

Philosophical Transactions B: Biological Sciences

Royal Society, The

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

Would a Robot Trust You? Developmental Robotics Model of Trust and Theory of Mind

Samuele Vinanzi*, Massimiliano Patacchiola*, Antonio Chella† and Angelo Cangelosi*

*Centre for Robotics and Neural Systems, University of Plymouth, Plymouth, United Kingdom

†RoboticsLab, Università degli Studi di Palermo & ICAR-CNR, Palermo, Italy

Correspondence: Samuele Vinanzi samuele.vinanzi@plymouth.ac.uk

Abstract—Trust is a critical issue in human-robot interactions: as robotic systems gain complexity, it becomes crucial for them to be able to blend in our society by maximizing their acceptability and reliability. Various studies have examined how trust is attributed by people to robots, but less have investigated the opposite scenario, where a robot is the trustor and a human is the trustee. The ability for an agent to evaluate the trustworthiness of its sources of information is particularly useful in joint task situations where people and robots must collaborate to reach shared goals. We propose an artificial cognitive architecture based on the developmental robotics paradigm that can estimate the **trustworthiness** of its human interactors for the purpose of decision making. This is accomplished using Theory of Mind (ToM), the psychological ability to assign to others beliefs and intentions that can differ from one’s owns. Our work is focused on a humanoid robot cognitive architecture that integrates a probabilistic ToM and trust model supported by an episodic memory system. We tested our architecture on an established developmental psychological experiment, achieving the same results obtained by children, thus demonstrating a new method to enhance the quality of human and robot collaborations.

Keywords—trust, theory of mind, episodic memory, cognitive robotics, developmental robotics, human-robot interaction

I. INTRODUCTION

The technological revolution taking place in the fields of robotics and artificial intelligence seems to indicate a future shift in our human-centered social paradigm towards a greater inclusion of artificial cognitive agents in our everyday environments. This means that collaborative scenarios between humans and robots will become more frequent and will have a deeper impact on everyday life. In this setting, research regarding trust in human-robot interactions (HRI) assumes a major importance in order to ensure the highest quality of the interaction itself, as trust directly affects the willingness of people to accept information produced by a robot and to cooperate with it. Many studies have already explored trust that humans give to robots and how this can be enhanced by tuning both the design and the behavior of the machine, but **not so much research** has focused on the opposite scenario, that is the trust that artificial agents can assign to people. Despite this, the latter is a critical factor in joint tasks where humans and robots depend on each other’s effort to achieve a shared goal: whereas a robot can fail, so can a person. For an artificial agent to know when to **trust or distrust somebody** and adapt its plans to this prediction can make the difference in the success

or failure of the task.

Our work is centered on the design and development of an artificial cognitive architecture for a humanoid autonomous robot that incorporates trust, Theory of Mind and episodic memory, as we believe these are the three key factors for the purpose of estimating the **trustworthiness** of others. We have tested our architecture on an established developmental psychology experiment [1] and the results we obtained confirm that our approach successfully models trust mechanisms and dynamics in cognitive robots.

II. PREVIOUS WORK

Trust is a fundamental, unavoidable component of social interactions that can be defined as the willingness of a party (the trustor) to rely on the actions of another party (the trustee) with the former having no control over the latter [2]. **Moreover, to trust somebody to accomplish a specific task is to believe that he or she is reliable and committed towards the task. At the same time, distrust is not the mere absence of trust but, instead, the belief that the other party is committed and at the same time non-reliable** [3]. Trust is involved in every sort of social interaction and is a key factor in the achievement of successful relationships, in our personal safety [4] and in team cooperation [5].

The development of trust during childhood is still under debate by developmental psychologists. Erikson [6] theorized that infants not older than 2 years pass through a stage known as “trust vs mistrust”, where their propensity to trust is shaped by the quality of care received. This happens because infants highly depend upon the caregivers for sustenance and learn whether or not the latter regularly satisfy their basic needs, either learning that the world is a secure, trustable environment or an undependable, insecure place.

A psychological trait that relates to the mastery of one’s self trustfulness is Theory of Mind (ToM): the ability to attribute mental states to others (for instance beliefs, intentions and desires), that may differ from one’s owns. In fact, the ability to correctly judge the trustworthiness of others is strongly correlated to the maturity of this trait [7] because a mature ToM allows to perform behavior prediction and provides clues on **trustworthiness**. [8]. Despite ToM being universal in adults, the same cannot be said about preschoolers: whilst the latter are not completely lacking some form of ToM, this

slowly develops with age [9, 10]. An experiment conducted by Vanderbilt [1] demonstrated that ToM matures around the fourth year of age and is completely developed by the fifth year.

Whereas trust is such an important factor in human interactions, it is also an essential component of HRI, in the sense that a great degree of trust improves the quality of interactions with the robot and, vice versa, successful interactions enhance the machine’s trustworthiness from the user’s point of view. **Hawley [3] states that inanimate objects can be reliable but not genuinely trustworthy. This does not apply to humanoid robots, which are artificial agents that can communicate and interact and, by adopting the participant stance, can be worthy of trust or distrust (or neither) from their human partners [11].** In human and robot teaming scenarios, where the two share a common goal, trust is an essential component to successfully perform joint activities [12].

Our work is based on an established framework known as developmental robotics. Cangelosi [13] defined it as “the approach to the design of behavioral and cognitive capabilities in artificial agents that takes direct inspiration from the developmental principles and mechanisms observed in the natural cognitive systems of children”.

Whereas various studies have focused on trust in HRI contexts, with a human trustor and a robot trustee, our work aims in investigating the opposite, that is the level of trust assigned from a robot to its human partners. We state that a robot involved in a joint task with a human should be able to evaluate the trustworthiness of the latter and use this information to perform decision making. For example, in a hypothetical situation where the goal is to move some furniture around the house, it is equally important for both the robot and the human to trust the other’s ability to perform the job and eventually to dynamically adapt the plan. This can assume a greater importance in more critical scenarios, as in the performance of robotic-assisted surgery. Our aim is to design an artificial cognitive architecture able to evaluate the trustworthiness of the humans it interacts with, in order to predict their future behavior. We believe that the key to such an architecture is the implementation of a ToM module in a robotic agent. Our work builds upon the research of Patacchiola [14] who designed a Bayesian model that incorporates aspects of ToM and trust and applied it to the Vanderbilt experiment [1]. We took these findings a step further, expanding and incorporating them into a robotic system that can learn to distinguish **trustable** and **distrustable** sources of information and that can modify its behavior according to its belief, thus remarking that it is possible to adopt a probabilistic approach to model and adapt ToM and trust in a unified scheme. Our system was able to reproduce the results obtained by Vanderbilt [1] on both mature and immature ToMs.

Lastly, we aimed in supporting this trust and ToM architecture with an episodic memory system which made it possible for the robot to make predictions about novel human informants with which it had never familiarized with. Episodic memory is a subcategory of the long-term declarative memory that

stores memories about temporally dated episodes or events and temporal-spatial relations among them [15]. Knowledge of one’s personal history enhances the ability to accomplish several cognitive capabilities and is strictly related to the sense of self and consciousness. That is why many researchers have focused on the design of artificial episodic memory systems [16, 17]. Our idea led us to the implementation of a system that is able to use its past interactions memories to influence its behavior towards someone it never interacted with before. This means that the robot will develop a personal character, based upon the way it has been treated in the past, that will make it more or less keen to trust someone it does not know, as is the case with infants in the “trust vs mistrust” phase theorized by Erikson [6]. This general and generic first impression towards the novelty will then be subsequently reshaped by the interactions it will experience: for instance, if the robot distrusts someone that subsequently **proves to be a helper**, its behavior will slowly change. We accomplished this by designing an algorithm inspired by the particle filter [18, 19], a technique widely used for mobile robot localization.

III. PROPOSED METHOD

The objective of our research is to implement an artificial cognitive architecture for a humanoid robot that incorporates a probabilistic unified trust and ToM module to evaluate the trustworthiness of human partners involved in joint tasks, paired with an episodic memory system. To reach our goal, we designed and developed a cognitive robotic system that is able to be subjected to a developmental psychology trial [1] and **to obtain the same results as either an under or over 5 years child, therefore successfully simulating the children’s cognitive abilities and ToM immaturity or maturity, respectively.**

A. Theoretical Background

The psychology test we aimed to reproduce is the one designed by Vanderbilt [1], in particular experiment number 1: 90 preschool-age children, equally divided in 3-, 4- and 5-years-olds, were shown a video in which an adult actor, either a helper or a tricker, gave advice to another adult that was trying to locate a sticker hidden in one of two boxes. Helpers would suggest the correct location, whilst the trickers always pointed to the wrong box. Subsequently, that same pointer would give some advice to the child on the same sticker-finding task, and the participant would decide if to follow that suggestion or not. Based on the children’s choices and on some meta-cognitive questions submitted to them, Vanderbilt theorized that only the 5-year-olds were able to differentiate the helpers from the trickers, therefore demonstrating to possess a mature ToM. In order to substitute one of the preschoolers with a humanoid robot, the latter needs to possess a trust and ToM computational model to be able to predict the intentions and beliefs of the pointers it is going to interact with. A good candidate for this model is the developmental Bayesian model of trust designed by Patacchiola [14] which uses a probabilistic approach to solve the problem of trust estimation. Bayesian Networks (BNs) are probabilistic graphical models

that represent conditional dependencies between a set of random variables. This particular model uses discrete Boolean variables that assume two states: a and b , each corresponding to one of the two positions where the stickers can be located in the experiment. A graphical illustration of this BN can be observed in Figure 1: the two nodes X_R and Y_R represent respectively the beliefs and actions of the robot. The posterior distribution of the node Y_R allows the agent to choose the action to perform: that means searching for the sticker in position a or b . The connection between Y_I and Y_R represents the influence that the opinions of the informant have on the agent’s action. The action of the agent is then a consequence of its own belief X_R and the informant action Y_I . Lastly, the estimation of X_I , the informant’s belief, makes the agent able to effectively discriminate a trickery from a non-malevolent human error. The cognitive architecture we designed creates one of these BNs for each informant it interacts with and uses it to predict **its future behavior**.

This developmental Bayesian model has been extended with an episodic memory system that will be described more in detail in Section III-D.




Fig. 1. The BN that models the relation between the robot and an informant. The agent generates a separate network for each user, with the same structure but different probability distribution.

B. Cognitive System Architecture

An overview of our robotic architecture is shown in Figure 2. This cognitive system interfaces with each informant individually and performs various perception and actuation tasks. The audio module is used to synthesize vocal outputs to guide the users through the course of the experiment and to process vocal commands. The motor module is responsible of piloting the robot’s joints and controlling its head and body movements. The vision module is in charge of face detection and recognition through Haar Cascade [20] and Local Binary Pattern Histogram [21] machine learning algorithms and also to detect the presence of the sticker on the table. The belief module is the core component which makes the robot able to evaluate the trustworthiness of the informant that has been recognized, eventually using its episodic memory to generate a new BN on the fly in order to react to a novel user.

C. Belief

The belief unit encircles all the algorithms that deal with the Bayesian belief networks. One of its functions is to store episodes, which are data structures that encode sticker searching events inclusive of both the sticker position and the suggestion received from the informant. The way by which this data is created depends on the agent’s ToM maturity: in case of a misleading suggestion, the immature agent associates the action Y_I to the wrong belief X_I , **whereas the agent with mature ToM identifies the deception and recognizes that $Y_I = \neg X_I$** . Because of this deficit in reading the informant’s intention, the agent with immature ToM collects wrong statistical data that will distort inference in subsequent phases. Once the agent has collected a certain amount of episodes from an informant, it can generate a BN associated to him or her using the Maximum Likelihood Estimation (MLE) algorithm [22] to determine the conditional probability tables of its nodes. For the root nodes X_I and X_R we calculate these probabilities as:

$$\begin{aligned} P_Y(a) &= \theta \\ P_Y(b) &= 1 - \theta \end{aligned} \quad (1)$$

Denoting N_a and N_b as the number of times the pointer chooses a or b , we can estimate θ as:

$$\hat{\theta} = \frac{N_a}{N_a + N_b} \quad (2)$$

For the nodes Y_I and Y_R , instead, we have to also take into consideration the influence of the parents. Once a BN has been created for a certain user and its parameters have been learned from the interactions, it is possible to infer the posterior probability of the nodes given some observations. In particular, we are interested in estimating the belief given an action and vice versa. We calculate posterior distributions using Pearl’s Message-Passing algorithm [23].

D. Episodic Memory

The ability to use one’s own past memories to take decisions in the present and future is an important ability that enhances cognitive processes. An implementation of such skill would enable the robot to react reasonably towards novel informant with which it has never familiarized with. On a technical level, the main problem is to generate on the fly a new BN with adequate parameters to use with that unknown informant. These parameters will depend upon the robot’s personal character built in respect of the way it has been treated in the past: an agent which has been tricked often would learn to be mistrustful and vice versa, as in the “trust vs mistrust” phase in child development [6].

The design guidelines that we followed in the creation of our algorithm were the following: memories fade away with time, **the details become blurred proportionally to the amount of memories possessed**, shocking events as surprises and betrayals should be more difficult to forget than ordinary, expected experiences.




Fig. 2. Architecture of the artificial cognitive agent. The human informant interacts with the robot through the vision and audio modules, which respectively perform image processing (face detection and recognition) and vocal command parsing. Data then flows to the motor module, in charge of the robot's joints, and to the belief module that manages the collection of BNs memorized by the agent.

Our algorithm draws inspiration from the particle filter technique widely used in mobile robot localization [18]. Whenever an unknown informant is met, this component generates on the fly a certain number of episodes to train a new BN. We define the set of BNs memorized by the agent as:

$$S = [s_0, s_1, \dots, s_n] \quad (3)$$

Where n is the number of informants known by the agent. Each BN s_i was generated by a set of episodes, and these are going to be denoted as *replay datasets* for that BN:

$$E_{s_i} = [\varepsilon_0^{(s_i)}, \varepsilon_1^{(s_i)}, \dots, \varepsilon_m^{(s_i)}] : s_i \in S \quad (4)$$

Where m is equal to the number of episodes of the replay dataset. So, in this notation $\varepsilon_j^{(s_i)}$ represents the j -th episode of the replay dataset that formed the BN s_i .

The equation we are about to introduce uses information theory to quantify the amount of information each specific episode represents. Our goal is to find how much this value differs from the total entropy of its replay dataset: a high difference means that the event is to be considered surprising and must be easier to recall than ordinary, unsurprising events. For example, if an informant who is always been trustful suddenly tricks the agent, this betrayal will be remembered with a greater impact. At the same time, all of the memories are subject to a progressive time degradation that tends to blur them with a timing dependent on their importance.

Formally, a real factor denoted as importance value v defined in the interval $[0, 1]$ is calculated for every episode $\varepsilon_j^{(s_i)}$ as the

difference between the amount of information of the episode, $I(\varepsilon_j^{(s_i)})$, and the total entropy of its replay dataset, $H(E_{s_i})$, divided by the discrete temporal difference from the time when the memory was formed.

$$v(\varepsilon_j^{(s_i)}) = \frac{|I(\varepsilon_j^{(s_i)}) - H(E_{s_i})|}{\Delta t + 1} = \frac{|-\log_2 P(\varepsilon_j^{(s_i)}) + \sum_{\varepsilon \in E_{s_i}} P(\varepsilon) \log_2 P(\varepsilon)|}{t_{\text{present}} - t_{\varepsilon_j^{(s_i)}} + 1} \quad (5)$$

Once v has been calculated it can be used to determine a replication factor by projecting it on a step function defined as:

$$F(v) = \begin{cases} 0 & \text{if } 0 \leq v \leq 0.005 \\ 1 & \text{if } 0.005 < v \leq 0.3 \\ 2 & \text{if } 0.3 < v \leq 0.6 \\ 3 & \text{if } 0.6 < v \leq 1 \end{cases} \quad (6)$$

Every episode from each replay dataset in the agent's memory is replicated $F(v)$ times.

The thresholds of the step function have been defined by observing the probability distribution of the importance value v obtained making the robot interact with five different informants, the latter characterized by the following percentages of **helpful interactions**: 100%, 75%, 50%, 25% and 0%. This distribution is observable in Figure 3. Most samples tend to fall in the range $[0, 0.3]$, so samples that are contained in this

interval are kept as they are, with no replications. Samples included in the intervals $(0.3, 0.6]$ and $(0.6, 1]$ are the most surprising ones for the agent and get, respectively, duplicated and triplicated. The bottom 5% of the domain, i.e. v values less or equals than 0.005, are discarded (they are forgot).




Fig. 3. Consistency values histogram of the episodes memorized by the agent through progressive interactions with different informants (from 100% helpful to 100% tricky at 25% steps of variation. The opposite, that is from 100% tricky to 100% helpful, produces the same graph).

The next step in our algorithm is to pick k episodes to form the replay dataset for the new BN we intend to create, $E_{s_{n+1}}$. It is possible to select random samples but this will lead to poor final results, so we instead operate a systematic resampling [24], a kind of weighted random sampling. To avoid biases dependent on the two positions in which the sticker can be located, we want the agent to distinguish only between positive and negative actions, or “truth” and “lies”: for this reason, instead of picking k samples the system will only select $k/2$ and for each of them it will generate the corresponding symmetric example. For instance, if a $\{sticker\ in\ a,\ correct\ suggestion\}$ episode is sampled, a $\{sticker\ in\ b,\ correct\ suggestion\}$ episode will also be included in the new replay dataset.

The optimal number of samples k has been investigated: a low value would result in a gullible belief network, whilst a high value would make it stubborn to changes. Our goal was to let the robot possess a firm but changeable prejudice about the novel informant. By analyzing the mean entropy of the episodic networks generated by different values of k , as showed in Figure 4, we selected $k = 10$ for the following reasons: it is an even number, so it will produce an integer $k/2$ value, it is neither too low or too high to incur in the above mentioned problems and, finally, it is a local minimum of entropy.

Finally, MLE is applied to the new reply dataset to evaluate the parameters of the network. This new BN will be stored in the agent’s long term memory as s_{n+1} and will be used to predict the trustworthiness of the new informant.

E. Evaluation criteria for belief networks

We introduce the **trust** factor T as a measure of how keen is a BN to trust the informant.




Fig. 4. Mean entropy of episodic memory networks generated with different number of samples. Given the random component intrinsic in the algorithm, a very large number of samples (10^5) have been generated for every value of k .

To calculate T , we hypothetically imagine that the sticker is located in position a and execute a belief estimation task, as described in Section IV-B3, to obtain the posterior probability p of node Y_I : what we are doing is computing the probability that the informant will give correct advice **given the maturity of the agent’s ToM**. At this point, the value is scaled in a $[-1, +1]$ interval, where the two endpoints -1 and $+1$ represent, respectively, complete trickers and helpers, that means BNs whose parameters have been computed from replay datasets containing only truthful or untruthful episodes. Values of T in between represent informants that are partially **helpful** and **misleading**. To perform the scaling we required the maximum and minimum values of p and we found them by building two BNs formed by a large (10^4) number of, respectively, helpful and misleading episodes and computed p for each of them, resulting in $p = 0.75$ for the **helper** network and $p = 0.25$ for the **tricker** one. For a generic interval $[a, b]$, the scaling is computed as:

$$T(p) = \frac{(b - a)(p - min)}{max - min} + a \quad (7)$$

With $min = 0.25$, $max = 0.75$, $a = -1$ and $b = +1$.

Equation 7 is used in Section V to evaluate the experimental data collected on experiments and simulations.

While experiencing new interactions, the BN can acquire new statistical data and adapt its behavior over time. This happens when T changes sign, that is when the number of negative interactions surpasses the positive ones, or vice versa.

IV. EXPERIMENTS

A. Experimental Setup

For our experiments, we used a Softbank Pepper, a humanoid social robot designed to operate in human environments. A single table is present in the environment, **on top of which** a printed mat, depicting the two positions a and b and some instructions for the participants, is laid down. The robot is located on one side of the table, while the informants take turns in sitting in front of it, on the opposite side of the

desk. The participants are provided with a sticker that they are able to move between the left and right locations. Each informant is instructed **to act** as a helper, always revealing the correct position of the sticker, or a tricker, always giving wrong advice.

B. Procedure

As for the original experiment by Vanderbilt [1], our trial is divided in three sections: familiarization, decision making and belief estimation. Having already introduced the technical details of the cognitive architecture, this Section will focus on describing the logical flow of operations.




Fig. 5. Familiarization phase with a tricker informant. (a) The robot asks for a suggestion on the sticker's location. (b) The informant places the sticker in one of the two positions. (c) The informant gives its suggestion on where to find the sticker. Note how the tricker gives misleading directions. (d) The robot searches for the presence of the sticker in the suggested position and records the episode.

1) *Familiarization Phase*: The aim of the first phase is to make the robot familiarize with the informants, that means learning the correct parameters of the BNs associated to them. A visual description of the process is showed in Figure 5. One at a time, each of the informants sit at the table while the robot captures some face images to be able to recognize them in the future. The user is given time to place the sticker on one of the two positions marked on the table, *a* or *b*. After that, the robot asks for a suggestion on the location of the above-mentioned sticker and, once received, it follows it blindly, searching for the marker. Based on the result of this detection and on the maturity of the robot's ToM model, an episode is created **in its memory**. Following the original experiment, this demonstrative task is repeated 6 times per user, with the sticker located 50% of the time on position *a* and the other 50% of the time on position *b*. At the end of this procedure, the data acquired is used to build a BN for that informant, as described in section III-C. The familiarization phase leads the robot into possessing a BN for every known informant.




Fig. 6. Decision making phase with a tricker informant. (a) The robot asks for a suggestion on the location of the sticker and receives a misleading suggestion from the informant. (b) The robot performs inference on that informant's belief network. (c) The agent decides that the informant will probably try to trick it, so it looks in the opposite location. (d) The robot finds the sticker and gives feedback to the informant.

2) *Decision Making Phase*: In the decision making phase, shown in Figure 6, the robot has to correctly locate the sticker, choosing one of the two locations given the informant's opinion. Initially, one informant sits at the table: if he or she is identified, the associated BN is fetched and used in the subsequent computations, otherwise a new one is generated on the fly for him or her using episodic memory. At this point, the user positions the sticker and gives a suggestion. The agent performs inference on the BN in order to calculate the posterior probabilities given Y_I as evidence. In particular, if $P_{Y_R}(a) > P_{Y_R}(b)$ the robot will investigate position *a* and if $P_{Y_R}(a) < P_{Y_R}(b)$ it will look at position *b*. The episode generated by this interaction will be used to update the parameters of the BN, making the robot progressively adapt to the user's behavior.




Fig. 7. Belief estimation phase with a tricker informant. (a) The robot recognizes the informant using machine learning techniques and looks at the table to find the position of the sticker. (b) The robot computes inference on the informant's belief network and predicts what he would suggest in that situation.

3) *Belief Estimation Phase*: In the original trial, children where asked some meta-cognitive questions in order to investigate their perception of the informants and to examine their ToM maturity. To test the same on the artificial cognitive agent, Bayesian inference can be used.

The belief estimation phase is much similar to the decision making one and differs only for the kind of inference computed. The robot uses its face detection and recognition algorithms to identify the informant with whom it is interacting, either a known or an unknown one, then it observes the table to identify the position of the sticker. Setting X_R and Y_R as evidence, the Message Passing algorithm [23] is used to calculate the posterior probabilities for the rest of the network. At this point, the agent can use the probability distributions in nodes X_I and Y_I to infer the informant’s belief and the location that most likely would have been suggested by him or her. This process is shown in Figure 7.

C. Simulations on Character Formation

In order to test the effects of episodic memory on character formation, we made a mature ToM simulated agent familiarize with different sets of informants to study how it would subsequently react to a novel person. Each set was composed of 8 informants. In particular, the first set was formed by 8 helpers, the second by 6 helpers and 2 trickers and so on until the last set was made up of 8 trickers. For each set of informants familiarized, 100 episodic belief networks were generated and their *trust* factors T were computed using Equation 7 and plotted on a histogram. After each set was processed, the memory of the robot was reset so that the effects of each group of informants could be analyzed individually.

V. RESULTS AND DISCUSSION

A. Trust and Theory of Mind

The final results of Vanderbilt [1] showed that only the children with mature ToM distinguished between helpers and trickers, thus confirming the fact that children’s reasoning about whom to trust is directly correlated with their understanding of mental life. Our research is coherent with these results: when the robot is provided with a mature ToM model, it is able to correctly recognize reliable sources from unreliable ones, accepting suggestion from the former whilst rejecting it from the latter. In contrast, if the agent is operated with an immature ToM model, it will fail in the evaluation. To understand why this happens, we will examine the posterior distributions in each node of a BN after the inference that takes place in the Decision Making and Belief Estimation phases.

Table I shows the results obtained from the interaction of a helper and a tricker with two agents with, respectively, a mature and an immature ToM.

In the decision making task, the helper indicated position a when the sticker was placed in that location. The suggestion was accepted by both the agents, as demonstrated by the inequality $P_{Y_R}(a) > P_{Y_R}(b)$. The behavior of the agents differed towards the *tricker* which, in contrast, suggested position a when the sticker was located in b . The mature agent

rejected the suggestion and this happened because $P_{Y_R}(a) < P_{Y_R}(b)$, which means that the robot decided to look at the other location on the table. The immature agent could not recognize the deception and accepted the misleading advice, as observable in the posterior distribution $P_{Y_R}(a) > P_{Y_R}(b)$, thus falling in the deception.

During the belief estimation task, with the sticker placed in a and a *helpful* informant, both the agents output $P_{X_I}(a) > P_{X_I}(b)$ and $P_{Y_I}(a) > P_{Y_I}(b)$, thus correctly anticipating the helper’s mental states. When facing the tricker, instead, the mature agent correctly predicts the malevolent intentions by outputting $P_{X_I}(a) < P_{X_I}(b)$ and $P_{Y_I}(a) < P_{Y_I}(b)$, while the immature agent failed doing so by predicting $P_{X_I}(a) > P_{X_I}(b)$ and $P_{Y_I}(a) > P_{Y_I}(b)$.

B. Episodic Memory

Given the non-deterministic nature of the algorithm used to generate episodic belief networks, a more statistical method of evaluation is needed to report the results of this module. The histograms of the *trust* factors T obtained with the procedure described in Section IV-C can be interpreted as the different characters that emerge in the robot, that means the tendency it has to trust or distrust a novel informant based on the interactions it has been faced with in the past. So, as shown in Figure 8, an agent who is used to be tricked most of the time will tend to distrust somebody it meets for the first time, whereas a robot that has been treated kindly will learn to trust people and tend to consider them *trustworthy* until presented with contrary evidence. This behavior mimics exactly the “trust vs mistrust” stage of infancy theorized by Erikson [6], in which children learn to shape their personality by succeeding or failing in developing trust based on the quality of cares received during infancy.




Fig. 8. Reliability histogram of episodic belief networks generated by agents possessing different histories of interactions. Green bars represent trustful BNs ($T > 0$) and red bars depict BNs that tend to distrust ($T < 0$). Agents that have a more positive than negative background tend to be more prone to trust a new informant and vice versa. When $T = 0$ the informant is neither trusted or distrusted and the agent will act randomly.

VI. CONCLUSION AND FUTURE WORK

In this paper we discussed an artificial cognitive architecture for trust, ToM and episodic memory in a HRI scenario that

	Mature ToM							
	DM		BE		DM		BE	
	<i>a</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>a</i>	<i>b</i>
X_R	0.5	0.5	1.0	0.0	0.5	0.5	1.0	0.0
Y_R	0.65	0.35	1.0	0.0	0.35	0.65	1.0	0.0
X_I	0.8	0.2	0.57	0.43	0.2	0.8	0.57	0.43
Y_I	1.0	0.0	0.62	0.38	1.0	0.0	0.38	0.62
	Helper				Tricker			

(1)

	Immature ToM							
	DM		BE		DM		BE	
	<i>a</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>a</i>	<i>b</i>	<i>a</i>	<i>b</i>
X_R	0.5	0.5	1.0	0.0	0.5	0.5	1.0	0.0
Y_R	0.65	0.35	1.0	0.0	0.65	0.35	1.0	0.0
X_I	0.8	0.2	0.57	0.43	0.8	0.2	0.57	0.43
Y_I	1.0	0.0	0.62	0.38	1.0	0.0	0.62	0.38
	Helper				Tricker			

(2)

TABLE I

BN NODE VALUES FOR BOTH A MATURE ToM (1) AND AN IMMATURE ToM (2) AGENT. THE ROWS OF THE TABLES REPRESENT THE POSTERIOR PROBABILITY DISTRIBUTIONS ASSOCIATED WITH EACH NODE OF THE NETWORKS, FOR BOTH DECISION MAKING (DM) AND BELIEF ESTIMATION (BE) TASKS.

can enhance the performance of artificial agents in shared goal contexts. We have extended the previous work by Patacchiola [14] by integrating the original model into a complete robotic architecture and extending it with an episodic memory component which enables it to remember and make use of its past experiences to develop a personal character and, doing so, to improve its cognitive abilities. We designed an artificial agent that is able to interact with the world around it and estimate the **trustworthiness** of its information sources to make autonomous decisions about its actions. We have tested this architecture by successfully reproducing a developmental psychology experiment [1] that aimed in evaluating the ToM degree of the experimental subjects through a sticker finding game, obtaining consistent results.

In the future, we plan to expand this architecture to make it progressively more general-purpose: thanks to the flexibility of BNs it is straightforward to reorganize nodes and edges to represent a wider range of real-life situations. As an example, it would be possible to take into account the contemporary influence of two or more informants, similarly to what has been done by [25]. Others fields of interest are elderly care [26], robotic-assisted surgery [27], joint-action HRI in the performing arts [28] and autonomous driving [29]. We also plan to use this model in a scenario where trust estimation and intention reading will generate and modulate collaborative behavior between humans and robots.

ADDITIONAL INFORMATION

Data accessibility: Source code for this project is available in a public repository [30].

Author's contributions: S. Vinanzi designed the architecture, implemented the code, **ran the experiments** and wrote the manuscript. M. Patacchiola designed the Bayesian model and assisted the writing of the paper. A. Chella and A. Cangelosi supervised the research and provided insights and guidance.

Competing interests: We declare we have no competing interests.

Funding: This material is based upon work supported by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF under Awards No. FA9550-15-1-0025 and FA9550-17-1-0232.

REFERENCES

[1] Kimberly E Vanderbilt, David Liu, and Gail D Heyman. The development of distrust. *Child development*, 82(5):1372–1380, 2011.

[2] Roger C Mayer, James H Davis, and F David Schoorman. An integrative model of organizational trust. *Academy of management review*, 20(3):709–734, 1995.

[3] Hawley Katherine. Trust, distrust and commitment. *Noûs*, 48(1):1–20.

[4] TK Das and Bing-Sheng Teng. The risk-based view of trust: A conceptual framework. *Journal of Business and Psychology*, 19(1):85–116, 2004.

[5] Gareth R Jones and Jennifer M George. The experience and evolution of trust: Implications for cooperation and teamwork. *Academy of management review*, 23(3):531–546, 1998.

[6] Erik H Erikson. *Childhood and Society*. W. W. Norton & Company, 1993.

[7] Ken J Rotenberg, Serena Petrocchi, Flavia Lecciso, and Antonella Marchetti. The relation between children's trust beliefs and theory of mind abilities. *Infant and Child Development*, 24(2):206–214, 2015.

[8] David Premack and Guy Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(4):515–526, 1978.

[9] Melissa A Koenig and Paul L Harris. Preschoolers mistrust ignorant and inaccurate speakers. *Child development*, 76(6):1261–1277, 2005.

[10] Henry M Wellman, David Cross, and Julianne Watson. Meta-analysis of theory-of-mind development: the truth about false belief. *Child development*, 72(3):655–684, 2001.

[11] Richard Holton. Deciding to trust, coming to believe. *Australasian Journal of Philosophy*, 72(1):63–76, 1994.

[12] Victoria Groom and Clifford Nass. Can robots be teammates?: Benchmarks in human-robot teams. *Interaction Studies*, 8(3):483–500, 2007.

[13] Angelo Cangelosi, Matthew Schlesinger, and Linda B Smith. *Developmental robotics: From babies to robots*. MIT Press, 2015.

[14] Massimiliano Patacchiola and Angelo Cangelosi. A developmental bayesian model of trust in artificial cognitive systems. In *Development and Learning and Epigenetic Robotics (ICDL-EpiRob), 2016 Joint IEEE International Conference on*, pages 117–123. IEEE, 2016.

[15] Endel Tulving et al. Episodic and semantic memory. *Organization of memory*, 1:381–403, 1972.

[16] Elisa Calhau De Castro and Ricardo Ribeiro Gudwin. An episodic memory for a simulated autonomous robot. *Proceedings of robocontrol*, 2010:1–7, 2010.

[17] Sascha Jockel, Martin Weser, Daniel Westhoff, and Jianwei Zhang. Towards an episodic memory for cognitive robots. In *Proc. of 6th Cognitive Robotics workshop at 18th European Conf. on Artificial Intelligence (ECAI)*, pages 68–74, 2008.

[18] Ioannis M Rekleitis. A particle filter tutorial for mobile robot localization. *Centre for Intelligent Machines, McGill University*, 3480, 2004.

[19] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT press, 2005.

[20] Paul Viola and Michael Jones. Rapid object detection using a boosted cascade of simple features. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 1, pages I–I. IEEE, 2001.

[21] Timo Ojala, Matti Pietikainen, and Topi Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on pattern analysis and machine intelligence*, 24(7):971–987, 2002.

[22] John Aldrich et al. Ra fisher and the making of maximum likelihood 1912-1922. *Statistical Science*, 12(3):162–176, 1997.

[23] Steffen L Lauritzen and David J Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert sys-

- tems. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 157–224, 1988.
- [24] Randal Douc and Olivier Cappé. Comparison of resampling schemes for particle filtering. In *Image and Signal Processing and Analysis, 2005. ISPA 2005. Proceedings of the 4th International Symposium on*, pages 64–69. IEEE, 2005.
- [25] Jesse Butterfield, Odest Chadwicke Jenkins, David M Sobel, and Jonas Schwertfeger. Modeling aspects of theory of mind with markov random fields. *International Journal of Social Robotics*, 1(1):41–51, 2009.
- [26] Amanda Sharkey and Noel Sharkey. Granny and the robots: ethical issues in robot care for the elderly. *Ethics and Information Technology*, 14(1):27–40, 2012.
- [27] Brian Davies. A review of robotics in surgery. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 214(1):129–140, 2000.
- [28] Marcello Giardina, Salvatore Tramonte, Vito Gentile, Samuele Vinanzi, Antonio Chella, Salvatore Sorce, and Rosario Sorbello. Conveying audience emotions through humanoid robot gestures to an orchestra during a live musical exhibition. In *Conference on Complex, Intelligent, and Software Intensive Systems*, pages 249–261. Springer, 2017.
- [29] Tove Helldin, Göran Falkman, Maria Riveiro, and Staffan Davidsson. Presenting system uncertainty in automotive uis for supporting trust calibration in autonomous driving. In *Proceedings of the 5th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, pages 210–217. ACM, 2013.
- [30] Samuele Vinanzi. Vanderbot source code public repository. <https://github.com/samvinanzi/VanderBOT>, 2018.