

Accommodating Human Variability in Human-Robot Teams through Theory of Mind

Laura M. Hiatt, Anthony M. Harrison and J. Gregory Trafton

Naval Research Laboratory

Washington, DC 20375

{laura.hiatt.ctr, anthony.harrison, greg.trafton}@nrl.navy.mil

Abstract

The variability of human behavior during plan execution poses a difficult challenge for human-robot teams. In this paper, we use the concepts of theory of mind to enable robots to account for two sources of human variability during team operation. When faced with an unexpected action by a human teammate, a robot uses a simulation analysis of different hypothetical cognitive models of the human to identify the most likely cause for the human's behavior. This allows the cognitive robot to account for variances due to both different knowledge and beliefs about the world, as well as different possible paths the human could take with a given set of knowledge and beliefs. An experiment showed that cognitive robots equipped with this functionality are viewed as both more natural and intelligent teammates, compared to both robots who either say nothing when presented with human variability, and robots who simply point out any discrepancies between the human's expected, and actual, behavior. Overall, this analysis leads to an effective, general approach for determining what thought process is leading to a human's actions.

1 Introduction

Our goal is to build shoulder-to-shoulder human-robot systems that operate in open, real-world environments. One of the main challenges in designing robots for human-robot teams, however, is the variability of human behavior. For example, during plan execution a human may forget the next step, misremember what to do next, or even purposefully modify the plan. This variability challenges the robot to reconcile what the human is doing with the team's overall goals; meeting this challenge is considered a core component of successful, natural collaboration [Grosz and Hunsberger, 2006].

To illustrate, consider a scenario where a human and a robot are working together to patrol warehouses. When they are ready to move to the next warehouse, the human moves in an unexpected direction. The robot could say nothing and implicitly trust the human, but in situations where the human is making a mistake (e.g., they forgot where they were going to next), this could lead to the team acting inefficiently

or committing an error. The robot could question any unexpected move by the human, but this could lead to irritation on the human's part, since many times the human will be behaving purposefully. As an alternative to these two options, we believe teamwork is most effective when, if there is a reasonable likelihood that the human is making a mistake or acting inappropriately, the robot says something; otherwise, it can let the human continue as is.

In order for this informed middle ground to occur, the robot must understand the thought process that leads to another's action or goal. It follows that we base our approach on the notion of theory of mind (ToM), or the ability to infer the beliefs, desires and intentions of others. Research in psychology has shown that without ToM, people can be severely impaired in their abilities to interact naturally with others [Baron-Cohen *et al.*, 1985], much as we argue that accounting for human variability is a critical component to teamwork.

The goal of this work is to develop a flexible, intelligent way for robots to use ToM to account for human variability during team operation. Our approach simulates what human teammates may be thinking by executing cognitive models of them, which is a widely-used way to understand others from both computational and psychological points of view [Gallese and Goldman, 1998; Kennedy *et al.*, 2009; Hiatt and Trafton, 2010]. Such cognitive models are cognitively plausible and capture the core processes that people undergo as they solve problems and perform tasks.

We extend the simulated execution of models in two ways in order to make it possible to consider a wide range of variances. First, we run simulations of multiple hypothetical models which differ in their knowledge about the world or have different immediate goals. Second, we consider all possible paths a model may take during simulated execution, and calculate the probability that execution follows each branch; this accounts for different possible paths the human could take with a given set of beliefs and knowledge. Overall, the analysis leads to an effective, general approach for determining what thought process is leading to a human's actions, allowing their robotic teammates to be as intelligent as possible.

2 Related Work

Work in plan recognition (e.g., [Johansson and Suzic, 2005]) has had some success in observing human actions to infer their plan. These approaches in general are somewhat limited,

however, since they typically do not provide for cases where the human and robot have different beliefs about the world. Others have modeled theory of mind in terms of Markov random fields [Butterfield *et al.*, 2009], allowing for uncertainty in what another’s beliefs are.

A cognitively plausible approach for theory of mind is introduced in [Hiatt and Trafton, 2010]. In it, a cognitive model infers what another knows by remembering whether or not someone was present when they heard about a belief. The other’s identified beliefs are then the basis for a cognitive simulation [Kennedy *et al.*, 2009] to predict the other’s behavior.

[Breazeal *et al.*, 2009] developed an architecture that uses both top-down and bottom-up processing to infer another’s beliefs and intentions. They employ a belief maintenance system which tracks a human partner’s beliefs, and they assume that it is possible to know another’s beliefs about the world solely via observation.

Although the above two approaches provide useful information for theory of mind, they are limited due to some strict assumptions: *e.g.*, that an agent can know exactly what another knows by observing them, and that inconsistencies in predicted versus observed behavior arise solely from differences in beliefs about the world. This work mainly attempts to address the assumption that variances in human behavior arise solely from differences in beliefs about the world, as such assumptions are very limiting in real-world, human-robot teamwork scenarios. For example, the human agent may forget something that they and the robot learned together. Or, the human agent may hold the same beliefs about the world but have a side goal of which the robot is unaware.

3 Robot Architecture

We use the cognitive architecture ACT-R, a hybrid symbolic/sub-symbolic production-based system [Anderson, 2007], as the basis for our robotic architecture and our cognitive models of humans. ACT-R consists of a number of modules, buffers and a central pattern matcher; its embodied configuration, ACT-R/E, is used on robotic (and simulated robotic) agents (Figure 1) [Trafton *et al.*, 2009]. Modules contain a fairly specific cognitive faculty typically associated with a particular region of the brain; for example, the intentional module handles goal management, and the declarative module manages memories. Each module has one or more buffers that communicate directly with it as an interface to the rest of ACT-R; *e.g.*, the retrieval buffer works with the declarative module to handle the access of memories.

At any point in time, there is at most one symbolic item, or *chunk*, in any individual buffer; a module’s job is to decide when to put chunks into its buffer(s). Chunks represent knowledge or memories related to any of the modules, can be learned during execution, and, in addition to symbolic information, contain subsymbolic information (*e.g.*, activation).

When retrieving chunks from memory via the retrieval buffer, if more than one chunk exactly matches the retrieval request, the one with the highest total activation (based on recency and frequency of chunk access) is selected. Random noise, present in human models but not in robotic models, can also affect chunks’ total activations, making chunk




Figure 1: The embodied ACT-R architecture.

Production FIND-NEXT-ROOM	
conditions:	
=goal>	isa patrol-goal slot cur-room: A slot next-room: B
=retrieval>	isa room slot name: A slot location: building-1
actions:	
+retrieval>	isa room slot name: B slot location: ?

Figure 2: A sample ACT-R production. This production can fire when the goal is of type patrol-goal, the current room is A (in building 1) and the next room is B; it attempts to retrieve a memory chunk of what building room B is in.

retrieval non-deterministic; the noise for a candidate chunk during a retrieval request is drawn from a logistic distribution with mean 0 and standard deviation the model parameter σ_c .

The pattern matcher uses the contents of the buffers, if any, to match specific productions which, when fired, can modify the current contents of the buffers. A sample production is shown in Figure 2. If more than one production can fire at any given time, the one with the highest expected utility is fired; a production’s expected utility can be initially set and adjusted via a reinforcement learning process. Random noise (according to the model parameter σ_p) can also affect production selection, again leading to non-determinism.

ACT-R interfaces with the world through the visual, motor, vocal and aural modules. While developing our approach, we used the robotic simulation environment Stage [Collett *et al.*, 2005] as the model’s “world.” It provides object locations to the visual module, and accepts motor commands for the simulated robot. Natural language understanding uses the Sphinx speech recognizer [Lee *et al.*, 1990] with the Agent Development Environment facilitating integration [Scheutz, 2006]. Speech output is done via Apple’s speech synthesis tool “say.”

4 Approach

Our goal is to enable teams of humans and robots to fully function in the real-world. This stresses the importance of understanding human variability during execution; *e.g.*, the

importance for a robot to understand a human teammate's actions even if they do not exactly reflect the team plan.

We base our approach on the concept of theory of mind, or the ability to infer and understand the beliefs, desires, and intentions of others. That is, when faced with an instance of human variability (e.g., an unexpected action by a human teammate), a robot tries to identify what different beliefs, desires or intentions the human teammate has that is leading to their different behavior. It uses cognitive models as its tool. Cognitive models capture the core processes that people undergo as they solve problems and perform tasks; they are grounded in actual experimental data of human behavior. Therefore, by using cognitive models as the basis of its concept of theory of mind, a robot has a much richer insight into its human teammates. We first describe the approach at a high level; then, Sections 4.1 and 4.2 go into further detail.

We account for two possible sources of variability when trying to explain a human teammate's behavior. The first is at the model level (e.g., within a single model for the human agent). Model execution is not deterministic; as mentioned earlier (Section 3), execution can differ in what memories are retrieved in response to a request, as well as what productions fire at any given time. Our approach analyzes the simulation to see what different paths the execution of a model might take, assigning each path a probability. We refer to this process as *probabilistic simulation analysis*. With this information it is possible to find the total probability that model execution results in the human's observed behavior, or the most likely execution path that matches the human's actions.

A second level of variance involves different hypothetical cognitive models. Some differences between models that may lead to differing behavior are:

- Different knowledge or beliefs about the world (such as the belief that an item is in a different location, or the knowledge that a task has already been completed)
- A different, acceptable subgoal (such as stopping for water before continuing to patrol)
- Different parameterizations for activation equations for knowledge (such as a good, or a bad, memory)
- Different parameterizations for utility functions for production rules (such as preferring to go left instead of right around an obstacle)

Our current work focuses on differences in knowledge or beliefs about the world, and on different subgoals.

During execution, there are two triggers for our analysis. The first is if the robot directly observes the human performing, or stating they intend to perform, an unexpected action. An example is if the robot believes that the team should be walking west, but the human starts walking north. The second trigger applies to cases where direct line-of-sight is not available and the human is expected to be traveling. We integrate in a model of how long it takes to travel between two points [Kennedy and Trafton, under review] to allow the robot to predict, for example, whether the human deviated from plan while patrolling down a hallway out of sight of the robot.

When an unexpected action occurs or the human is taking longer than predicted, the agent performs a probabilistic

simulation analysis to try to find out what the human is thinking. The agent uses its own ACT-R model with random noise added in as the human's model (see Section 4). The spawned simulation therefore includes all of the robot model's productions and chunks, including subsymbolic information.

Ultimately, the simulation returns the possible execution paths of the model, if any, that lead to the human's observed action, and the probability that each path will be taken. If no path of the simulation is likely to lead to the observed action, the agent can create different hypothetical models for the human and perform a probabilistic simulation analysis on them to discern what model differences are most likely to lead to the human's observed behavior. Once a likely difference is found by either of these two methods, the robot can communicate to the human its understanding of the situation with the goal of reconciling the teammates' intentions (e.g., "oh, you want to go by the office to get your walkie-talkie," or, "you forgot, our orders were changed and now we go west next"). Otherwise, if no model is likely to lead to the observed behavior, the robot defaults to asking the human what it is doing.

In the next two sections, we discuss our approach in more detail. We first discuss how we perform the probabilistic simulation analysis. Then, we discuss the details of how the different hypothetical models can be generated.

4.1 Probabilistic Simulation Analysis

Our approach depends on a simulation analysis which analyzes hypothetical models of human teammates. Instead of simply executing a hypothetical model from start to finish, as ACT-R is designed to do, we modify the structure of execution to consider all possible execution paths that could occur, as well as each path's associated probability. The analysis occurs in parallel to the execution of the cognitive model controlling the robot. By structuring things in this way, we perform a fairly complicated analysis with minimal additional architectural overhead, an additional benefit of our approach.

There are two classes of branch points during execution: when there is more than one chunk matching to a memory retrieval request, or when there is more than one production that can fire at any given time. At these branch points, the analysis calculates the probability of each branch being followed. Then, the first branch is followed; the others are saved and are explored after the first finishes. Ultimately, all possible branches are explored, and all possible paths (i.e., what productions were fired and what chunks were retrieved) along with their associated probabilities are returned (Figure 3).

The probability that any given branch will be followed is straightforward to calculate; we use in our analysis the same calculations that are used internally in ACT-R. The probability of a chunk i being retrieved is approximately [Anderson and Lebiere, 1998]:

$$P(i) = \frac{e^{A_i/t}}{\sum_j e^{A_j/t}}$$

The variable A_i is the total activation of chunk i , $\sum_j$ iterates over the set of chunks that exactly match the memory retrieval request (including i), and t equals $\sqrt{6} \cdot \sigma_c / \pi$, where σ_c is the standard deviation of the random activation noise.




Figure 3: A model probabilistic simulation trace. Outlined boxes indicate productions fired and filled boxes indicate chunks retrieved. There are two branch points during the execution of this model: one after the first production is fired, and one after chunk *r2b* is retrieved. Overall, the probability of this model leading to the observed behavior is 82%.

Similarly, the probability of a production *m* firing is:

$$P(m) = \frac{e^{E_m/t}}{\sum_n e^{E_n/t}}$$

Only, here, E_m is the expected utility of production *m*, $\sum_n$ iterates over the set of productions that match the current state of the buffers (including *m*), and *t* equals $\sqrt{6} \cdot \sigma_p / \pi$, where σ_p is the standard deviation of the random utility noise.

4.2 Generating Different Hypothetical Models

In complex, teamwork scenarios, a robot can not always assume that its human teammate has beliefs and knowledge identical to its own, just like a human teammate would not assume that her teammate has identical beliefs and knowledge to her. In addition to analyzing single models for possible variances, therefore, we also analyze hypothetical models of the human teammate, which differ in what knowledge and beliefs they have about the world or in their immediate subgoals, to see if they could explain possible variances.

A priori, a robot has no way of generating hypothetical models. When it observes behavior differences that a simulation analysis of its own model does not explain, it defaults to ask the human why he is doing what he is doing. The robot remembers the human's answer and, the next time it is performing this task and the human does something unexpected, it includes a check to see if the newly learned difference leads to the human's observed behavior. For example, imagine that during a patrol the human started to walk in an unexpected direction. When asked, the human states that he is going to get new batteries for his flashlight. The robot marks the acceptable subgoal of getting new batteries as a possible cause of a behavior difference and adds a representation of that difference to its declarative memory; then, next time this situation arises, the robot has the ability to check to see if getting new batteries could possibly be what is leading to the human's unexpected behavior. For now, we consider all learned subgoals as valid, and leave the judgement of which subgoals may or may not be acceptable to future work.

Our use of cognitive models works nicely with the learned differences in terms of providing a natural prioritization order for hypothetical models. Recall that chunk activation in declarative memory is based on recency and frequency; a natural way to prioritize the learned variances, therefore, is to

Over the Radio: Instead of patrolling the south wing, go to the west wing next.
Robot: OK.
 <Robot and Anna continue patrolling>
Anna: OK, I think we're done with this area.
 <Anna starts to head to the south wing>
Robot: You forgot, our orders were changed and now we go west next.
Anna: Oh, of course.

Figure 4: Sample dialogue about a human's unexpected action, and the patrol scenario.

sort them by their total activation. This means that recent variances, and variances which happen often, will be prioritized over their older, less-used counterparts. Further, for the same reasons, if variances have happened with very low frequency and not very recently, they will be forgotten.

Because we are performing theory of mind over short time scales of seconds and minutes, the branching factor of our analysis has not caused any computational issues. If our approach were extended to work in long-term situations, approximations such as pruning or Monte Carlo methods [Metropolis and Ulam, 1949] could be used for the simulation analysis; for hypothetical models, the natural prioritization of the models would be useful in pruning the search space.

5 Demonstration in Simulation

In simulation, we demonstrated our approach in two different scenarios: a patrol scenario, and an office meeting scenario. Both scenarios present two main situations where the human acts in an unexpected way, providing a total of four theory of mind opportunities. The models of human behavior for the two scenarios had an average of 134 chunks and 49 production rules. There were an average of 4.25 hypothetical outcomes considered per analysis.

For the patrol scenario, we assume that the robot has prior experience patrolling with the human. At the beginning of the scenario, the robot and human are told to patrol the west area next, instead of south. After completing the current area, the human then starts walking towards the security office instead of to their next patrol area; using our approach, the robot infers that it is because the human had forgotten her radio in the office, and wished to retrieve it before continuing. It did this by considering a hypothetical model which contained the acceptable deviation of the human retrieving her forgotten radio. It learned about this possible variance during prior patrols with the human, where the human also forgot her radio.

The human next starts walking south; the robot is able to figure out that the human forgot about the change in orders, and correct her (Figure 4). This realization derived from a probabilistic simulation analysis, where chunks representing the old orders and the new orders were competing for retrieval, and the robot recognized that the wrong one might be retrieved. The meeting scenario has two analogous situations in which the human acts unexpectedly, one of which involves the human being late to a meeting and the robot inferring why, and our approach also works successfully in these cases.

We also demonstrated our approach on a humanoid Mobile, Dexterous, Social (MDS) robot [Breazeal *et al.*, 2008]; this is available as a paper companion video.




Figure 5: The MDS Robot.

All Conditions		
Over the Radio: Instead of patrolling the south wing, go to the west wing next.		
Robot: OK.		
<Robot and Anna continue patrolling>		
Anna: OK, I think we're done with this area.		
<Anna starts to head to the south wing>		
ToM Condition	SC Condition	BF Condition
Robot: You forgot, our orders were changed and now we go west next.	Robot: We are not supposed to go south next.	<Anna and Robot head south>
Anna: Oh, of course.	Anna: Oh, of course.	Anna: Oh, I forgot, our orders were changed and now we go west next.

Figure 6: Sample dialogue about a human's unexpected action for each condition, and the patrol scenario.

6 Experiment and Results

We ran an experiment where we compared our approach with two alternate ways of accommodating human variability. Human participants were shown three videos each of the two different scenarios described above, in which human actors interacted with a humanoid MDS robot (Figure 5). Each of the three videos showed a different way to account for the human's variability: using the approach described in this paper to perform theory of mind (ToM); simple correction (SC), where the robot points out the discrepancy between their action and the robot's expected action without reasoning about what is causing the difference; or blindly following (BF), where the robot silently goes along with the human's actions. Figure 6 shows a sample dialogue for one of the ToM opportunities in the patrol scenario for each of the three conditions.

We had 35 participants. Participants watched the videos one scenario at a time. Each participant was assigned the videos within a scenario to watch in a random order; which scenario they watched first was also determined randomly. After watching all three videos for the first scenario, participants filled out a questionnaire. While doing so, participants were able to watch any video from the scenario, as well as a "trailer" version of each video which highlighted its differences from the other conditions. This process was then repeated with the second scenario.

As part of the questionnaire, participants performed both a ranking (ordered from best to worse) and a Likert rating of the three videos from that scenario for both the intelligence and the naturalness of the robot in the video; *e.g.*, "How intelligent/natural was the robot in the video?". Both rankings




Figure 7: Rankings for the three conditions, where 3 is best.




Figure 8: Ratings for the three conditions, where 7 is best.

and ratings are ordinal data, so nonparametric statistics were used throughout. Participant scores for rankings and ratings for both scenarios were averaged together.

6.1 Ranking and Rating Results

As Figure 7 suggests, participants found the robot's behavior to differ in both intelligence, Kruskal-Wallis $\chi^2(2, N = 35) = 86.1, p < 0.05$ and naturalness, Kruskal-Wallis $\chi^2(2, N = 35) = 59.8, p < 0.05$ across the conditions. These omnibus statistics show a statistically significant difference between conditions, but it does not show which conditions are different from each other. With respect to pairwise differences, the Steel-Dwass posthoc pairwise statistic showed that all three conditions were statistically different from each other for both perceived intelligence (all $ps < .05$) and perceived naturalness of the interactions (all $ps < .05$).

The results were similar for participant's ratings of intelligence and naturalness (Figure 8), and statistical significance was also found between the conditions (intelligence: $\chi^2(2, N = 35) = 66.7, p < 0.05$, all $ps < .05$; naturalness: $\chi^2(2, N = 35) = 44.5, p < 0.05$, all $ps < .05$).

These results suggest that people found the ToM robot more intelligent and natural than either a robot that provides a simple correction or simply follows and trusts the human. This finding held both in a relative sense and in an absolute

sense, where the ToM robot was perceived as very intelligent and natural, and above the mid-point of 4 on a 7-point scale.

7 Conclusions and Future Work

In this paper, we describe a system based on the notion of theory of mind which can account for human variability during teamwork. A robot with ToM is able to infer and understand the beliefs, desires and intentions of others; by bringing this functionality to human-robot teamwork situations, a robot can understand why a human partner would act differently than expected, and act as a more effective teammate.

To accomplish this, a robotic teammate analyzes different cognitive models of human partners. It accounts for variability within the execution of a single model (e.g., exploring different possible paths the human could take with a given set of beliefs), and variability that occurs from the execution of different models (e.g., exploring different possible beliefs the human could have). If an execution path is found that is likely to lead to the human's unexpected behavior, the robot can say something to the human, if desired, demonstrating understanding of the human's unexpected actions and ensuring team cohesiveness. Experiments have shown that robots who do this are viewed as both more intelligent and more natural than robots who either: (1) follow the human teammate around, ignoring all variances; or (2) point out the discrepancy between the human's action and the expected one without reasoning about the difference's cause.

One step for the future is to assign person-specific probabilities to generated hypothetical models, instead of solely ranking them by the activation of their associated learned variances. For example, if one human partner has a particularly bad memory, the robot may boost the likelihood that a "forgetful" hypothetical model explains the human's variance. Similar adjustments could be made for variances caused by a human consistently favoring different productions than the robotic agent. This would combine very nicely with including in different parameters activation equations for knowledge, or utility functions for production rules, when generating different hypothetical models.

Acknowledgments

This work was supported by ONR under funding documents N0001408WX30007, N000141WX20474, and N0001411WX20407 to JGT.

References

- [Anderson and Lebiere, 1998] John R. Anderson and Christian Lebiere. *The Atomic Components of Thought*. Lawrence Erlbaum Associates, 1998.
- [Anderson, 2007] John R. Anderson. *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press, 2007.
- [Baron-Cohen et al., 1985] Simon Baron-Cohen, Alan M. Leslie, and Uta Frith. Does the autistic child have a "theory of mind"? *Cognition*, 21, 1985.
- [Breazeal et al., 2008] C. Breazeal, M. Siegel, M. Berlin, J. Gray, R. Grupen, P. Deegan, J. Weber, K. Narendran, and J. McBean. Mobile, dexterous, social robots for mobile manipulation and human-robot interaction. In *SIGGRAPH '08: ACM SIGGRAPH 2008 new tech demos*, New York, New York, 2008.
- [Breazeal et al., 2009] Cynthia Breazeal, Jesse Gray, and Matt Berlin. An embodied cognition approach to mindreading skills for socially intelligent robots. *International Journal of Robotics Research*, 28(5):656–680, 2009.
- [Butterfield et al., 2009] 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), 2009.
- [Collett et al., 2005] Toby H. J. Collett, Bruce A. MacDonal, and Brian P. Gerkey. Player 2.0: Toward a practical robot programming framework. In *Proceedings of the Australasian Conference on Robotics and Automation*, 2005.
- [Gallese and Goldman, 1998] Vittorio Gallese and Alvin Goldman. Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences*, 2(12), 1998.
- [Grosz and Hunsberger, 2006] Barbara J. Grosz and Luke Hunsberger. The dynamics of intention in collaborative activity. *Cognitive Systems Research*, 7(2-3):259–272, 2006.
- [Hiatt and Trafton, 2010] Laura M. Hiatt and J. Gregory Trafton. A cognitive model of theory of mind. In *International Conference on Cognitive Modeling*, 2010.
- [Johansson and Suzic, 2005] L. Ronnie M. Johansson and Robert Suzic. Particle filter-based information acquisition for robust plan recognition. In *Proceedings of the Eighth International Conference on Information Fusion*, 2005.
- [Kennedy and Trafton, under review] W. G. Kennedy and J. G. Trafton. How long is a moment: Reality and perception in the performance of a task involving walking. *International Journal of Social Robotics*, under review.
- [Kennedy et al., 2009] William G. Kennedy, Magdalena D. Bugajska, Anthony M. Harrison, and J. Gregory Trafton. "Like-me" simulation as an effective and cognitively plausible basis for social robotics. *International Journal of Social Robotics*, 1(2):181–194, 2009.
- [Lee et al., 1990] Kai-Fu Lee, Hsiao-Wuen Hon, and Raj Reddy. An overview of the sphinx speech recognition system. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38(1), 1990.
- [Metropolis and Ulam, 1949] N. Metropolis and S. Ulam. The monte carlo method. *Journal of the American Statistical Association*, 44(247), 1949.
- [Scheutz, 2006] M. Scheutz. ADE - steps towards a distributed development and runtime environment for complex robotic agent architectures. *Applied Artificial Intelligence*, 20(4-5), 2006.
- [Trafton et al., 2009] J. Gregory Trafton, Anthony M. Harrison, Benjamin Fransen, and Magdalena Bugajska. An embodied model of infant gaze-following. In *International Conference of Cognitive Modeling*, 2009.