Emulating the consistency of human behavior with an autonomous robot in a market scenario

Saad Ahmad Khan, Saad Arif and Ladislau Bölöni
Dept. of Electrical Engineering and Computer Science
University of Central Florida
4000 Central Florida Blvd, Orlando FL 32816
{skhan, lboloni}@eecs.ucf.edu, saadarif@gmail.com

Abstract

Mobile robots moving in a crowd need to conform to the same social standards as the human participants. Imitating human behavior is a natural choice in these situations - however, not every human behaves in the same way. On the other hand, it is known that humans tend to behave in a consistent way, with their behavior predictable by their social status.

In this paper we consider a marketplace where humans perform purposeful movement. With many people moving on intersecting trajectories, the participants occasionally encounter *micro-conflicts*, where they need to balance their desire to move towards their destination (their mission) with the requirements of the social norms of not bumping into strangers or violating their personal space. We model micro-conflicts by a series of two-player games.

We show that if all humans are using consistent strategies which are aware of their own social status and can infer the social status of their opponent, the overall social costs will be lower compared to scenarios where the humans perform inconsistent strategies (even if those strategies are adaptive). We argue that robots acting in social environments should also adopt consistent strategies and align themselves with the ongoing social structure.

Introduction

Robots interacting with humans in social settings must obey the social rules of the specific culture in which they operate. At the same time, the robots have a specific goal or mission to achieve. There is often a conflict between obeying the social rules and the most efficient way to achieve their mission. The only exception is the very specific case where the only goal of the robot is to behave in a socially acceptable way.

We are tempted to think that making a robot behave in a socially acceptable way is equivalent for the robot to mimic "human behavior". However, if we observe human social settings, we find that not all humans behave in the same way in all social encounters. First, human social behavior has a certain randomness even for seemingly identical settings. Second, humans vary their behavior in function of the opponent and the circumstances of the encounter. And finally, not every human choose to obey the social rules. On the other hand, it is a well known fact of psychology that the overall functioning of the social life depends on the *consistency* of behavior. One of the principal requirements of human social interaction is that the participants form a theory of mind of each other (Adolphs 2003). This allows them to predict the beliefs, goals and actions of the interaction partner. This allows for a significant variance on allowed behavior. However, a certain consistency in the behavior is required, as we cannot model or predict the mind of an erratically behaving interaction partner. The agents need to take this into consideration about their human interaction partners; furthermore they need to act such that the humans can form a predictive model of them. Research show that humans are willing to treat agents as social actors (Nass, Steuer, and Tauber 1994) although in some situations they will treat humans differently from agents or robots (Sanfey et al. 2003). In most social settings, it is not the individuals pursuing aggressive or defensive strategies who are causing the most social disturbance but the ones who are erratically switching between the two.

The objective of this paper is to develop models of human behavior in a specific social setting (movement in a marketplace). We develop models of the mission cost and social cost of movement, and introduce a framework where humans need to balance between the two in social encounters called micro-conflicts. The micro-conflicts are modeled as a series of two-player games, in which the participants must deploy specific strategies. The consistency of the behavior does not mean that every human deploys the exact same strategy every time (in fact, such an overly uniform strategy creates problems in which symmetrical strategies can be broken only by one party abandoning the game). Rather, a consistent behavior means that the behavior in the micro-conflict can be predicted from observable attributes of the participant.

The scenario considered in this paper is as follows. In a busy marketplace a number of customers perform a purposeful movement. They visit various landmarks such as stores and stalls where they spend a certain amount of time, then they move to other landmarks. We assume that agent movements are independent, *i.e.*, group movement patterns aren't under consideration for this particular paper. The movement of the humans between the landmarks follow planned trajectories, which avoid obstacles, but try to get from one landmark to the other in the shortest amount of time. Reaching their destination in the planned time is the *mission* of the individual human. Delays represent *mission costs*, which the agent tries to minimize. At the same time, the humans need to obey social norms, which require them not to bump into other humans, violate their personal space or block their movement. If they violate these norms, humans incur *social costs*. If two humans are about to collide with each other, they need to take actions to avoid this by one or both of them changing their speed and or trajectory. We call such an encounter a micro-conflict. The strategies of the two agents in a micro-conflict must balance mission costs and social costs. We use the term "micro" to illustrate the fact that such conflicts are normally resolved very quickly (in matter of seconds).

The objective of this paper is to study different types of strategies which humans might deploy in such scenarios. We are particularly interested in how a consistent strategy would look like in this setting, and how a robot might emulate this. One of the important insights is that as important as it might seem to obey all the social rules, in a sufficiently dense crowd it is impossible to completely avoid incurring any social cost.

A multidimensional cost model

One way to quantify the decision making process of humans in social settings is by taking into consideration the costs and benefits of certain actions. While we can theoretically construct an integrated cost/benefit function in the form of a single scalar which captures the decisions of an agent, this function will have a complex and opaque expression which changes from agent to agent. In practice it is more convenient to consider a *vector of costs*, each of them being tied to a well-defined social norm, physical measurement or satisfaction level of a certain mission. These values are often traded off against each other, but they can not, in general converted into each other in an arbitrary and linear way. Similar multidimensional cost models have been used to model social scenarios (Bhatia, Khan, and Bölöni 2012). The peculiarity of the scenario we are considering is that the agents must consider a combination of social and mission costs. As the scenario we are considering considers only physical movement of the humans and the robot, ignoring, for the time being, gestures and verbal communication, all the costs are expressible through geometric models.

We will assume that humans and the robots use the same set of costs. For humans, the social costs are rooted in the psychology and social conditioning. While robots, naturally, do not have these factors, in order to be accepted by humans as natural participants in the crowd, they need to emulate the human behavior as closely as possible (Nakauchi and Simmons 2002; Walters et al. 2005; Pacchierotti, Christensen, and Jensfelt 2005).

The various cost dimensions can be grouped into two large categories: *social costs* depend on the social norms governing the environment and the participants while *mission costs* depend on the specific goals of the human or robot.

Modeling the social costs of moving in crowd

We will model the social costs of moving in the crowd by a number of geometrical zones associated with the opponent agents. An agent incurs costs whenever it enters into one of these zones. The zones are not necessarily circular, they move and change orientation with the agents. The costs associated with these zones are justified by psychological models of human perception, and they must be calibrated for the individuals as well as for the culture.

For the work in this paper we consider three zones:

Physical contact zone: represented by the actual physical size and shape of the human or robot agent. Violating this zone means physical contact and carries a large social cost.

Personal space: is the spatial region which a person (and by extension, a robot) regards as psychologically his (Hall and Hall 1969). Within the personal space, we model the personal distance (1-1.5 ft) and the social distance (3-4 ft). The cost decreases towards the outside of the area, becoming zero outside the social distance perimeter.

Movement cone: the movement cone represents the space where the human or the robot made public its intention to move. For the purpose of this paper, we consider the movement cone as circular pie extending from the agent in the current direction of movement, for a radius equal of 3 seconds movement with the current speed. The movement cone is only relevant for a mobile agent. By violating the movement cone, the opponent forces the agent to change its movement, unless it accepts a high social cost by violating the personal space or even the physical space.

We are using a model where the social costs are additive across the cost types and for the multiple agents. For instance, if the agent violates more than one agent's personal space, it will occur the sum of the costs. On the other hand we retain only the maximum social cost for each micro-conflict.

Modeling mission costs

We assume that the participants in the crowd have tasks to accomplish, thus any delay caused by a micro-conflict comes with a mission cost.

For a *non-urgent* mission, the mission can still be achieved at an arbitrarily later time - thus the mission cost of a delay is proportional with the delay. For *urgent* missions, the delay reduces the probability of mission success, thus the cost of the delay escalates in time.

For the purpose of this study we assume that the human participants have non-urgent missions. The mission of the robot is to follow its owner in the crowd. Repeated delays in the resolution of a micro-conflict make the robot fall more and more behind, modeled with a mission cost which for every new second of delay considers the full amount of time the robot is behind its owner. On the other hand, the robot is able to catch up with its owner between micro-conflicts (or equivalently, the owner will wait for the robot to catch up). Other assumptions are, of course, possible, but they are beyond the scope of this paper.

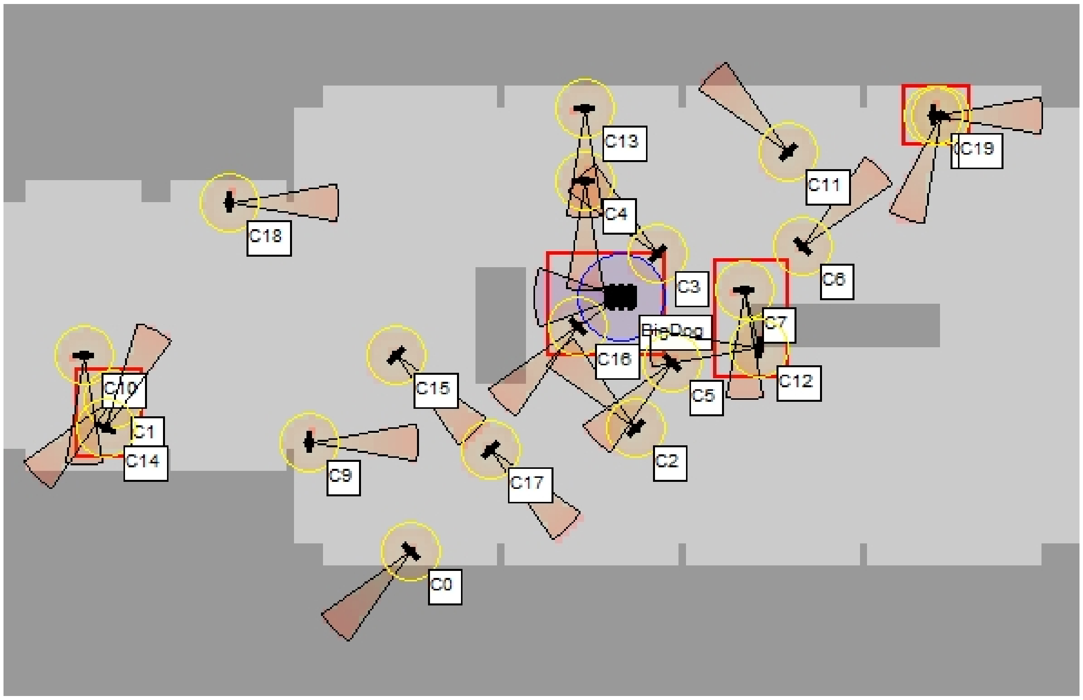

Figure 1: A moment in the scenario of the robot navigating a crowd of people on the market. The screenshot shows the visualization of the scenario in the simulator at time $t = 21sec$.

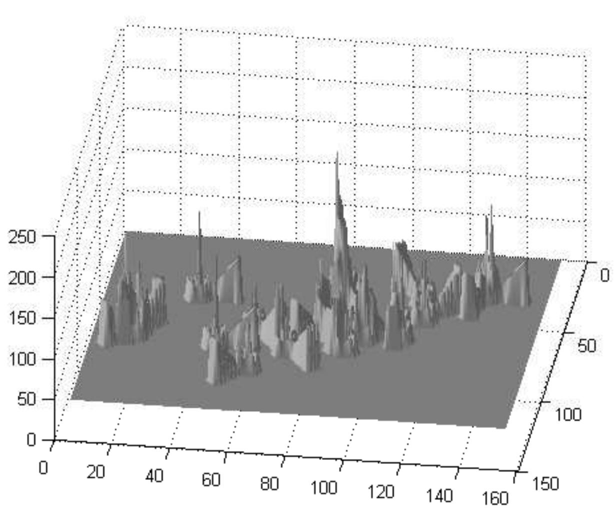

Figure 2: The diagram shows the cummulative social cost at that particular moment. The goal of the robot can be interpreted as an attempt to move while keeping to the "valleys" of this constantly changing surface.

Modeling of human strategies

Let us now discuss the strategies used by humans in resolving micro-conflicts. As we have discussed in the introduction, humans use a variety of strategies, and even the individual humans use different strategies from encounter to encounter. Nevertheless, human behavior must be consistent and predictable, erratic behavior is considered socially undesirable. To model these observations, we will make the assumption that humans classify their opponents into recognizable types, and then adapt different strategies against the various classes. We divide the population into six classes based on gender (male and female) and age (children, youth and elder).

The consistent behavior of humans, means that a person of a certain class (*e.g.* a young man) will deploy consistent behavior towards other classes (*e.g.* old man or young woman). Figure 3 describes the behavior.

Robot's strategy for humans in social context

The intent of an agent is not physically observable unless observations are made from past experience of micro-conflicts. The motivation behind the robot's strategy is to be consistent with its behavior during its interaction with different types of agents in social context. The robot's strategy uses a two-fold approach: there is a passive phase and an active phase. In the passive phase the robot classifies the intent of the agent and in the active phase it selects the best strategy (aligned with the social context). The term "passive" refers to the offline learning, i.e., the robot is trained with various examples of

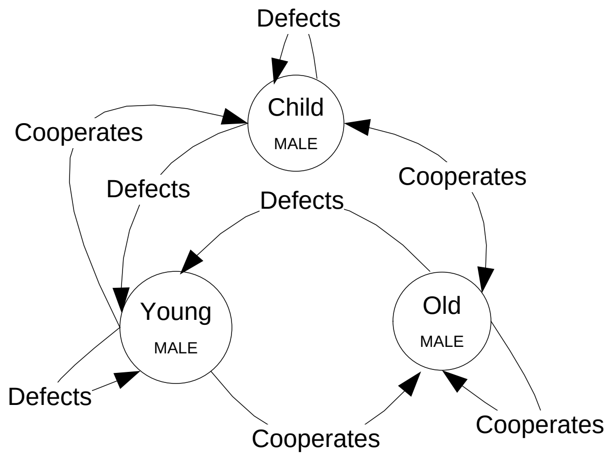

Figure 3: The social interaction between males belonging different age groups

human-agents. In the modeled system, some of the physical features of the human agents partially overlap. Hence, this fuzziness in the physical attributes of humans helps in introducing noise to the robot's identification system.

Humans tend to undergo erratic behavior in chaotic situations, e.g., humans have hysterical behavior patterns under panicking situations (Pelechano, Allbeck, and Badler 2007). In our model, the robot is trained to adapt with human behavior at micro-level (which is further linked with social-context at macro-level). Hence, for learning strategies to counter the heterogeneity in human behavior, the robot trains itself under the social context of the micro-conflict. For our work, we consider the micro-conflicts between humans and robot in dense crowds with regular human patterns. Therefore, the limited scope of this paper makes the passive strategy conditionally independent of the social context, i.e., the Naïve Bayes classifier takes the form

$$p(X = x_i \mid Y = y_j, Z = z_k) = p(X = x_i \mid Y = y_j) \quad (1)$$

where x_i is the category of the human-agent, y_i is the set of physical attributes and z_k is the social-context.

In this paper, the strategies of human agents, as shown in Table 1, are modeled assuming the social setting of routine life behaviors. The robot, after classifying the human-agent, further classifies the opponents' strategy using Table 1. The active part of the robot's strategy is to select an appropriate counter strategy contemplating the urgency of its mission. For example, for urgent missions the robot would try its best to minimize time-cost whereas for normal missions it would try to minimize the social-costs. For experimental analysis, we modeled the following strategies for the robot:

- **Respectful: (Maximal cooperate)** When an agent A is respectful to agent B, it means that it will give space wide enough to agent B to avoid any social cost. Therefore, agent A will cooperate over the time with agent B. This strategy is adopted by the robot when the urgency of mission cost is low.
- **Bump: (Maximal defect)** The agent A would minimize its mission costs and would not consider any social costs. The assumed response of agent B is that it will cooperate in any case. The strategy is adopted by the robot when the urgency of mission cost is high.
- **Consistent: (Naïve Bayes)** The agent A would classify agent B based on its social status. Therefore the assumed intent of either to cooperate or defect is based on social traits. The strategy is adopted by the robot for its consistent behavior within human interaction.

Experimental Setup

In the following we describe the results of a series of experiments involving the behavior of a robot in a crowd. The experimental scenario involves a marketplace in a Middle-Eastern country. The area is a narrow space surrounded with shops whose entrances serve as landmarks, as well as internal obstacles. A number of shoppers perform purposeful movement, which involves visiting shops for a shorter or longer times. The path chosen by the individuals balances the shortness of the path with the avoidance of the obstacles and large groups of people. Micro-conflicts are resolved through a succession of games using two different set of experiments as discussed afterwards.

In this baseline scenario we consider the presence of a patrol of peacekeeping soldiers traversing the market while being accompanied by a Boston Dynamics Big Dog robot (Raibert et al. 2008). The mission of the robot is to follow the soldiers through the crowd as closely as possible with consistent behavior towards the population. The soldiers can change their movement at any time, triggering frequent path re-plannings, for which we use the D*-lite algorithm (Koenig and Likhachev 2005). The robot participates in micro-conflicts in the same way as the human participants. Naturally, the robot's personal space and physical space is different from that of a human (a Big Dog robot is larger than a human).

We consider two different set of experiment for the the behavioral simulation of the robot. In the first set of experiments we evaluate different set of populations against consistent strategy of the robot. In other other set of experiments, we evaluated different strategies of robot against single population set of the humans. In both of the scenario, we humans agents play consistent sets of strategies against each other based on their social status.

Experiment Set I: Varying Population Density and Social Status

We consider three different sets of population each dominated by a particular social status. For each set we consider dense male population: the population is uniformly distributed with 70% males. For modeling various sets of population we consider three different times of the day. Therefore, the first phase of the day consists of majority of the agents belonging to the old age group, the second phase of the day (that is the evening time) is mostly populated by the children and in the last phase of the day the population is

Physical Features		ImmediateStrategy	Classification
Chin-type	Height (ft)		
Square, Pointed	1 - 4	Pass-front	Child (Male, Female)
Square	4 - 7	Pass-behind	Young Man
Pointed	4 - 7	Respectful	Young Woman
Square, Pointed	4 - 6	Bump	Old (Man, Woman)

Table 1: Modeled attributes of human agents for middle eastern social context

concentrated with young people. The distribution statistics for each time of the day is as follows

- **Morning-Time** - 10% Children, 20% Young
- **Evening-Time** - 30% Old, 30% Young
- **Night-Time** - 70% of Young, 10% Children

Experiment Set II: Varying strategies with same population distribution

In these sets of experiment, we vary different strategies of the robot for its micro-conflicts with the agent. The robot would either be very polite, bully or would play a strategy against another agent based on the social status of the opponent. The age group dynamics for the population in this particular experiment were 70% males, 40% young and 30% children.

Training the classifier of the robot

The robot uses a Naïve Bayes classifier which has been trained with a set of hundred examples using the data-set generated from Table 1. The height attributes of training set used for Naïve Bayes training set has the following statistics: After classification of the social status of the human, the robot selects the strategy using the social graph as shown in Figure 4.

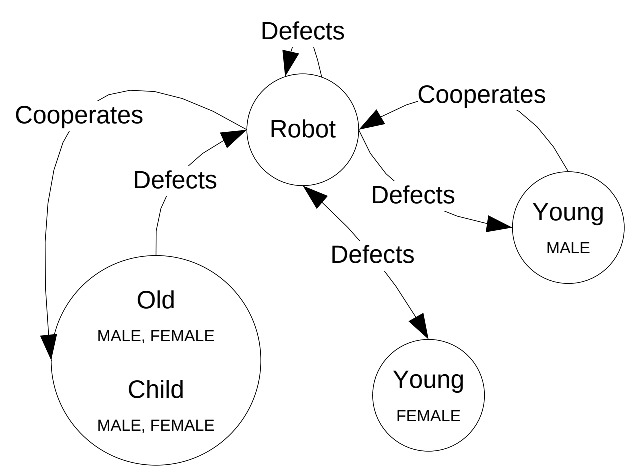

Figure 4: Social Behavior Graph for the robot

Results

For the first set of experimental results as given in Figure 5 we observe that on average the highest social cost was incurred by the robot when it was maneuvering for its mission during the night time. This is due to the fact that the robot defects for both gender types of the young people as given in Figure 4. Similarly, minimal social costs occur during the daytime and during the evening when the population consists mainly of the children and elderly people. Hence in this case the robot has minimal penalty for social costs when cooperates and others defect. This is direct reflection of the robots consistent cooperative policy of interaction with the children and elderly people (see Figure 4).

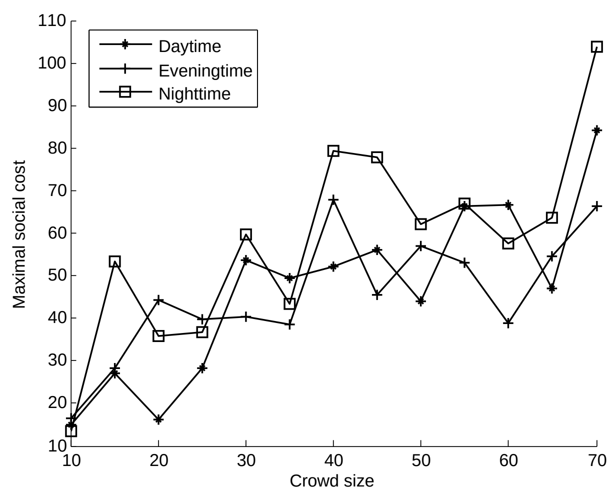

Figure 5: Experimental Results I. Sum of the maximum social costs / micro-conflict.

From Figure 6, we observe that the most expensive missions were the ones that took place during the daytime and the evening. The trend of high mission costs is associated with those duration most of the population was constituted of children and elderly people (above 70% in this case).

For the next experimental setup (see Experiment set II in Section), we had modeled the population using majority of agents as defectors: they will defect when they will face the micro-conflict with the robot. In this scenario more than 60% of agents are defectors (against robot strategy of cooperation) and hence, we see an increasing trend mission costs (see Figure 7).

From Figure 7 we can see that if our robot knew before-

Attribute	Class-Type					
	Male Child	Female Child	Young Man	Young Woman	Old Man	Old Woman
Height (mean)	2.49	2.48	6.06	5.96	5.06	5.01
Height (std. dev.)	1.063	1.0722	0.7851	0.8237	0.8224	0.8426

Table 2: Height attribute from the training-set for the human classes

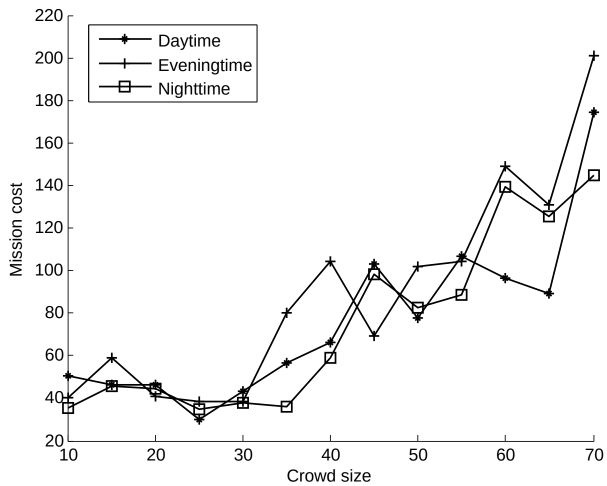

Figure 6: Experimental Results I: Sum of the mission costs.

hand to be polite during all of the micro-conflicts, then the conflicts with those opponents whom defect would result in lower mission cost. This trend can be observed for the "respectful" strategy, where the mission cost still remains low as compared to the "classification strategy". The reason behind this low cost using "respectful" strategy arises due to the fact that no-collision whatsoever may take place as the robot would be cooperative during all of those micro-conflicts.

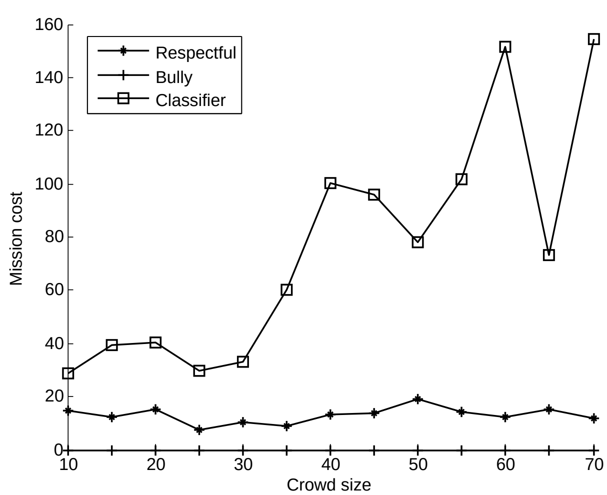

Figure 7: Experimental Results II: Sum of the mission costs.

Related work

The works described in this section are related to the contributions which consider similar or related scenarios as of ours. agent in HiDAC (Pelechano, Allbeck, and Badler 2007) is modeled individually with unique behavior without the need of a modeling using a common centralized controller. The authors use two-tier model where the high-level contributes to the performance of agent behavior such as navigation, learning, decision-making and communication. The low-level module incorporates the perceptions required to model the reactive behavior of agent for collision avoidance. In (Guy et al. 2011) the authors present a unique way to generate heterogeneous crowd behavior using the personality trait theory. Their model was constructed by performing statistical analysis of humans-behavior's collected through user study. For the indexing of human-behavior, the participants were shown the baseline video which was used in comparison with another human-behavior simulated video. Afterwards the participants were asked to describe the differences in behavior as being more or less "Aggressive", "Shy", "Assertive", "Tense", "Impulsive" and "Active". Using this data they mapped the crowd simulation parameters to the perceived behavior's of agents in their simulations. In (Mehran, Oyama, and Shah 2009) the authors use a computer vision technique to analyze the social force model for detection of abnormal behavior in crowd scenes. There unique approach technique was also able to capture the dynamics of the crowd behavior without tracking individual behavior.

An interesting class of game theoretic approaches governing encounters between mobile agents are based on modeling the human adversaries using Stackelberg games. Most of these approaches consider a patrolling strategy, where the goal is not the avoidance of a collision, rather a facilitation of patrolling, where opponent agents actively try to avoid the patrol (Basilico, Gatti, and Amigoni 2009; Amigoni, Gatti, and Ippedico 2008; Paruchuri et al. 2008). This hide and seek game can be modeled as the zero-sum strategic game where the hider selects the cell from the grid, and the seeker seeks (selects) the cell chosen by the hider. Modeling in terms of Stackelberg game with repeated interactions, the strategy selection by follower (hider) is assumed to be optimal based on the leader's (seeker) strategy. The possibility for hider to observe the seeker's strategy before committing its own strategy radically influences the outcome of the game. But as humans deviate from optimal selection due to irrational behavior, its necessary for the leader to incorporate such irrational behavior in its strategic model. In (Pita et al. 2009) three such algorithms are introduced, based on mixed integer linear programming which effectively handles the uncertainties due to bounded rationality and limited observations of adversary. Some of these algorithms are currently being actively deployed (GUARDS(Pita et al. 2011), PROTECT(Shieh et al. 2012)).

Conclusion

In this paper we developed a model for the behavior of an autonomous robot in a social setting where certain types of behaviors incur social costs. While pursuing their own goals or missions, the social agent will occasionally encounter micro-conflicts, where a suitable balance between social and mission costs must be found for each agent. We argued that the right approach for a robot is not a strategy to avoid all of social costs. Instead, the robot must present a consistent strategy against specific types of human interaction partners. This would allow the humans to form a mental model of the robots behavior (a "theory of the robot mind") and adjust their own behavior accordingly. Our future work will be directed in applying the proposed model in more complex interaction scenarios and validate them through various experimental, simulation and survey-based methods.

Acknowledgements: The research reported in this document/presentation was performed in connection with Contract Number W911NF-10-2-0016 with the U.S. Army Research Laboratory. The views and conclusions contained in this document/presentation are those of the authors and should not be interpreted as presenting the official policies or position, either expressed or implied, of the U.S. Army Research Laboratory, or the U.S. Government unless so designated by other authorized documents.

References

Adolphs, R. 2003. Cognitive neuroscience of human social behaviour. *Nature Reviews Neuroscience* 4(3):165–178.

Amigoni, F.; Gatti, N.; and Ippedico, A. 2008. A game-theoretic approach to determining efficient patrolling strategies for mobile robots. In *Proc. IEEE/WIC/ACM Int. Conf. on Web Intelligence and Intelligent Agent Technology*, volume 2, 500–503. IEEE.

Basilico, N.; Gatti, N.; and Amigoni, F. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proc. of the 8th Int. Conf. on Autonomous Agents and Multiagent Systems*, 57–64.

Bhatia, T. S.; Khan, S.; and Bölöni, L. 2012. Towards an operational model for the propagation of public perception in multi-agent simulation. In *13th International Workshop on Multi-Agent Based Simulation (MABS-2012)*, 1–12.

Guy, S.; Kim, S.; Lin, M.; and Manocha, D. 2011. Simulating heterogeneous crowd behaviors using personality trait theory. In *Symposium on Computer Animation, ACM*, 43–52.

Hall, E., and Hall, E. 1969. *The hidden dimension*. Anchor Books New York.

Koenig, S., and Likhachev, M. 2005. Fast replanning for navigation in unknown terrain. *IEEE Transactions on Robotics* 21(3):354–363.

Mehran, R.; Oyama, A.; and Shah, M. 2009. Abnormal crowd behavior detection using social force model. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, 935–942. IEEE.

Nakauchi, Y., and Simmons, R. 2002. A social robot that stands in line. *Autonomous Robots* 12(3):313–324.

Nass, C.; Steuer, J.; and Tauber, E. R. 1994. Computers are social actors. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 72–78. ACM.

Pacchierotti, E.; Christensen, H.; and Jensfelt, P. 2005. Human-robot embodied interaction in hallway settings: a pilot user study. In *IEEE Int. Workshop on Robot and Human Interactive Communication*, 164–171. IEEE.

Paruchuri, P.; Pearce, J.; Marecki, J.; Tambe, M.; Ordonez, F.; and Kraus, S. 2008. Playing games for security: an efficient exact algorithm for solving Bayesian Stackelberg games. In *Proc. of the 7th Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems*, 895–902.

Pelechano, N.; Allbeck, J.; and Badler, N. I. 2007. Controlling individual agents in high-density crowd simulation. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation*, 99–108. Eurographics Association.

Pita, J.; Jain, M.; Ordóñez, F.; Tambe, M.; Kraus, S.; and Magori-Cohen, R. 2009. Effective solutions for real-world Stackelberg games: When agents must deal with human uncertainties. In *Proc. of the 8th Int. Conf. on Autonomous Agents and Multiagent Systems*, 369–376.

Pita, J.; Tambe, M.; Kiekintveld, C.; Cullen, S.; and Steigerwald, E. 2011. Guards-game theoretic security allocation on a national scale. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 3.

Raibert, M.; Blankespoor, K.; Nelson, G.; Playter, R.; et al. 2008. Bigdog, the rough–terrain quadruped robot. In *International Conference of Automatic Control*, 10823–10825.

Sanfey, A. G.; Rilling, J. K.; Aronson, J. A.; Nystrom, L. E.; and Cohen, J. D. 2003. The neural basis of economic decision-making in the ultimatum game. *Science* 300(5626):1755–1758.

Shieh, E.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; DiRenzo, J.; Maule, B.; and Meyer, G. 2012. PROTECT: A deployed game theoretic system to protect the ports of the united states. In *Proc. of the 11th Int. Conf. on Autonomous Agents and Multiagent Systems*, 13–20.

Walters, M.; Dautenhahn, K.; Koay, K.; Kaouri, C.; Boekhorst, R.; Nehaniv, C.; Werry, I.; and Lee, D. 2005. Close encounters: Spatial distances between people and a robot of mechanistic appearance. In *5th IEEE-RAS Int. Conf. on Humanoid Robots*, 450–455.