

MONITORING INTERPERSONAL RELATIONSHIPS THROUGH GAMES WITH SOCIAL DILEMMA

Roman Gorbunov, Emilia Barakova, Rene Ahn and Matthias Rauterberg

Department of Industrial Design, Technische Universiteit Eindhoven, P.O. Box 513, Eindhoven, The Netherlands

Keywords: Evolution, Game theory, Social preferences, Social dilemma, Cooperation, Fairness.

Abstract: In this paper we introduce a method to monitor interpersonal relations through a game with a social dilemma. In the game players can interact with each other through negotiations and by exchanges of resources. To enable the monitoring of interpersonal relations this environment confronts players with specially selected instances of the game, where strategies based on different social factors (like helpfulness or fairness) will enforce different choices in the game. An evolutionary inspired optimization was used to find the games with the special social setting. The special selection of the games helps us to relate the observed actions directly to parameters that model strategies that the players are likely to adopt. Through an estimation of these parameters we are able to observe quantitative differences in the social preferences by different players. Moreover, we demonstrate that players play differently depending on whom they are interacting with. This strongly indicates that the observed playing styles can reveal certain aspects of the relations between players.

1 INTRODUCTION

The recent rapid development of technologies for computer mediated human-human interaction can improve the quality of communication and open up new possibilities for interpersonal interaction. Among the computer based interaction mediums we can distinguish computer games where people can experience new kinds of interactions and decision making problems. The development in this field requires a theoretical support in form of models describing the interaction between humans in computer based environments. Computer based systems for communication and interaction provide a lot of quantitative information about interaction between humans. The understanding and correct interpretation of these data could stimulate the development of new computer based environments for human-human interactions. The data collected by informational systems are interesting from both theoretical and practical perspectives. From a theoretical point of view, these data can be used to better understand how people interact with each other under different conditions. From a practical point of view, the data generated by interaction mediums can be used for an automatic monitoring of interpersonal relations. Social atmosphere in different kinds of task oriented teams is very important for the performance of the team (Rauterberg et al., 2008).

In this context, the methods for automatic monitoring of social relations may be useful for detection and prevention of social problems on earlier stages. In the present work we focus on the usage of computer based strategic multiplayer games for monitoring of interpersonal relations in small groups of people.

The computer games are of particular interest in the context of the monitoring interpersonal relations since they cover many different aspects of real life interactions between people. Collaboration is one of the aspects of computer games, which has a particular interest for monitoring interpersonal relations in goal oriented teams. Many games designers and producers are currently exploring the addition of cooperative patterns within their games. A special attention was given to identification of design patterns of cooperative game (El-Nasr et al., 2010). Limited resources are considered as one of the design patterns of collaborative games. In the present work we pay a special attention to the design and implementation of a game which utilizes this pattern to induce cooperation within the game.

In behavioral game theory it has been demonstrated that people do not play like game theoretical rational players. In particular, it has been shown that the behavior of humans in game environments can be better described if social factors such as altruism, self-interest and fairness are taken into account (Gal

and Pfeffer, 2006; Gal et al., 2004). The influence of social factors on the game strategies of humans makes the behavior of human players dependent on social relations with other players. In particular, it has been demonstrated that a humans way of playing depends on whom they are playing with: "friends" and "strangers" are treated differently (Marzo et al., 2004). Moreover, it has been shown that humans treat computer agents differently depending on the game strategies embedded into these agents (van Wissen et al., 2009). In particular, collaborative, altruistic behavior in agents induces more collaborative and altruistic behavior in human subjects.

Different models have been developed to describe human-human interaction in game environment. For example, relations of human subjects to other players can be modeled by retrospective and prospective thinking which reflects humans' memories about past behavior of other players as well as certain expectations about future behavior of these players (Gal and Pfeffer, 2007). With this approach the reciprocal behavior of humans can be modeled. It has been demonstrated that human players do not base their decisions solely upon the options they have. Instead, humans also reason about the preferences of other players in the game (Ficici and Pfeffer, 2008a; Ficici and Pfeffer, 2008b).

The models describing human-human interactions in game environments were proven to improve performance of computer agents when interacting with human subjects. For example, it has been demonstrated that agents using purely axiomatic rules are outperformed by human like computer agents which decide whether to help other agents based on partial information about the environment and their partners' plans (Kamar et al., 2006). Similarly, computer agents which, like humans, were able to estimate the degree of helpfulness of other agents and to adopt to it, outperformed agents playing classical game theoretic strategies (Talman et al., 2005). The fact that behavior of human players depends on relations to (beliefs about) other players makes games into an attractive tool for automatic monitoring of interpersonal relations.

In the presented work we use the Colored Trail (CT) game as a tool for monitoring interpersonal relations. Colored Trails (CT) is a game designed to enable investigation of properties of decision making strategies in multi-agent situations of varying settings and complexity (Grosz et al., 2004). Different variations of the game have been used to study human-human and human-agent interactions (van Wissen et al., 2009; Gal and Pfeffer, 2006; Talman et al., 2005; Grosz et al., 2004; Gal et al., 2004; Ficici and

Pfeffer, 2008a; Marzo et al., 2004; Kamar et al., 2006; Gal and Pfeffer, 2007; Hennes et al., 2009; Ficici and Pfeffer, 2008b; Gal et al., 2005).

2 METHODS

2.1 Colored Trails Game

The colored trails is a multiplayer negotiation game that can be classified as a computer version of a board game involving strategic and logical reasoning. The game provides a flexible setting which allows researcher to vary the complexity of the game, the level of dependence between players, and the accessibility of information about the state of the game. The colored trails game has been used in the MARS-105 experiment to investigate its potential in monitoring interpersonal relations in small isolated groups of people performing long term missions (Hennes et al., 2009; Voynarovskaya et al., 2010; Rauterberg et al., 2008).

The CT game resembles real life situations in which people have different goals and need some resources to reach these goals. In the initial state of the game players can be in different situations and, as a consequence of that, they require different kinds of resources to reach their goals. In other words, resources of different kinds can have different values to different players. A redistribution of the resources can be done as a result of negotiations between players. The game is interesting for analysis since it contains both competitive and collaborative components.

More specifically, the game is played on a rectangular board composed of colored squares (see Figure 1). Colors are chosen from a fixed set of colors. In the




Figure 1: Example of the colored trails game on the response phase.

beginning of the game every player is placed on one of the colored squares of the board. Additionally to that one square is assigned to be the goal-square. In

the beginning of the game every player also receives a set of colored chips (resources). The colors of the chips are taken from the same set as the colors of the board squares. Players can move on the board using their chips. In more detail, they can make a horizontal or vertical move to one of the neighboring squares if they have a chip of the same color as the new square. By making a step on the board a player irreversibly spends a chip. The goal of the player is to move as close as possible to the goal-square, spending a minimum of chips. Before making their moves players are allowed to exchange some of their chips with another player. Any exchange of chips is possible if both participants of the exchange agree to do this exchange.

During the previous research on monitoring of interpersonal relations during the MARS-105 isolation experiment the game was played by three persons with predefined roles (Hennes et al., 2009). Two of the subjects played the role of proposer, in which they could propose a chips exchange to the third player who played the role of the responder. The responder, in his/her turn could accept one of the two propositions or reject both of them. The accepted chips exchange was realized so that two of three players got new sets of chips. After that the game engine calculated and performed the best possible move for every player. In the present study we use a modified version of the game which is also used in the new MARS-500 isolation experiment. The following modifications have been proposed:

- Every subject plays the role of a proposer with the freedom to choose to whom to make a proposition. In this way we increase the number of proposers per game (from two to three) and, as a consequence, we improve the statistics of our experiment. Moreover, giving a freedom to chose a partner we get additional information for the analysis.
- As a consequence of the first modification, every player can potentially receive a proposition and play the role of a responder. In this way we increase the number of the responders per game (from one to two or three).
- In the modified version of the game responders can have from one to three propositions and, in this way, they experience a broader range of situations in comparison to the previous version of the game. In this way we can access more aspects of the responders behavior.

2.2 Generation of Games

The game is given by the state of the board (how it is colored), the positions of the three players, the loca-

tion of the goal, and the three sets of chips owned by the three players.

Any exchange of chips between two players is characterized by payoffs earned by the proposer and responder. Different exchanges can correspond to identical pairs of payoffs. We can say that a set of exchanges represents one "option" if all exchanges from the set are characterized by the same pair of payoffs. In this context, the situation in which a proposer needs to make a decision is characterized by a set of available options. Options available to a proposer can be visualized as dots (a "cloud" of options) in a two dimensional space of payoffs in which the x- and y-axis correspond to the payoffs of proposer and responder, respectively (see Figure 2). Among the available options we count the one which corresponds to no exchange (a "default" option).

We assume that proposers value options based on pairs of payoffs associated with the option. A proposer, searching for "the best" exchange, should try to maximize his/her own payoffs (p) and, at the same time, he/she should also try to maximize the payoff of the responders (r) since it determines the probability that the proposed exchange will be accepted by the responder. For this reason, any option which increases both his/her own payoff and the payoff of the responder with respect to another option should be considered as a "better" option. We will say that the first option (characterized by its payoffs (p_1, r_1)) dominates the second option (p_2, r_2) if $p_1 \geq p_2$ and $r_1 \geq r_2$ and $(p_1 \neq p_2$ or $r_1 \neq r_2)$.

An option which is not dominated by any other option is called a Pareto optimal option. In a general case there can be many Pareto optimal options (forming a Pareto set). A rational proposer is expected to choose one of the options from the Pareto set. So, the interesting question is what algorithm players use to choose from a Pareto set. In any case it should constitute some kind of a balance between the benefit of the proposer and the benefit of the responder.

It is very well known that people can interact with each other in fair and/or cooperative way. Fairness is usually understood as an equal (or as close as possible) redistribution of goods. Cooperativeness, in its turn, is usually oriented to the performance of a team (not of a single player). In other words, it can be understood as a strategy which maximizes the total payoff of the team. So, it is obvious that the two above mentioned strategies maximize different utility functions. However, players playing with different strategies will not always make different choices since it can happen that there is an option which provides both the maximal total payoff and, at the same time, the most fair distribution of payoffs. Keeping this in mind

we intentionally generated games in which collaborative and fair strategies lead to different non overlapping set of options. In this way we can clearly distinguish between fair and collaborative behavior in players. The last property is important since we could expect that people tend to play more fairly or more collaboratively depending on the interpersonal relations between them. Moreover, in task oriented groups of people the collaborative behavior is very important. So, we need to have a reliable way to identify collaborative interactions in the game. Additionally to that we have assumed that there are different kinds of fair behavior. In more detail, when people judge different options they might care more about what everybody gets additionally to what everybody already had or, alternatively, they might care more about what everybody will have in the end. As an example, people might think that those who have less in the beginning deserve to get more extra. Keeping this in mind we generated games in which the set of options providing fair gain does not overlap with the set of options that provides a fair final result.

The utility functions corresponding to the three above introduced strategies can be expressed mathematically in the following way:

- The cooperativeness : $p + r$, where p and r are final scores of proposer and responder, respectively.
- The fairness of the final score: $|p - r|$
- The fairness of the gain from the exchange: $|(p - p_0) - (r - r_0)|$, where p_0 and r_0 are payoffs obtained by proposer and responder in the case where no exchange happens (default payoffs).

An example of the payoff space is shown in figure 2. As we can see in the figure, the game provides three non overlapping sets of options. Every set corresponds to one of the predefined strategies: the collaborative strategy maximizing the total payoff of proposer and responder, the strategy leading to a fair gain from the exchange and the strategy leading to a fair final score.

We have found that only a very small percentage of the games fulfills the above introduced criteria. So, a random search of the desired games would be too time-consuming. To overcome this problem an evolutionary search was performed. In more detail, we started from a game generated randomly. In other words, the state of the field, the positions of players and the goal as well as redistribution of chips was done randomly. The only restriction that we applied is that every player can have from three to five chips. Then we started an iterative process in which either the state of the field or the position of the player or goal, or a set of chips was modified. During the mod-




Figure 2: Options available to the proposer in the space of payoffs. Options corresponding to different strategies form non-overlapping sets.

ification of the field we randomly changed the color of a randomly chosen square. Changing a position of a player or goal we made one horizontal or vertical step to one of the neighboring squares (under the obvious restriction that elements of the game should stay on the field). Additionally to that we did not allow two players to occupy the same square. Modifying set of chips we randomly chose a player and added or removed a chip of a randomly chosen color. Doing that, we kept the restriction that the number of chips owned by one player should be between three and five.

On every step we calculated the overlap between sets of options representing different strategies. The size of the overlap was calculated for every potential pair of players and the total size of the overlap was taken into account. In this way we could guaranty that the game will be interesting independent of who negotiates with whom. A considered mutation was accepted only if it decreases the overlap. The mutation process was continued until no overlaps were found for any pair of players.

In our study we have used a four-by-four board with a palette containing five colors. With this setting we could have a reasonable combinatorial complexity of the game (not too complex, not too trivial). The score was calculated in the following way. Every step between the goal and the final position of the player caused a penalty equal to 25 points. For every chip remaining after the move each player received a reward equal to 10 points.

In our web based experiment 27 participants took part. 9 teams consisting of 3 players were formed. Teams played different numbers of games (ranging from 4 to 18). In total 93 games were played. The teams were formed by self-selection and mutual agreement of team members. The player in a team knew each other since the main intention of our study

was to develop a method for monitoring interpersonal relations.

We need to notice that with the present work we do not try to make any conclusions about a particular social group. We also do not try to present any results that could be dependent on social factors (like race, gender, education). In this context we do not need to worry if the group of participants is large enough to be representative. What we do want to do is to show that different players adopt different strategies that are also dependent on their partners.

3 RESULTS

3.1 Population of Strategies

As has been mentioned in the previous section, we generated games in such a way that the three above defined strategies correspond to three non-overlapping sets of options. In other words, if a proposer uses one of the four predefined strategies, we can determine the strategy if we know what proposition he/she made.

It's possible that a player chooses an option which does not belong to any of the four sets representing the predefined strategies. If this happens we check if the chosen option belongs to the Pareto set since any reasonable strategy leads to a Pareto optimal solution. If it is in the Pareto set (and not in one of the predefined sets) we can assume that player either tried to find a balance of the four above defined criteria or that he/she used other criteria to evaluate options. If the chosen option is not in the Pareto set, we can conclude that the player made a mistake (for example, he/she miscalculated payoffs corresponding to some exchanges or he/she did not see all available exchanges). The portion of non Pareto optimal propositions can be used as a measure of the complexity of the game. If all propositions are Pareto optimal, the game is trivial. Likewise, if a majority of propositions are not Pareto optimal, the game is too complex. A game should maintain an appropriate level of challenge in order to be entertaining (Yannakakis and Hallam, 2006; van Lankveld et al., 2010).

The analysis of the collected data has shown that 29% of the propositions were not Pareto optimal. Among the Pareto optimal propositions 41% of the propositions did not belong to the three pre-defined set of strategies (see Table 1). The cooperation was proposed in 29% of the Pareto optimal propositions. In 18% and 12% of the Pareto optimal propositions the proposers chose to be fair with respect to the gain and final score, respectively (see Table 2).

Table 1: Portion of Pareto optimal propositions.

Non Pareto	29%
Pareto	71%

Table 2: Population of different strategies on the Pareto frontier.

Cooperation	29%
Fair Gain	18%
Fair Final Score	12%
Mixed Strategies	41%

3.2 Angle of Benefit

A direct way to describe a proposition is to give a relation between the points gained by proposer and responder in the case where the exchange is realized. This relation can be described by an angle given through the following function:

$$\alpha = \arctan\left(\frac{r - r_0}{p - p_0}\right). \quad (1)$$

The introduced angle has a clear interpretation. It gives direction of the line connecting the default option with the proposed one. In particular an exchange which is equally beneficial to proposer and responder is characterized by angle equal to 45°. A value between 0° to 90° corresponds to an exchange which is mutually beneficial to proposer and responder. The angles in the range between 90° and 180° correspond to propositions which are not beneficial to the proposers but are beneficial to the responder (a proposer provides help to a responder). In a similar way, the angles between -90° and 0° describe situations when proposers ask for help (i.e. a proposer offers exchanges in which he/she gain points and responder loses his/her points).

As in the previous section, we analyze how the distribution of the introduced parameter is influenced by the decisions of the proposer. In Figure 3 we show the distribution of the introduced angle over all available options as well as over options chosen by the proposers.

In this figure we can see a clear and expectable influence of the proposers decisions on the distribution of the angle. Proposers tend to choose options which are beneficial to both participants of the exchange (proposer and responders). We can see that options which are more beneficial to the proposers but still beneficial to the responder (angles between 0° and 45°) slightly dominate options which are more beneficial to the responders but still beneficial to the proposers (angles between 45° and 90°). We also can see that proposers ask and propose help. It is interesting to notice that proposers propose help significantly




Figure 3: Distributions of the angle of benefit calculated over all available options as well as options chosen by proposers.

more frequently than they ask for help. It also should be noticed that proposers practically never propose (or ask for) help for which the cost of the help exceeds the benefit from the help.

3.3 Fairness, Cooperation and Symmetry of Cooperation

We found out that decisions of different proposers generate essentially different distributions of the angles. To describe the shape of the distribution we split the whole range of the angles into three regions corresponding to quantitatively different situations. The first region corresponds to situations in which proposers propose help (i.e. an offer which is beneficial to the responder and requires some sacrifice from the proposer). The second region corresponds to propositions which are beneficial to both proposer and responder. The third region corresponds to propositions which are beneficial to the proposer and require some sacrifice from the responder (i.e. the proposer asks for help). For every proposer we calculate the number of propositions falling into every of the above defined regions. We will use the following notation: n_{ph} and n_{ah} indicate the number of proposition in which the proposer proposed or asked for help, respectively. n_{mb} is the number of mutually beneficial propositions. Further we make use of the fact that the strategy of a proposer is described by the relative populations of every region. In this way we can decrease the number of parameters describing the behavior of proposer (from three to two). In other words we can remove the degree of freedom which represents the number of played games while preserving the information about behavior of the proposers. Accordingly, we introduce

the following two parameters:

$$v_{cl} = \frac{n_{ah} + n_{ph}}{n_{mb} + n_{ah} + n_{ph}}, \quad (2)$$

”level of cooperation”, gives the portion of propositions in which proposer proposed or asked for help.

$$v_{gr} = \frac{n_{ph}}{n_{ah} + n_{ph}}, \quad (3)$$

”give ratio”, indicates how frequently a proposer proposes help in all those situations when help is involved.

We assume that the values of the two introduced parameters (v_{cl} and v_{gr}) depend on the identity of the proposer as well as on who is the responder. So, the behavior of every proposer can be described by two pairs of the considered parameters (corresponding to two responders). An example of the diagram visualizing interpersonal relations in a team is shown in figure 4. Every player is represented by two points connected by a red line. Blue lines are used to indicate a responder corresponding to a point representing behavior of a proposer. In other words, red lines represent subjects and blue lines represent interpersonal relations between the subjects. As we can see in the figure, players exhibit significantly different behavior in terms of the two above introduced parameters. Moreover, we can see that the behavior is different depending on who is the responder. In other words, we have demonstrated that players adopt different strategies depending on who they interact with. So, we can conclude that the proposed method provides quantitative measurements of interpersonal relations.

We would like to emphasize that propositions made by a given proposer to a given responder are not only conditioned by utility function of the proposer but also by his/her believes about the utility function of the responder. Within this work we do not try to study the mechanism of making decisions based on utility function of the proposer and his/her believes about the utility function of the responder. Instead, we propose a method to measure and describe the phenomenon with the intention to access interpersonal relations using computer games.

In terms of the introduced diagram the healthiest relations in the group would be represented by a point located on the very top of the figure (the highest cooperation) and in the middle of the horizontal direction (proposer proposes help as frequent as he asks for it).

The clear meaning of the introduced parameters provides a simple interpretation of the diagrams. For example, in the figure 4 we can see that one of the players (the C_2 -player) shows the maximal level of cooperation. However, the give-take ration of this player is too high indicating that this player almost

always proposes help. For this player we could recommend to ask for help when it is more beneficial for the whole group to do so. In the same team we can see that the second player (B_2 -player) exhibits different behavior depending on who is the responder. This player is more cooperative with the cooperative player (C_2 -player) and less cooperative with the least cooperative player (A_2 -player).




Figure 4: Strategies of players in team 2 given by values of the "give/take" and "help" ratios.

The limited number of the games raises the question of the accuracy of the presented parameters. To address this question we have numerically calculated probability density function (PDF) of the proposer models for every observed population of decisions (every model is given by values of the give and help ratios). In more details, we have randomly generated models of proposers. Every model was given by the probabilities of the three different kinds of decisions (propose help, ask for help, propose mutually beneficial exchange). Every probability was generated by the random number generator implementing uniform distribution between 0 and 1. The generated probabilities were normalized to make the sum of the three probabilities to be equal to 1. Every model (random decisions generator) was used to generate the same number of decisions as in the considered observed set of decisions (made by a player). The model was accepted if the generated and observed decisions give the same population of the three kinds of decision. In this way, for every observed population of the decisions (given by the give and help ratios) we have generated 1 million models that were used to construct the PDF of the models. For every generated PDF we have calculated the cutoff value such that integration of the part of the distribution with the values larger than the cutoff should give 0.9. In other words we found the region around the maximum of the PDF such that the probability for the real model to be

within this region is equal to 90%. These regions are shown in the figure 4 with the same colors as those of the observed points.

As we can see on the figures the difference between the observed values is very unlikely to be produced purely by the inaccuracy of the measurements.

4 CONCLUSIONS

In the present work we have proposed a way to generate game instances with a special set of options in the space of payoffs. By confronting players with these game instances we are able to access human strategies based on their choices in these games. We have formulated a set of parameters to characterize all options available to the players. With the introduced parameters we could draw some conclusions about a player's strategy even though the choices of the player did not belong to any of the predefined sets of strategies and were not even Pareto optimal. In particular, in this way we have a method to extract useful information about social preferences of players even in situations when not all options are noticed by the player because of the combinatorial complexity of the game. Additionally to that we have introduced two ratios which describe behavior of players. The help ratio indicates how a player chooses between fair and collaborative options. In that way we got a direct measure of the level of collaboration. The give/take ratio is used to describe the ways of collaboration. In particular we can see if collaboration between people is exploiting or sacrificing (or balanced between the two extremes). Using the introduced ratios we were able to see clear differences in behavior of the different players. Moreover, we could show that different players play differently depending on who is their partner. In that way we have demonstrated a possibility to access characteristics of interpersonal relations with the usage of collaborative games.

ACKNOWLEDGEMENTS

The research reported in this paper is supported by NWO User Support Program Space Research. The project number is ALW-GO-MG/07-13. The authors would like to thank the management and the work group of the Mission Execution Crew Assistant (MECA) project, and particularly Mark Neerincx, for the support and collaboration. We also are grateful to Karl Tuyls and Daniel Hennes for their help and support. We would like to thank our partners from the

Moscow Institute of Biomedical Problems, and particularly Vadim Guschin, for support with preparation and conduction of experiment within the MARS-500 study. We also would like to thank Jeffrey Goldstein, Peter Suedfeld, Jack van Loon and Kees Nieuwenhuis for their support during the project. We are grateful to Natalia Voynarovskaya for her support in preparation and conduction of the presented experiment.

REFERENCES

- El-Nasr, M. S., Aghabeigi, B., Milam, D., Erfani, M., Lameman, B., Maygoli, H., and Mah, S. (2010). Understanding and evaluating cooperative games. *Proceedings of the 28th international conference on human factors in computing systems*, pages 253–262.
- Ficici, S. G. and Pfeffer, A. (2008a). Modeling how humans reason about others with partial information. *Proceedings of the 7th international joint conference on autonomous agents and multiagent systems*, 1:315–322.
- Ficici, S. G. and Pfeffer, A. (2008b). Simultaneously modeling humans’ preferences and their beliefs about others’ preferences. *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, 1:323–330.
- Gal, Y., Grosz, B. J., Kraus, S., Pfeffer, A., and Shieber, S. (2005). Colored trails: a formalism for investigating decision-making in strategic environments. *Proceedings of the 2005 IJCAI workshop on reasoning, representation, and learning in computer games*, pages 25–30.
- Gal, Y. and Pfeffer, A. (2006). Predicting people’s bidding behavior in negotiation. *Proceedings of the 5th international joint conference on autonomous agents and multiagent systems*, pages 370–376.
- Gal, Y. and Pfeffer, A. (2007). Modeling reciprocal behavior in human bilateral negotiation. *Proceedings of the 22nd national conference on artificial intelligence*, 1:815–820.
- Gal, Y., Pfeffer, A., Marzo, F., and Grosz, B. J. (2004). Learning social preferences in games. *Proceedings of the 19th national conference on artificial intelligence*, pages 226–231.
- Grosz, B. J., Kraus, S., Talman, S., Stossel, B., and Havlin, M. (2004). The influence of social dependencies on decision-making. initial investigations with a new game. *Proceedings of the 3rd international joint conference on autonomous agents and multiagent systems*, 2:782–789.
- Hennes, D., Tuyls, K. P., Neerincx, M. A., and Rauterberg, G. W. M. (2009). Micro-scale social network analysis for ultra-long space flights. *The IJCAI-09 workshop on artificial intelligence in space, Pasadena, California, USA*.
- Kamar, E., Gal, Y., and Grosz, B. (2006). Incorporating helpful behavior into collaborative planning. *Proceedings of the 8th international conference on autonomous agents and multiagent systems*, pages 875–882.
- Marzo, F., Gal, Y., Grosz, B. J., and Pfeffer, A. (2004). Social preferences in relational contexts. *Proceedings of the 4th conference in collective intentionality*.
- Rauterberg, M., Neerincx, M., Tuyls, K., and van Loon, J. (2008). Entertainment computing in the orbit. *International federation for information processing*, 279:59–70.
- Talman, S., Gal, Y., Hadad, M., and Kraus, S. (2005). Adapting to agents’ personalities in negotiation. *Proceedings of the 4th international joint conference on autonomous agents and multiagent systems*, pages 383–389.
- van Lankveld, G., Spronck, P., van den Herik, H. J., and Rauterberg, M. (2010). Incongruity-based adaptive game balancing. *Advances in computer games*, 6048:208–220.
- van Wissen, A., van Diggelen, J., and Dignum, V. (2009). The effects of cooperative agent behavior on human cooperativeness. *Proceedings of the 8th international conference on autonomous agents and multiagent systems*, 2:1179–1180.
- Voynarovskaya, N., Gorbunov, R., Barakova, E., Ahn, R., and Rauterberg, M. (2010). Nonverbal behavior observation: Collaborative gaming method for prediction of conflicts during long-term missions. *ICEC 2010, LNCS, 9th international conference on entertainment computing*, 6243:103–114.
- Yannakakis, G. N. and Hallam, J. (2006). Towards capturing and enhancing entertainment in computer games. *Advances in artificial intelligence*, 3955:432–442.