

MICRO-SCALE SOCIAL NETWORK ANALYSIS FOR ULTRA-LONG SPACE FLIGHTS

D. Hennes¹, K. P. Tuyls¹, M. A. Neerincx², and G. W. M. Rauterberg¹

¹*Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands*

²*Delft University of Technology, P.O. Box 5031, 2600 GA Delft, The Netherlands*

ABSTRACT

This paper proposes to leverage the mathematical means of game theory to analyze on-board social crew dynamics. We describe how game theory facilitates capturing the essence of interactive decision making, thereby raising the potential for a fully automated and unintrusive monitoring and diagnosis tool. Finally, we present preliminary findings based on the base-line data collection and the first phase of the ground based Mars-500 isolation experiment.

Key words: Game theory, social network analysis, Mars 500.

1. INTRODUCTION

A manned mission to Mars is considered to be one of the ultimate endeavors of mankind. Though not feasible with state-of-the-art technology, various space agencies are exploring plans to make a human mission to Mars a viable venture. Recent proposals anticipate launch dates as early as 2035. Besides highly reliable technical equipment, crew members on their mental and physical peak are of utmost importance in order to assure success. The isolated space environment during ultra-long space flights affects a number of physiological, psychosocial and mental processes critically involved in human performance. It is vital to a mission's success to understand these psychological limits. Past experiences have shown that the mental health of the crew can have a great effect on the success or failure of a mission (e.g. long-term isolation can lead to sleep deprivation, depression, irritability, anxiety, impaired cognition, and even hostility). Latent and overt stress factors are mental strain, lack of capability to rescue crew members, isolation, monotony, tedium of life aboard an autonomous shuttle, and interpersonal problems. These issues develop slowly over time and are very difficult to detect and remedy for observers on the ground. Computer-interactive intervention programs show promise for adaptation to long-duration space flights. Software-based delivered prevention and intervention information may be more comfortable for crew members than disclosing highly per-

sonal information to others [7, 8]. Programs applying cognitive-behavioral and self-help instructions show the potential to be as effective as face-to-face intervention if dealing with mild to moderate depression, anxiety, and other types of psychopathology [2, 6, 7, 11]. In order to trigger the use of such countermeasures effectively one needs to assess the situation onboard continuously, thus constituting the need for an automated tool to measure the individual mental capacity as well as interpersonal dynamics.

Interpersonal relationships can be captured and analyzed using the representation of a social network; persons are represented by nodes while links between nodes describe or label the relationship features. Considering the small number of crew members during ultra-long space flights, we speak of "micro-scale networks". In order to determine the characteristics of links (e.g. cooperation or fairness) we make use of a strategic game, called Colored Trails - a three-player negotiation game [3]. It is played on a board of 4 times 4 squares, colored in one of five colors. Each player possesses a piece located on the board and a set of colored chips. A colored chip can be used to advance a player's piece to an adjacent square of the chip's color. The general goal is to position pieces onto or as close as possible to a common goal location. Colored Trails has two distinct roles: two proposers and one responder. Proposers can suggest a chip exchange to the responder. The responder can accept exactly one - or no proposal at all. All players are aware of the board state; proposers only have information about their own chip set and the one of the responder; the responder can see all chip sets. A single game of Colored Trails is played as a one-shot game, meaning that proposers can only make a single proposal and the responder can only accept or reject proposals and is not allowed to make a counter offer. Furthermore, players do not actually have to move their pieces step-by-step toward the goal location. Once the responder has reacted to the proposals (by either accepting one or rejecting both) the chips are exchanged according to the winning proposal, or stay fixed if the responder rejected both proposals. With the post-exchange chip sets to their disposal, players individually advance on a "colored trail" toward an optimal location and score accordingly.

In this work, we use the mathematical foundation of game theory and evolutionary game theory (EGT) to model and




Figure 1. Illustration of the Mars-500 facilities (Image credits: IBMP).

study the crew member interactions in the game of Colored Trails. Traditional game theory mainly analyzes rational decision making, while EGT emphasizes adaptation and learning. Both the ability to think rationally and to adapt to changing situations are deciding factors for mission success.

The experiment presented here is part of the Mars-500 study to be carried out at the Institute for Biomedical Problems (IBMP) in Moscow. The European Space Agency and the Russian Academy of Sciences jointly plan and conduct this study in order to simulate a manned mission to Mars. It provides the unique opportunity to study crew member interactions while collecting data about the human subjects' health and performance during experimental isolation. The confinement study imitates all key peculiarities expected to be present during future missions to Mars (i.e. ultra-long duration, need for autonomy, affected communication due to signal delay, and limited stock of expendables). A pretest of 105 days has recently launched; the full scale experiment with a duration of 520 days will start in the last quarter of 2009.

The rest of this article is organized as follows. Section 2 provides background information about Mars 500, the game theoretical framework as well as Colored Trails. In Section 3 we discuss related work. Section 4 outlines the methodology used to analyze data collected during isolation and Section 5 presents preliminary findings. Section 6 concludes this article with discussion and plans for future work.

2. BACKGROUND

This section addresses the specifics of the Mars 500 program in greater detail (Section 2.1). Besides general setting and operational conditions of Mars-500, this section

also provides the particular prerequisites of the here presented research. Therefore, we summarize required background knowledge from the fields of game theory and evolutionary game theory by means of a simplified example (Section 2.2) and hereafter consider more complex interactions in the game of Colored Trails (Section 2.3).

2.1. Mars-500

In 2004 the Institute for Biomedical Problems (IBMP) in Moscow and the European Space Agency have started to plan a full-scale ground based simulation of a manned mission to Mars. Such a full scale mission requires between 520 to 700 days of isolation. Referring to the lower end of this time frame the initiative was named *Mars 500*.

Ground based isolation studies, including bed rest studies as well as human missions in low Earth orbit provide crucial experience and insights to evaluate the feasibility of a manned Martian mission and to list general requirements for long duration space flights. Human missions beyond the Earth orbit are affected by radiation hazards, the impact of microgravity on human physiology and various medical and psychological issues. While the former two rely on research onboard the International Space Station (ISS) and future missions to Moon, the affect of long term isolation can be effectively addressed by ground based studies such as the Mars 500 program.

The goal of the Mars 500 study is to gather data, knowledge and expertise required to prepare a real mission to Mars. Hence, all key peculiarities expected to be present during future missions to Mars are reflected:

- ultra-long duration
- need for autonomy

- affected communication due to signal delay
- and limited stock of expendables.

This ensures that psychological and physiological impacts of isolation through such an extended period of time are observed as close to reality as possible.

A crew of six candidates (four Russians and two Europeans) are sealed inside the facilities of the Institute for Biomedical Problems in Moscow. This crew is chosen to aggregate a very diverse field of expertise and skills, including knowledge in the field of aviation, biology, engineering, medicine and physiology. A initial 105-day isolation period has recently started in March 2009 and is currently underway. Once completed successfully, a full 520-day study is planned to launch in the first quarter of 2010 in which all elements of a Martian mission are being simulated: transit to Mars, orbiting the planet, landing maneuvers, Mars surface operations and return to Earth. The facilities at IBMP accommodate for these different simulation episodes with a number of interconnected isolation chambers. An illustration of this setup is depicted in Figure 1. Part of the isolation facilities resemble the hermetically sealed mock-up space craft: a habitable module, a utility module and the medical unit. Furthermore a landing ship and the Martian surface itself are being simulated.

Throughout the isolation period and all episodes of simulation, the crew needs to be self-reliant to a great extent. This includes monitoring life support, control resource consumption, maintaining technical equipment as well as performing a number of scientific experiments in which the participants act as subjects themselves. In this work we present an experiment where crew members encounter in strategic game interactions while data is being gathered about their interpersonal dynamics. The next section describes the mathematical framework of game theory and evolutionary game theory that allows to quantify these interactions.

2.2. Game theory and evolutionary game theory

Traditional game theory studies strategic interactions where multiple players act rational in order to maximize their expected payoffs. These players can for instance represent companies in the economical area, single individuals in the psychological or social science domain as well as computer agents situated in a multi-agent setting.

To provide an intuitive understanding of the methodology of game theory we shall first consider a particular strategic encounter. The *Prisoners' Dilemma* is a prevalent example to showcase game theoretical modeling and analysis practice. This chapter first introduces the classical predicament of two suspects, while gradually extending the problem to cover iteration, evolution and finally multi-state settings.

The dilemma presents itself in the following situation. Two subjects are under arrest. The police cannot provide enough evidence to bring in a severe accusation. Thus the two prisoners are separated for interrogation and are both confronted with the same choice: either confess the crime and hence *defect*, or to deny any accuse and hence *cooperate* with the other suspect. If one suspect testifies while the other remains silent, the defector is released while his loyal fellow is convicted to serve full sentence. If both defect, both receive the full sentence. If both deny all accuses, they serve a short sentence for a minor charge. How should the prisoners choose rationally in order to minimize penalty?

Seemingly, the best option is to cooperate, both parties only serve a short sentence and are set free after a relatively short time. However, this intuition suggests a non-rational decision. Game theory provides the tools to reason about such a strategic situation and identify the rational choice. Let us consider the options of a suspect under the following set of assumptions. A suspect is self-interested and therefore only cares about his own well-being (or degree of punishment in that perspective) and thus does not profit from lowering the sentence of the other. Furthermore, any communication between suspects is inhibited and both are ensured that the other is provided with the same information and faces the same choice. Both have the same option to choose exactly one of the two actions *defect* or *cooperate*. In other words, all action sets and corresponding payoffs are equivalent, therefore we call such a game symmetric. Since no communication is allowed we might as well assume that players act simultaneously. The payoff of each player is not solely dependent on its own choice but rather on the collective joint action. Suppose a player chooses to *cooperate*. If the other one cooperates as well the player is convicted to serve a short sentence; if the other one betrays him he serves the full sentence. If he instead chooses to *defect*, he is set free in the former case while as well serving the full sentence in the later case. We observe that cooperating is dominated by defecting, i.e. for each adversarial response, *defect* yields at least as much payoff as *cooperate*. Thus players will make the rational decision to mutually defect and hence end up serving both the maximum sentence. While intuition suggests otherwise, in the perspective of rationality the suspects face an unwinnable situation. This paradox clearly shows the difficulty in analyzing non-zero-sum situations (i.e. the gain of one party is not necessarily a loss of another).

We now revisit our initial example, the *Prisoners' Dilemma*, in order to bridge the gap between traditional and *evolutionary game theory*. The classic dilemma situation can be extended in different ways - one of them is to play the game for multiple stages. In the *Iterated Prisoners' Dilemma* the two suspects encounter the same predicament over and over again. Clearly, this can affect the equilibrium strategies. If the game is repeated infinitely often, players can obtain the Pareto efficient equilibrium reward by invariably cooperating. If however the game is played N times and this information is known to the players, the rational strategy is to always defect which

can be shown by backwards induction [5].

In [1] a tournament version of the *Iterated Prisoners' Dilemma* is proposed. Each player has to choose a strategy, possibly depending on memory of previous encounters. Interactions between pairs of individuals occur on a probabilistic basis while natural selection favors individuals that have performed better in the past. The derived model as well as results of the tournament show how cooperation based on reciprocity can evolve in a mechanism that advocates self-interest. The simplest deterministic strategy, called *Tit-for-Tat*, has won the tournament. *Tit-for-Tat* starts by cooperating and hereafter responds by playing the action the opponent has played in the last stage.

The primary objective of evolutionary game theory is to model the dynamics of strategy changes in iterated games such as the *Iterated Prisoners' Dilemma*. In particular, interactions are modeled using biological-inspired genetic operators such as natural selection and mutation. At each time step two randomly matched individuals from a population play a particular pure strategy while their payoff determines the replication success of the represented strategies. The process of population change is captured by the *replicator dynamics*.

2.3. Colored Trails

While the prisoners dilemma is a good example to introduce traditional game theory and evolutionary game theory it is very limiting and not suitable to capture more complex inter-social dynamics as present during ultra-long space flights. For the first iteration of the Mars-500 experiment, the 105-day pretest, we are primarily interested in games that feature the following properties:

- Simple enough for analysis
- Rich enough to reflect features of real life interactions
- Grounded in a situated task domain
- Strategic, i.e. partial information to promote complex reasoning
- Suited to measure social factors such as fairness

The Colored Trails framework developed at Harvard University, School of Engineering and Applied Sciences serves all these requirements [4]. Specifically, we use a three-player negotiation variation [3]. It is played on a board of 4 times 4 squares, colored in one of five colors: blue, green, red, yellow and grey. Each player possesses a piece located on the board and a set of colored chips. A colored chip can be used to move a players piece to an adjacent square (diagonal movement is not allowed) of the same color. The general goal is to position pieces




Figure 2. Colored Trails board

onto or as close as possible to a goal location indicated by a flag. Although, there is a single goal, each player receives points purely based on its own performance. Colored Trails has three distinct roles: proposer 1 (P1), proposer 2 (P2) and responder (R). Figure 2 shows an example of the board, goal and player locations.

Proposers can propose a chip exchange to the responder. The responder can accept exactly one - or no proposal at all. All players are aware of the board state; proposers only have information about their own chip set and the one of the responder while the responder can see all chip sets. A single game of Colored Trails is played as a one-shot game. This means that proposers can only make a single proposal and the responder can only accept or reject proposals and is not allowed to make a counter offer. Furthermore, players do not actually have to move their pieces step-by-step toward the goal location. Once the responder has reacted to the proposals (by either accepting one or rejecting both) the chips are exchanged according to the winning proposal; or stay fixed if the responder rejected both proposals. Then, the best possible sequence of moves is automatically computed and each player receives a personal score.

Game phases

A game of Colored Trails is divided into a sequence of three phases and final evaluation. During each phase a timer indicates how much time is left in the current phase.

Initial phase

In this phase the game board and the chip stacks are presented to the players. This initial phase allows the player to locate its own piece on the board and get acquainted with the chip distribution.

Proposal phase

In this phase the two proposers (i.e. P1 and P2) can make chip exchange offers to the responder. The Proposers chips are always located on the left hand side of the window, the responders chip stack on the right hand side.

Reaction phase

In the reaction phase, the responder is presented with the two proposals. After evaluating the situation he can either accept one of the proposals or reject both.

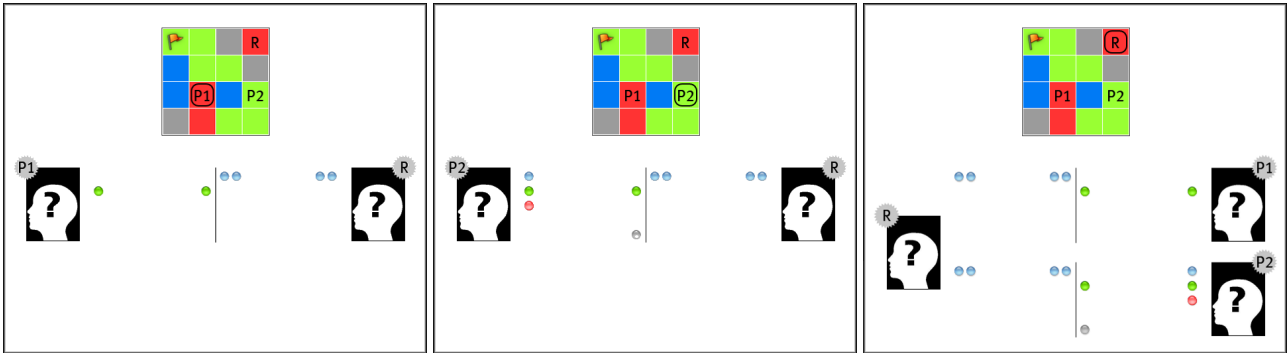


Figure 3. Colored Trails: three player negotiation game. From left to right: view of proposer 1 (P1), proposer 2 (P2) and responder (R). Note that placeholder figures are exchanged for pictures of the corresponding players.

Evaluation

The final evaluation is fully automatic; players do not need to act in this phase. The game server automatically computes the best possible sequence of moves for each player and assigns personal scores. Points are calculated as follows. For reaching the goal location a player receives 100 points. If he does not reach the goal, 25 penalty points are subtracted for every square between the goal and the players position. In addition, for every chip the player has not used, he receives 10 extra points.

3. RELATED WORK

This section presents a selection of related work. In particular, we refer to the bioengineering system *Homeostat* (Section 3.1) and the analyses of complex strategic interactions using evolutionary game theory as proposed by Walsh et al. [13] (Section 3.2).

3.1. Homeostat

Homeostat is one of the experiments participating in the Mars 500 study. The "Homeostat apparatus" is a system of interconnected devices, each operated by one of the subjects. Each device features an indicator and a control knob; rotating the knob influences the position of the indicator. The indicator is not directly linked to the control knob of the corresponding unit but rather responds to the operation of other devices as well. Every subject is instructed to fix the position of the indicator to a specific position. Due to the interaction between control loops, actions of each individual mutually induce interference. The collective group problem can only be solved through well coordinated actions. The selection of interaction structures (e.g. tree, chain or cycle) and interaction coefficients determine the degree of mutual influences and allow for a variety of experimental setups.

The Homeostat system has been previously used during a number of isolation and confinement experiments, in-

cluding the 60-day ESA EXEMSI campaign [9], the 135-day isolation during the "Human Behaviour in Extended Spaceflight" (HUBES) study in the *Mir* orbital station mock-up at IBMP [12] and the "Simulation of the Flight of the International Crew on Space Station" (SFINCSS).

Homeostat has proven to be suitable for on-line inferring of interaction effectiveness of an isolated crew with complex communication structure. Furthermore, leadership and "follower" tendency can be detected; the internal balance between these complementary behaviors is a deciding factor for the effectiveness of the group.

3.2. Evolutionary game theoretic analysis of complex games

In Section 2.2 we have introduced the notion of expected payoffs to quantify the success of actions taken during strategic interactions between players. Once these strategic encounters become more complex (i.e. increased number of players or actions), traditional game theoretic analysis becomes intractable. Walsh et al. [13] propose to use heuristic strategies rather than atomic actions and to compute a so called heuristic payoff table for these primitives. A *heuristic payoff table* captures the average payoff of each heuristic strategy for all possible mixtures of strategies for a finite number of players. This approach is suitable for games where the underlying rules are well-specified and common knowledge though they may describe very complex, repeated interactions between multiple agents. Rules specify particular actions agents may take depending on the the state of the game. Each agent follows a specific strategy selected from the set of heuristic strategies. Hence, the heuristic payoff table constitutes an abstract representation of the underlying game with the focus on heuristic strategies rather than atomic actions. Utilizing this methodology allows to apply for even more complex encounters the same game theoretic analysis as for simple games, e.g. the prisoners' dilemma. Walsh et al. [13] apply this approach to the auction domain and analyze a simulated continuous double auction with predefined trading strategies as heuristics.

Heuristic payoff tables have been proven to be beneficial beyond the domain of auctions and trading strategies. In [10] the same approach has been used to analyze the game of Poker (No-Limit Texas Holdem) using empirical data collected during human play. The authors use expert domain knowledge found in Poker literature to cluster complex atomic actions into a small set of heuristics, called *meta strategies*. This allows to construct a heuristic payoff table and perform an evolutionary game theoretic analysis; which enables to discover switching behavior between meta strategies.

4. METHOD

In this section we concisely explain the methodology we propose to perform the evolutionary game theoretic analysis of the data collected during Mars-500.

1. Collecting interaction data

The game Colored Trails serves as a strategic, interactive and situated task domain. We record full state information, proposals, responses as well as response times for each action.

2. Clustering atomic actions into meta strategies

Similar to the case of Poker, atomic actions in the game of Colored Trails are far too complex to analyze directly. However, clustering atomic actions into a small set of meta strategies is not straight forwarded either: expert domain knowledge for this specific game is lacking. Furthermore, the efficiency of automated clustering algorithms strongly depends on an appropriate distance measure. This distance measure needs to be defined over (a set of) features that describe atomic actions within Colored Trails. Both, the possibility to devise "hand-picked" rules clustering actions into meta strategies and the development of an adequate distance metric are the subject of current investigations.

3. Computing the heuristic payoff table

The heuristic payoff table represents the payoff table of the Colored Trails game for the different meta strategies the subjects can employ during play. It summarizes the success of a certain meta strategy depending on meta strategies chooses by other players. Hence, it simplifies the complexity of the game by using an abstract representation based on a small set of meta strategies.

4. Hypothesis testing and correlation

Once the heuristic payoff table is computed, evolutionary game theory analysis can be used to test different hypothesis, e.g. (a) the strategic behavior of subjects is influenced by personal preferences and not only guided by rational decision making; (b) subjects form coalitions while excluding other subjects from the group; (c) if the task becomes more complex, time pressure will effect the

Table 1. Preliminary data collection for the 105-days study. For both groups the number of session and games are listed as well as the playing time.

Group	Sessions	Games	Time
pre isolation			
Group 1	2	41	40 min
Group 2	2	42	35 min
in isolation			
Group 1	2	31	23 min
Group 2	2	29	20 min
total			
Group 1	4	72	63 min
Group 2	4	71	55 min

degree of influence induced by personal preferences. Furthermore, it needs to be shown if these findings correlate with other psychological investigations during Mars-500, e.g. questionnaire studies and the Homeostat experiment, as well as information obtained during debriefing.

5. PRELIMINARY RESULTS

The detailed analysis and discussion of data obtained during the 105-days study remains for future work. However, this section augments the general description of our methodology with preliminary results from the first weeks of the experiment.

Table 1 lists the number of games played prior to the isolation period as well as during the first weeks of isolation. Each group played four sessions with a total of 72 and 71 games respectively. In the proceeding analysis we identify subjects using a group prefix "1." or "2." coupled with a label "A", "B", "C" to distinguish between the three subjects within a group.

Clustering of atomic actions in Colored Trails is not straight forward. Therefore, let us restrict the analysis to the behavior of the responder. Note that roles rotate in each group, thus all subjects become the responder at equal intervals. Table 2 shows the absolute number of proposals accepted by the different responders as a function of the proposing individual. Subject 1.A accepted 14 proposals from 1.B, 7 proposals from 1.C and in three games both proposals got rejected. The chipset distribution does not favor one of the two proposers and therefore we can suspect that this imbalance hints at either

- subject 1.B is more generous while proposing, or
- a social preference for 1.B over 1.C.

To test further investigate this hypotheses we compute the payoffs associated with the meta strategies (1) "always




Figure 4. Illustration of the Colored Trails payoff table for the two groups of players. Arcs are labeled with the average payoff received by an individual (end of arrow) while accepting a proposal from another individual (origin of arrow).

accepting proposals from 1.B” and (2) ”always accepting proposals from 1.C”.

Meta strategy (1) yields an average payoff of 115.7 ± 18.0 and meta strategy (2) results in an average of 117.9 ± 17.3 . While both strategies award a comparable payoff, subject 1.A still favors proposals from 1.B over the proposals from 1.C with a ratio of 2:1. The full payoff tables for both groups is illustrated in Figure 4. In Group 2 we find a similar imbalance between proposal acceptance rate and associated average payoff. Subject 2.A accepted far more proposals from subject 2.B despite the higher average payoff obtained by accepting proposals from subject 2.C. An extreme case of such coalition forming is reflected by the fact that subject 2.C has not accepted a single offer from subject 2.B.

Once data is recorded over a longer period, payoff values can be recomputed over different time intervals. This allows to discover behavioral changes using the evolutionary game theoretic analysis as described in Section 4. With the aim to detect when coalitions form or brake up and how crew dynamics evolve over time.

Finally, preliminary data in Figure 4 shows that Group 1 succeeds to cooperate more efficient than Group 2 as indicated by a higher overall payoff flow, to be precise the weighted average payoff compares 112.9 (Group 1) to 96.4 (Group 2).

6. DISCUSSION AND CONCLUSIONS

While this paper has only presented results based on preliminary data, it is clear that the crew dynamics during the play of Colored Trails show interesting coalition forming behavior. This strengthens our hypothesis that play in Colored Trails is not purely based on rational decision making but rather influenced by personal preferences. It has yet to be investigated if these findings correlate with psychological investigations of various aspects of crew behavior during Mars-500, e.g. questionnaire studies and

Table 2. Absolute number of proposals accepted by the different responders as a function of the proposing individual (within brackets).

Responder	Proposal 1	Proposal 2	Reject
1.A	14 (1.B)	7 (1.C)	3
1.B	7 (1.A)	12 (1.C)	5
1.C	15 (1.A)	7 (1.B)	1
2.A	18 (2.B)	3 (2.C)	1
2.B	7 (2.A)	15 (2.C)	0
2.C	17 (2.A)	0 (2.B)	4

the Homeostat experiment, as well as information obtained during debriefing.

To summarize, we proposed to leverage the mathematical means of game theory to analyze interpersonal crew dynamics during ultra-long duration space flights. We describe how game theory facilitates capturing the essence of crucial interactive decision making processes, thereby raising the potential for a fully automated and unintrusive monitoring and diagnosis tool; such a tool is the prerequisite to effectively trigger computer-based self-help countermeasures to depression or anxiety. Finally, we have presented preliminary findings based on the first 105-day test phase of the ground-based isolation study Mars-500 and outlined our research agenda for the full 520-day isolation period.

ACKNOWLEDGMENTS

This research was partially funded by the Netherlands Organization for Scientific Research (NWO) and the Netherlands Institute for Space Research (SRON).

REFERENCES

- [1] Axelrod, R. and Hamilton, W. (1981). The evolution of cooperation. *Science*, 211:1390–1396.
- [2] Cavanagh, K. and Shapiro, D. (2004). Computer treatment for common mental health problems. *Clinical Psychology*, 60:239–251.
- [3] Ficici, S. and Pfeffer, A. (2008). Modeling how humans reason about others with partial information. In *Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems*.
- [4] Gal, Y., Grosz, B. J., Kraus, S., Pfeffer, A., and Shieber, S. (2005). Colored trails: a formalism for investigating decision-making in strategic environments. In *Proceedings of the IJCAI Workshop on Reasoning, Representation, and Learning in Computer Games*.
- [5] Gintis, H. (2000). *Game Theory Evolving. A Problem-Centered Introduction to Modelling Strategic Interaction*. Princeton University Press.
- [6] Greist, J., Marks, I., Baer, L., Kobak, K., Wenzel, K., Hirsch, M., Mantle, J., and Clary, C. (2002). Behavior therapy for obsessive-compulsive disorder guided by a computer or by a clinician compared with relaxation as a control. *Clinical Psychiatry*, 63:138–145.
- [7] Kanas, N. et al. (2006). Psychology and culture during long-duration space missions. Technical report, International Academy of Astronautics Study Group on Psychology and Culture During Long-Duration Space Missions.
- [8] Leon, G. (1999). Select-in and countermeasure considerations for long duration crews. In *Proceedings of the 29th International Conference on Environmental Systems*.
- [9] Novikov, M., Bystritskaya, A., Eskov, K. N., Vasilyev, V. K., Vinokhodova, A. G., and Davies, C. (1993). Homeostat - a bioengineering system. In *Proceedings of the 23rd International Conference on Environmental Systems*.
- [10] Ponsen, M., Tuyls, K., de Jong, S., Ramon, J., Croonenborghs, T., and Driessens, K. (2008). The dynamics of human behaviour in poker. In *Proceedings of the 20th Belgium-Netherlands Conference on Artificial Intelligence (BNAIC 2008)*, pages 225–232, University of Twente, Enschede, The Netherlands.
- [11] Proudfoot, J., Goldberg, D., Mann, A., Everitt, B., Marks, I., and Gray, J. (2003). Computerized, interactive, multimedia cognitive-behavior program for anxiety and depression in general practice. *Psychological Medicine*, 33:217–227.
- [12] Savilov, A., Bayevsky, R., Bystrytskaya, A., Gushin, V., Manovtsev, G., Nichiporuk, I., Novikov, M., Ponomareva, I., and V.P.Salnitsky (1997). On-board equipment-based study of the dynamic of psycho-physiological and biomedical responses of the operators during 135-day isolation in the mir orbital station mock-up. *Aerospace and Ecological Medicine*, 5:28–38.
- [13] Walsh, W., Das, R., Tesauro, G., and Kephart, J. (2002). Analyzing complex strategic interactions in multi-agent systems. In *Proceedings of the AAAI Game Theoretic and Decision Theoretic Agents Workshop*.