Predictability in Human-Agent Cooperation: Adapting to Humans' Personalities

Sebastian Ahrndt
DAI-Laboratory
Technische Universität Berlin
Ernst-Reuter-Platz 7
10587 Berlin, Germany
ahrndt@dai-labor.de

Benjamin Breitung
DAI-Laboratory
Technische Universität Berlin
Ernst-Reuter-Platz 7
10587 Berlin, Germany
breitung@dai-labor.de

Johannes Fähndrich
DAI-Laboratory
Technische Universität Berlin
Ernst-Reuter-Platz 7
10587 Berlin, Germany
faehndrich@dai-labor.de

Sahin Albayrak
DAI-Laboratory
Technische Universität Berlin
Ernst-Reuter-Platz 7
10587 Berlin, Germany
albayrak@dai-labor.de

ABSTRACT

Making artificial agents a constituent part of human activities leads to more affiliated teamwork scenarios and at the same time introduces several new challenges. One challenge is the team members' ability to be mutually predictable, which is required to effectively plan own actions, e.g., in the field of human-aware planning. This work approaches the question whether or not agents are able to learn the personality of a human during interaction. In particular, we developed an agent model able to learn human personality during repeatedly played rounds in the Colored Trails Game. Human personality is described using a psychological theory of personality types known as the Five-Factor Model. The results indicate that some characteristics of a personality can be learned more accurately/easily than others.

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Learning; J.4 [**Computer Applications**]: Social and Behavioural Sciences

Keywords

Information systems, User/Machine systems, Human factors, Software psychology

1. INTRODUCTION

Human-human teamwork has already been studied decadeslong and several properties distinguishing good and effective from flawy teamwork were identified [18]. One of them is the characteristic of team members to be (mutually) predictable to each other [5]. In human-agent teamwork, predictability addresses the circumstance that an agent can only plan its own actions effectively—which includes coordination activities—when it is assessable what the others, including the human, will do [4, 13]. To address this challenge a combination of theory- and data-driven approaches was postulated as being beneficial (cf. [17]). In particular, the use of human behavioural models providing insights into the human nature from the psychological point of view is desired [3, 12]. These models form the theoretical basis for predicting human behaviour and can provide information about the personality, habits and capabilities of humans, e.g. in terms of hand-coded rules. Although these models are a good starting point, they must be adapted to the human's individual preferences during the actual interaction [12]. Hence, the motivation for this work is to proof that agents are able to learn the personality of a human during interaction. In order to do so, we applied a scientific game as a testbed and a human personality model derived from psychology as the theoretical vehicle. The findings indicate that some characteristics of a personality can be learned more accurately/easily than others, and that this information can be used within the interaction with humans to predict their next actions more accurately.

In the following we will introduce the scientific game we used within this work (see Section 2) and the psychological theory that was applied (see Section 3). Subsequently, the topic of this work will be compared to the previous state-of-the-art (see Section 4). Afterwards, we will introduce the agent model used (see Section 5) and the experimental results (see Section 6). Finally, we will wrap up the work providing final remarks (see Section 7).

SAC'15, April 13-17, 2015, Salamanca, Spain.

http://dx.doi.org/10.1145/2695664.2695702

2. THE COLORED TRAILS GAME

The Colored Trails Game [8] (CT) is a multi-agent testbed to investigate cooperative decision making within a chess-like setting. The basic settings of the game are the following: The board is a $N \times M$ grid consisting of coloured squares with a predefined set of available colours. Each player has

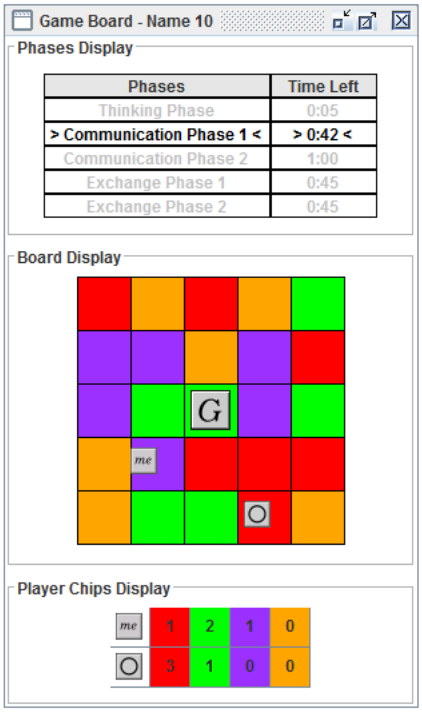

Figure 1: A sample screenshot of the game environment, showing the different phases (without movement), the current board and the available chips for each player.

a specific starting position and an amount of coloured chips that can be used to move to a square of the same colour. The main goal for each player is to reach the 'goal square'. An arbitrary amount of chips can be exchanged with (or given to) other contestants in a communication phase arranged for this purpose, consisting of a proposal, a decision, and the actual exchange. Proposed offers can be accepted or refused, but it is not required to give the amount of promised chips. A lot of definitions, e.g. whether or not the goal square is identical for everyone or if the other players' chips can be seen are left open and can be alternated to meet each researcher's expectations. While some rules of the game — like the size of the board — are less important than others, especially the distribution of chips as well as the scoring evaluation are essential when examining the players' behaviour.

The limited availability of the resources force cooperation if not all of the players have the possibility of reaching the goal on their own or if some can while others need assistance. Nevertheless, the highest significance lies in the calculation of the game's final score. Commonly, points for reaching the goal will be granted but everything else can be defined to provoke or decrease cooperation.

Fig. 1 shows a screenshot of the game environment. The 'phases display' indicates the current phase of the game. The thinking phase has been inserted for human players to get an overview of the situation. The communication phase can be used by the players to create a proposal in an additional window. In the exchange phase the players transfer some amount of chips, which might be related to the priorly accepted proposal.

3. FIVE-FACTOR MODEL

The Five-Factor Model of personality [14, 15] (FFM) is a psychological theory that can be used to model human personality types and their influences on the decision-making process of humans. As suggested by the name, the FFM introduces five dimensions characterising an individual, which are briefly described in the following:

- ***O**penness to experience* describes a person's preference to vary their activities over keeping a strict routine and is also related to their creativity (e.g., inventive, emotional and curious behaviour vs. consistent, conservative and cautious behaviour).
- ***C**onscientiousness* describes a person's preference to act duteously over spontaneously. This directly relates to the level of self-discipline when aiming for achievements (e.g., efficient, planned and organised behaviour vs. easy-going, spontaneous and careless behaviour).
- ***E**xtraversion* describes a person's preference to interact with other people and to gain energy from this interaction over being more independent of social interaction (e.g., outgoing, action-oriented and energetic behaviour vs. solitary, inward and reserved behaviour).
- ***A**greeableness* describes a person's preference to trust others, to act helpful and to be optimistic over an antagonistic and sceptical mindset. This trait directly influences the quality of relationships with other individuals (e.g., friendly, cooperative and compassionate behaviour vs. analytical, antagonistic and detached behaviour).
- ***N**euroticism* describes a person's preference to interpret external stimuli such as stress as minatory over confidence and emotional stability. Neuroticism addresses the level of emotional reaction to events (e.g., sensitive, pessimistic and nervous behaviour vs. secure, emotionally stable and confident behaviour).

The characteristic of each dimension is defined as a variation from the norm, whereas each dimension is an overarching container subsuming different lower-level personality traits. For example, neuroticism is associated with subordinated traits such as anxiety, hostility and impulsiveness [15]. Taking this observation into account, one can argue that the FFM theory is a conceptual framework about human personality traits that can, for example, be used to integrate other theories about human personalities into its structure [10, 16]. The reason for using the FFM of personality instead of other popular personality theories can be found elsewhere [1].

4. RELATED WORK

Among others, predictability requires deliberating about the other agents' actions. In particular it requires to deliberate what the next actions are that the team members will execute in order to approach the joint goal of the collaboration. In human-agent interaction predictability can be found in works reaching from the question whether people reason about other people's actions [6] to questions of how to recognise inter-player relationships in multi-player games [21]. The learning of social preferences to enhance predictability is another topic. Here, *L. Hoog* and *N. Jennings* [9] present a work where agents use a weighted sum of the other agents' expected outcomes as a utility function. The examined behaviour is called socially rational decision-making and is based on the idea of social welfare functions. A comparable work is presented by *Gal et al.* [7] introducing social preferences in terms of the three dimensions self-interest, social welfare and inequity aversion. Agents build knowledge in these dimensions about the other agents and integrate this knowledge into their own decision-making process. A detached approach is presented by *Bradshaw et al.* [4] using hand-crafted policies to adjust the decision-making of agents to the respective use-case.

Talman et al. [20] present a work that illustrates the use of a rather simple abstraction of personality types. Personalities of agents are determined through the two dimensions cooperation and reliability. The agents play the CT game and try to optimise a utility function incorporating whether the player reaches the goal, the distance to the goal and the number of chips left. During repeatedly played games the agents reason about each other's helpfulness along the two dimensions. As an effect they try to respond more effectively by customizing their behaviour appropriately for different personalities. For example, an agent tries to avoid collaboration with another agent recognised as selfish, meaning that the other agent is neither cooperative nor reliable. Otherwise, the selfish agent would always win the game at the others' expense. The extensive evaluation carves out that the agent that adapts its own behaviour with regard to the personalities of the opponents outperforms agents who do not adapt. Furthermore this adaptation leads to an increased social welfare for the group in the long term.

This paper differs from the ones presented above in that we apply an existing psychological theory of human personality and evaluate our agent model using games played with humans and artificial agents. In addition it is not our goal to produce optimal group behaviour but to prove that we can learn and use information about the personality of the human. It is comparable to the mentioned works as it also applies a multi-attributed utility function for the decision-making process. In fact, it is motivated by the work of *Talman et al.* and transfers their ideas from agent-agent to human-agent cooperation.

5. APPROACH

In the following we will construct the agent model. It is restricted to three of the five dimensions (conscientiousness, extraversion, agreeableness) of the FFM. That is because the action space of CT makes it difficult to associate all actions with the traits of the FFM. Still, the agent model is general enough to vary the game's complexity by means of the grid size, the number of players and the visibility parameters. Our approach is based on the idea to link the personality traits to the available actions by interpreting the meaning of the trait, taking into consideration the effect of the action. For the remaining traits (openness, neuroticism) this is hard to accomplish, as there is no possibility to reward creative or punish conservative behaviour. Furthermore, as CT is a scientific game it is not constructed to evoke emotional reaction in its players.[1]

[1]One might argue that repeatedly losing in the game leads to an emotional reaction. But this effect is an ordinary one and solely considered no indication for the emotional stability.

To build an estimate of the personality of the human an agent i refers to a human k using the tuple $P_k = \{p_c^k, p_e^k, p_a^k\}$, where each $p \in P$ represents a personality trait. As the traits in the FFM are declared as variation from the norm, the range of each p is $[0, 1]$ and the initial value is set to 0.5. This set is one of the features used by the agent to build the expected utility of taking action a while playing against a human.

To improve the estimates of the personality the agent adapts each p during the interaction in the following way:

- p_c — denotes the estimate of the conscientiousness of the human and is interpreted as how reliable a player is. Therefore fulfilling a trade increases and not fulfilling it decreases this value. As failing to predict the reliability of a player can lead to significant score losses for the agent, this trait is of utmost importance. To update the estimate after each trading agreement, we compute the conscientiousness of a human by increasing/decreasing it with a constant factor x_c using the following equation:

$$p_c \leftarrow \begin{cases} p_c + x_c & \text{if successful exchange} \\ p_c - x_c & \text{if successful exchange but fraud} \\ p_c - 2 \cdot x_c & \text{if fraud} \end{cases} .$$

  The first case applies when the proposed set of chips is equal to the one received. The second case applies when the set of proposed and received chips is not equal, but in the set of received chips exist some chips that are useful for the agent. The last case applies if the agent was fooled. This is the case when there is no exchange or when the agent only receives useless chips. Thus, bailing out an agreed trade is punished harder, as it is a greater break of trust and might critically damage the agent's chance to reach the goal square.

- p_e — denotes the estimate of the extraversion of the human and is interpreted as how sociable the player is. Therefore it is increased when the player makes a proposal of exchanging chips, which is the most extroverted action possible in the game. It is decreased when the player acts passively by not proposing anything. To update the estimate after each round, we compute the extraversion of a human by increasing or decreasing it with a constant factor x_e using the following equation:

$$p_e \leftarrow \begin{cases} p_e + x_e & \text{if proposed} \\ p_e - n \cdot x_e & \text{otherwise} \end{cases} .$$

  The first case applies when the player offers a proposal, the second case otherwise. The multiplicator n is growing until the player offers something and corresponds to the number of rounds played:

$$n \leftarrow \begin{cases} 0 & \text{if proposed} \\ n + 1 & \text{otherwise} \end{cases} .$$

- p_a — denotes the estimate of the agreeableness of the human and is interpreted as how friendly/altruistic a player is. Therefore it is increased when the player accepts and decreased when the player declines offers. It is increased/decreased twice when the offers are favourable for the opponent. To update the estimate after each active communication phase, we compute the agreeableness of a human by increasing/decreasing it with a constant factor x_a using the following equation:

$$p_a \leftarrow \begin{cases} p_a + 2 \cdot x_a & \text{if accepted and altruistic} \\ p_a + x_a & \text{if accepted} \\ p_a - x_a & \text{if not accepted} \\ p_a - 2 \cdot x_a & \text{if not accepted but favorable} \end{cases}.$$

  This equation rewards generous offers and exchanges as they might be harmful to the players own score. At the same time it reduces the agreeableness estimate when the exchange of important chips was declined. Thus the level of agreeableness is a kind of measure of the selfishness of the player.

The constants x_c, x_e and x_a were adjusted and determined in test games played prior to the experiment. For reasons of readability we omitted the edge cases when the estimates reach the minimal/maximal value of the interval. In these cases a positive/negative adjustment was no longer applied.

We use the estimates p_e and p_a to calculate the expectation that a proposal will be accepted, as the weighted sum $e^{acc} = p_e \cdot w_e + p_a \cdot w_a$. The weights are used to adjust the influence of the traits. A second value indicates the expectation whether an agreed exchange indeed takes place and is represented as $e^{exc} = p_c$.

The second feature to build the expected utility is the score that is reachable with the current set of coloured chips (r^c), the score that is reachable after a successful trade (r^t) and the score that is reachable falling for a betrayal (r^f). Here betrayal means accepting a trade and transferring own chips without getting the promised response. All three can be easily calculated when knowing (1) that CT controls the movement phase by applying the A* algorithm to determine the best option to move towards the goal square and (2) the scoring function of the game, which sums the following parameters:

- 100 starting points;
- 50 points for reaching the goal square and ending the round as winner;
- 10 additional points for all coloured chips left;
- 10 penalty points for each proposal made by the player; and
- 20 penalty points for each tile between the final position and the goal square calculated using the Manhattan distance.

Both features are then used to calculate the expected value (reward) of executing action a given the current state of the game s using the following multi-attribute utility function when making a proposal:

$$u_a^i(s, P_k) = e^{acc} \cdot e^{exc} \cdot r^t + (1 - e^{acc}) \cdot r^c + e^{acc} \cdot (1 - e^{exc}) \cdot r^f$$

When the agent receives a proposal the likelihood that it will be accepted is not of relevance since the agent can choose its answer and only has to consider that the exchange truly takes place. Therefore we remove e^{acc} when building the utility for an action in this case.

Given this function the estimate of the personality of the human influences the policy of the agent, which tries to maximise the utility. That means that for each agent i playing against a human k an optimal action a_k^* exists, that maximise the utility in state s where $a_k^* \in \underset{a}{argmax}\ u^i(s, P_k)$, which is executed. If equally valued actions exist, the one is selected that was found first. Indeed, in the implementation the agent has no knowledge that there exists more than one action that maximise the utility. A more elaborate behaviour here would be to evaluate whether a chain of actions would lead to a higher score, leading to an agent that acts 'farsighted' instead of 'myopic'.

6. EXPERIMENT AND RESULTS

For the experiment we implemented the introduced agent model for the CT environment and invited 22 participants, which were mainly students. At the beginning, the participants were asked to describe their personality using a questionnaire derived from the IPIP[2]. Afterwards, the game environment was explained and each participant got a 10 minutes tutorial on how to play the game. Here we explained the rules and the scoring function, and actually played the game with the subjects. The scoring function consisted of reaching the goal, the distance to the goal and the chips left as described earlier. In the initial stage the participants played against an agent that did not adapt to the opponent. Afterwards the attendees played 30 games in a row against the adapting agent. The goal of the participants was to reach the maximum score in as many games as possible.

Table 1 lists the data collected within the experiment. The scoring results listed in columns 2 and 3 show the mean value of the points of all 30 games determined for each human player and the agent playing against the participant. It shows that the agent outperforms the human players on average, even though the difference is fairly small. We tried to minimise random effects by setting up the CT environment in a way that the same number of chips was given to the opponents and the central field was chosen as the goal square. Taking that and the total number of 660 games played into consideration, the scoring difference can be seen as significant.

However, that only shows that an agent adapting to its opponent can compete with the human player and illustrates one use-case for applying information about personality to human-agent cooperation. Table 1 also lists the deviation between the agents' estimates of the personality traits (columns 3 to 5) of their opponents and the actual personality assessment derived from the questionnaire including the

[2]IPIP — International Personality Item Pool: A Scientific Collaboratory for the Development of Advanced Measures of Personality and Other Individual Differences — `http://ipip.ori.org/`. For the experiment the 100-Item Set of IPIP Big-Five Factor Markers was used.

Table 1: Listing of the average scores reached by the opponents (human and agent) within the games and the average score and deviation over all games. Also includes a listing of the deviation between the agent's estimate of the humans personality trait and the one derived from the questionnaire.

#	Human	Agent	Extra.	Agree.	Consc.
1	111	99	0.18	0.225	0.265
2	107	154	0.09	0.09	0.12
3	98	113	0.02	0.28	0.245
4	121	127	0.075	0.025	0.215
5	118	140	0.035	0.06	0.4
6	105	113	0.05	0.19	0.175
7	132	134	0.03	0.235	0.09
8	100	107	0.14	0.335	0.24
9	88	154	0.015	0.05	0.425
10	142	102	0.045	0.225	0.11
11	104	106	0.055	0.195	0.075
12	105	112	0.07	0.295	0.37
13	99	144	0.06	0.17	0.425
14	121	120	0.065	0.19	0.12
15	126	111	0.16	0.095	0.215
16	145	137	0.04	0.05	0.22
17	86	141	0.025	0.215	0.065
18	102	107	0.015	0.06	0.13
19	138	132	0.145	0.075	0.47
20	154	110	0.05	0.21	0.275
21	101	124	0.02	0.165	0.215
22	97	138	0.125	0.285	0.23
μ	**113.64**	**123.86**	**0.07**	**0.17**	**0.23**
σ	**15.84**	**14.77**	**0.04**	**0.08**	**0.09**

average deviation. It shows that the smallest variation is found for the extraversion parameter, while agreeableness and conscientiousness are drifting further apart (Fig. 2 depicts the spreading of the values in a boxplot). A value of zero would mean that both characterisations are perfectly equivalent, which is only a theoretical option that also can not be observed when assessing an individual using self-assessment, a questionnaire and a professional assessment.

Since the CT game has a very limited action space available to evaluate and analyse the behaviour of the other player, it is difficult to associate these actions with factors of the FFM. Thus, the environment/action-space might have to be more complex. Another possible explanation for the parameters not depicting the survey results might be that our interpretation of how the traits influence the actions taken does not precisely fit the humans' behaviour. Since the goal of the game is to reach the best possible score, it might be beneficial to use a more generalised trait just indicating how cooperative the human is (as done by *Talman et al.* [20]). Despite these imaginable hindrances the outcome is still considerably good, especially for the value of extraversion.

7. CONCLUSION

In this work we presented an agent model that uses a representation of a human's personality to reason about the outcome of cooperative actions in the Colored Trails Game.

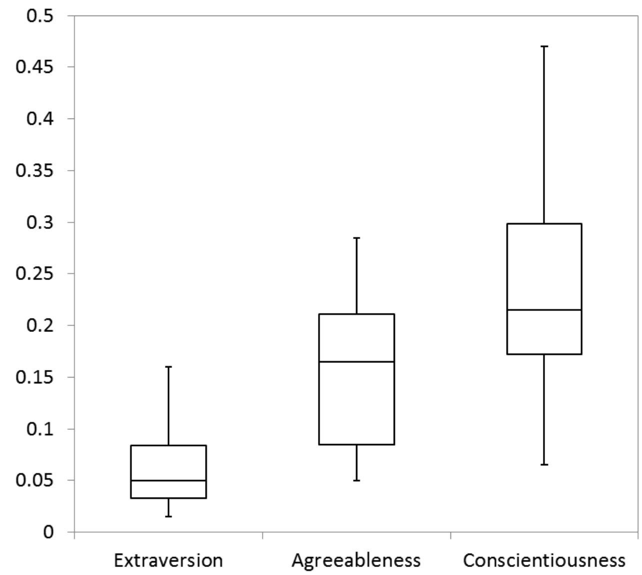

Figure 2: Boxplot of deviation between questionnaire and agent's estimate of the player's personalities.

The personality representation is based on the Five-Factor Model of human personality. The agent is able to adapt its estimates of the personality traits while observing the actions of the other player. The basic idea here was to link the traits to the action space by interpreting the meaning of the trait, taking into consideration the effect of the action. The experimental results showed that the different personality traits vary in the difficulty to be observed/learned. Nevertheless, the evaluation shows that this information can be used within the interaction with humans to predict the next actions of such humans more accurately. Even though it was not the goal of the work to implement an agent able to outperform the human players, the agent's performance was actually equal or even better than that of most human players. Here further investigations must be made to gather information why some of the human players outperformed the agent even if it was able to build a model of them. As this is only related to the game played, the overall findings are interesting for fields such as human-agent negotiation and human-agent cooperation. In fact, it is one of the requirements postulated for joint human-agent activities (cf. [4, 11, 12, 13]).

In future work, we want to integrate these findings into HPLAN [2], an extension of the JIAC V agent-framework, which facilitates the implementation of joint human-agent activities. Here developers will be enabled to annotate agents' actions and their influences on the personality trait. Using such a development environment will enable us to conduct further research and user studies in environments inhabited by robots and humans. Furthermore, it is planned to replace the currently used approach of adapting the estimate by replacing it with existing works from the reinforcement learning (RL) community, such as Q-learning [19]. This can be done, as predictability addresses the transition probability of Markov Decision Processes (MDPs), which is in our work influenced by the personality of the user. Here future work includes a theoretical part about integrating the personality model as a preference vector that can be used to learn in MDPs and a practical part evaluating whether or not model-based RL can be used under this circumstances and whether or not model-based RL is fast enough for direct interaction with human users.

8. REFERENCES

[1] S. Ahrndt, A. Aria, J. Fähndrich, and S. Albayrak. Ants in the OCEAN: Modulating agents with personality for planning with humans. In N. Bulling, editor, *Multi-Agent Systems (EUMAS 2014)*, Lecture Notes in Artificial Intelligence, pages 1–16. Springer, December 2014.

[2] S. Ahrndt, P. Ebert, J. Fähndrich, and S. Albayrak. Hplan: Facilitating the implementation of joint human-agent activities. In Y. Demazeau, F. Zambonelli, J. M. Corchado, and J. Bajo, editors, *Advances in Practical Applications of Heterogeneous Multi-Agent Systems. The PAAMS Collection.*, volume 8473 of *Lecture Notes in Computer Science*, pages 1–12. Springer Int. Publishing, 2014.

[3] S. Ahrndt, J. Fähndrich, and S. Albayrak. Human-aware planning: A survey related to joint human-agent activities. In J. Bajo, J. M. Corchado, P. Mathieu, A. Campbell, A. Ortega, E. Adam, E. M. Navarro, S. Ahrndt, M. Moreno, and V. Julián, editors, *Trends in Practical Applications of Heterogeneous Multi-Agent Systems. The PAAMS Collection.*, volume 293 of *Advances in Intelligent Systems and Computing*, pages 95–102. Springer Int. Publishing, 2014.

[4] J. M. Bradshaw, P. Feltovich, M. Johnson, M. Breedy, L. Bunch, T. Eskridge, H. Jung, J. Lott, A. Uszok, and J. Diggelen. From tools to teammates: Joint activity in human-agent-robot teams. In M. Kurosu, editor, *Human Centered Design*, volume 5619 of *Lecture Notes in Computer Science*, pages 935–944. Springer Berlin Heidelberg, 2009.

[5] M. E. Bratman. Shared cooperative activity. *The Philosophical Review*, 101(2):327–341, April 1992.

[6] S. G. Ficici and A. Pfeiffer. Modeling how humans reason about others with partial information. In Padgham, Parkes, Müller, and Parsons, editors, Proc. of the *7th Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, volume 1, pages 315–322. IFAAMAS, May 12–16 2008.

[7] Y. Gal, A. Pfeffer, F. Marzo, and B. J. Grosz. Learning social preferences in games. In *Proc. of the 19th National Conf. on Artificial Intelligence*, pages 226–231. The AAAI Press, 2004.

[8] B. J. Grosz, S. Kraus, S. Talman, B. Stossel, and M. Havlin. The influence of social dependencies on decision-making: Initial investigations with a new game. In *Proc. of the 3rd Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, volume 2 of *AAMAS*, pages 782–789. ACM Press, 2004.

[9] L. M. Hoog and N. R. Jennings. Socially intelligent reasoning for autonomous agents. *IEEE Transactions on Systems Man and Cybernetics*, 31(5):381–393, September 2001.

[10] O. P. John and S. Srivastava. The big-five trait taxonomy: History, measurement, and theoretical perspectives. In L. A. Pervin and O. P. John, editors, *Handbook of Personality: Theory and Research*, volume 2, pages 102–138. The Guilford Press, 1999.

[11] A. Kirsch, T. Kruse, and L. Mösenlechner. An integrated planning and learning framework for human-robot interaction. In *4th Workshop on Planning and Plan Execution for Real-World Systems (held in conjuction with ICAPS 09)*, pages 1–6, Thessaloniki, Greece, September 2009.

[12] A. Kirsch, T. Kruse, E. A. Sisbot, R. Alami, M. Lawitzky, D. Brscic, S. Hirche, P. Basili, and S. Glasauer. Plan-based control of joint human-robot activities. *KI – Künstliche Intelligenz*, 24(3):223–231, September 2010.

[13] G. Klein, D. D. Woods, J. M. Bradshaw, R. R. Hoffmann, and P. J. Feltovich. Ten challenges for making automation a 'team player' in joint human-agent activity. *Human-Centered Computing*, 19(6):91–95, November/December 2004.

[14] R. R. McCrea and P. Costa. Reinterpreting the Myers-Briggs Type Indicators from the perspective of the Five-Factor Model of personality. *Jounal of Personality*, 57(1):17–40, March 1989.

[15] R. R. McCrea and O. P. John. An introduction to the Five-Factor Model and its applications. *Journal of Personality*, 60(2):175–215, 1992.

[16] B. P. O'Connor. A quantitative review of the comprehensiveness of the Five-Factor Model in relation to popular personality inventories. *Assessment*, 9:188–203, 2002.

[17] R. Prada and A. Paiva. Human-agent interaction: Challenges for bringing humans and agents together. In *Third Int. Workshop on Human-Agent Interaction Design and Models*, Proc. of the 13th Int. Conf. on Agents and Multiagent Systems (AAMAS), pages 1–10. IFAAMAS, 2014.

[18] K. Sycara and G. Sukthankar. Literature review of teamwork models. Technical Report CMU-RI-TR-06-50, Robotics Institute, Pittsburgh, PA, November 2006.

[19] R. S. Sutton and A. G. Barto Reinforcement Learning: An Introduction, MIT Press, May 1998 .

[20] S. Talman, M. Hadad, Y. Gal, and S. Kraus. Adapting to agents' personalities in negotiation. In M. Pechoucek, D. Steiner, and S. Thompson, editors, *Proc. of the 4th Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 383–389, ACM, New York, NY, USA, 2005.

[21] B. Wilson, I. Zuckerman, and D. Nau. Modeling social preferences in multi-player games. In L. Sonenberg, P. Stone, K. Tumer, and P. Yolum, editors, *Proc. of the 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, volume 1–3, pages 337–344. IFAAMAS, Taipei, Taiwan, 2011.