Using Personality Models as Prior Knowledge to Accelerate Learning About Stress-Coping Preferences*

(Demonstration)

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

Marco Lützenberger
DAI-Labor, Technische Universität Berlin
Ernst-Reuter-Platz 7
10587 Berlin, Germany
marco.luetzenberger@dai-labor.de

Stephen M. Prochnow
German Turkish Advanced Research Centre for ICT
Ernst-Reuter-Platz 7
10587 Berlin, Germany
stephen.prochnow@gt-arc.com

ABSTRACT

The management of (dis)stress is an important factor for a long and healthy life. Yet, stress affects people differently and everyone manages stress in different ways. In this paper we introduce *PeSA*, the *Personality-enabled Stress Assistant*, an agent-based application that accounts for this individualism. PeSA merges several agent techniques: Reinforcement learning is used to learn about preferences of the users, prior knowledge and knowledge transfer is applied to accelerate the learning process, agent mirroring helps to enable communication and offline functionalities. Based on these mechanisms, PeSA guides through stressful phases by proposing coping strategies that are tailored to the personality of each individual user. Users can assess these advices and thus provide a reward or punishment signal that helps PeSA to improve its suggestions.

Keywords

Human-Agent Teamwork; Human-Behaviour Models; Reinforcement Learning

1. INTRODUCTION

Psychological stress is a well-known trigger for several physical diseases and has a significant impact on our health and health care costs (cf. [5]). The problem that we address with PESA is that the perception of stress is highly individual—as are the actions that may help to relieve stress [2]. PeSA accounts for this very fact by adapting its coping-strategy recommendations to the individuality of the user. In doing so, individual preferences, characteristics, and personalities are supported.

The implementation of this capability is based on a combination of a data-driven and a theory-driven mechanism and directly addresses one of the unsolved human-agent interaction challenges that were described by Prada and Paiva [9]. The theoretical component acts as an accelerator for learning individual preferences and is based on models that were established by human-factor psychology. The data-based component learns by receiving rewards from the users. Knowledge from other PeSA agents is also used to accelerate the learning process, however, to make this work, their users have to have compatible personality models.

PeSA for itself is basically a combination of agent-based software engineering and learning algorithms that use the humans' feedback as reward signals. From a technical perspective, any PeSA instance is a software agent that autonomously collects data, individually detects situations in which stress-relieving actions are required and adapts its recommendations to the personality of its user. To make the learning process faster, we extended the agents' capability to share knowledge with other PeSA agents. The basic idea behind this is that users with similar personalities prefer similar countermeasures (cf. [2, 3]), thus, it is possible that relatively new PeSA agents can learn from agents that have already collected experiences in relieving stress levels—given that their users show similar characteristics.

In our demo, we show how an experienced PeSA agent recommends coping-strategies in order to relieve the stress level of its user and how this agent learns from user feedback. We also show how fresh PeSA agents request knowledge from 'mature' PeSA agents with similar personality profiles in order to become better assistants very quickly. To make this tangible, we show that the user of the new agent is treated in a similar fashion, once the knowledge of an available agent has been transferred and once a stressful situation has been detected. We also show that such knowledge transfer is only done when the personality profiles are actually compatible.

2. PESA IN A NUTSHELL

PeSA was implemented as an application for the Android platform with a cloud-based backend that hosts the multi-agent system. Within the Android container we implemented PeSA as a software agent using the multi-agent-framework *JIAC V* [7]. The backend hosts mirrors of the individual agents, enabling the communication within the agent-system. In more detail: The module that observes a user and triggers actions is separated from the one that communicates to other modules in order to exchange already learned information about how to treat particular person-

*Visit `http://pesa.dai-labor.de` for more information.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

ality profiles. Each module was implemented as a JIAC V agent, such that a PeSA agent actually comprises two agents (one mandatory agent on the Smartphone and one optional agent in the backend). The reason for this design decision is data security. Exchanging information about a user's psychological profile is critical in terms of data security. We offer this feature to accelerate the learning process, yet, only optionally. Agents on the user's Smartphone are capable to learn all required information by themselves. We proceed by explaining how PeSA agents can learn and subsequently elaborate on how they exchange their collected knowledge.

Single Agent – Local learning. To enable the PeSA agent to perform stress management and to adapt its recommendations to the particular requirements of its user, we implemented a reinforcement learning (RL) cycle using classical Q-learning [10] to learn about the action preferences of each user, whereas the actions are the recommended coping-strategies. Since tabula-rasa RL is not fast enough in direct interaction with users, we integrated information about correlations between personality and preferred coping strategies as prior knowledge. These correlations were retrieved from psychological studies [2, 3] and integrated into our PeSA agent as an initial feedback model $\hat{H}$ (henceforth: human-behaviour model) using action biasing [6]. Applying action-biasing adjusts the Q-values only during the action selection, thus, it does not affect the learned Q-values directly. Action-biasing reads as follows: $Q'(s,a) = Q(s,a) + (\beta * \hat{H}(s,a))$, where β is a decaying parameter. In other words, one can say that the personality information is used as an initial heuristic for the action selection. We described this idea in prior work [1]. To complement the human-behaviour model, we determine the personality of the user by means of a questionnaire, which has its roots in human psychology and which aims to assess the *Big-Five* [8] personality traits. In PeSA, we use the determined Q-values to sort the list of available coping-strategies.

Multi-agent system – Knowledge Sharing. To accelerate the learning cycle, we (optionally) allow inexperienced PeSA agents to request knowledge that was already learned by existing agents that assist users with similar personality profiles. If knowledge-sharing is enabled, each PeSA app communicates to one PeSA agent that is running in the backend. Backend agents store information about the learned Q-values and the state of the human-behaviour model. With each new installation, one PeSA agent is deployed in our backend—this agent represents the new user. The agent receives the personality profile from the PeSA app agent after the assessment was completed and broadcasts this information to other available agents in the system, asking for similar personalities and already learned preferences. Responses are collected and ranked for compatible profiles. Therefore, the sum of absolute differences over all measured personality traits is calculated and the one, which is closest to zero but below a certain threshold is selected. Based on this ranking, action preferences are updated for the new user and β is set to the same value that the more experienced agent is using.

3. INTERACTION DETAILS

The goal of the demonstrator is twofold, first, we show how the individual PeSA agent supports a user, secondly, we show how the learning procedure is accelerated when new PeSA agents join an existing PeSA multi-agent system.

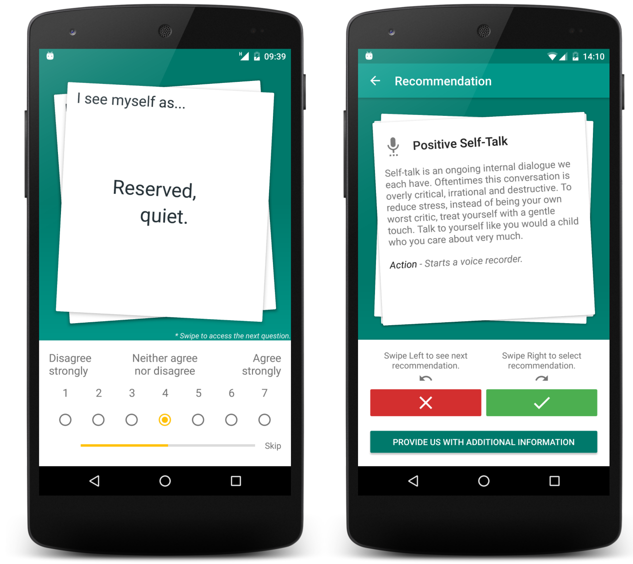

(a) Personality assessment **(b)** Recommendation

Figure 1: Screenshots showing parts of the PeSA app.

In the first step PeSA determines a user's personality using a 10-item personality assessment based on the work of *Gosling et al.* [4]. The personality is measured according to the Big-Five personality theory. Figure 1a shows a screenshot of the questionnaire. Based on the user's answers we calculate the personality profile, which is completed by the sex of the user and which is used to determine the initial strategy. The workflow proceeds with the stress diary, which is the landing page of the PeSA app. Here the user receives information about the goal of PeSA and the history of recorded stress events. If such an event is recognised, PeSA recommends a list of coping strategies. Figure 1b shows a screenshot of the recommendation dialogue; listing a coping-strategy from the cognitive restructuring category that shows a positive correlations with higher extraversion values (cf. [3, p. 1096]). For each advice, the user has the option to accept it (swipe right) or reject it (swipe left). These votes are used as reward signals to learn the actual Q-values for each action (coping-strategy).

To show how new PeSA agents enter the system, we use *ASGARD* [11], a graphical network monitoring tool. ASGARD allows us to show a visualisation of the multi-agent system at runtime. It is able to visualise and analyse the communication between JIAC agents and enables us to analyse the inner processes of PeSA agents.

4. FINAL REMARKS

PeSA combines several agent-techniques in order to help users to manage stress and to live healthy. In doing so, PeSA demonstrates that human-behaviour models improve the interaction with real human beings, i.e. making informed decisions during planning for joint activities.

We are currently evaluating this improvement through the A/B test feature that is provided by the *Google PlayStore*. Our goal is to empirically verify our hypothesis that personality-based information accelerates the learning stage.

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, 12th European Conference, EUMAS 2014, Prague, Czech Republic, December 18-19, 2014, Revised Selected Papers*, number 8953 in Lecture Notes in Artificial Intelligence, pages 3–18. Springer, April 2015.

[2] C. S. Carver and J. Connor-Smith. Personality and coping. *Annual Review of Psychology*, 61:679–704, January 2010.

[3] J. Connor-Smith and C. Flachsbart. Relations between personality and coping: A meta-analysis. *Journal of Personality and Social Psychology*, 93(6):1080–1107, 2007.

[4] S. D. Gosling, P. J. Rentfrow, and W. B. Swann Jr. A very brief measure of the big-five personality domains. *Journal of Research in Personality*, 37:504–528, 2003.

[5] U. Hapke, U. Maske, C. Scheidt-Nave, L. Bode, R. Schlack, and M. Busch. Chronic stress among adulty in Germany. *Bundesgesundheitsbl*, (5/6):1–5, 2013.

[6] W. B. Knox and P. Stone. Reinforcement learning from simultaneous human and MDP reward. In Conitzer, Winikoff, Padgham, and van der Hoek, editors, *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pages 475–482. IFAAMAS, 2012.

[7] M. Lützenberger, T. Konnerth, and T. Küster. Programming of multiagent applications with JIAC. In P. Leitao and S. Karnouskos, editors, *Industrial Agents – Emerging Applications of Software Agents in Industry*, pages 381–400. Elsevier, 2015.

[8] 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.

[9] R. Prada and A. Paiva. Human-agent interaction: Challenges for bringing humans and agents together. In *Proc. of the 3^{rd} Int. Workshop on Human-Agent Interaction Design and Models (HAIDM 2014) at the 13^{th} Int. Conf. on Agent and Multi-Agent Systems (AAMAS 2014)*, pages 1–10. IFAAMAS, 2014.

[10] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. Adaptive Computation and Machine Learning. MIT Press, May 1998.

[11] J. Tonn and S. Kaiser. ASGARD – A graphical monitoring tool for distributed agent infrastructures. In Y. Demazeau, F. Dignum, J. M. Corchado, and J. B. Pérez, editors, *Advances in Practical Applications of Agents and Multiagent Systems*, volume 70 of *Advances in Intelligent and Soft Computing*, pages 163–175. Springer, 2010.