

Using Reinforcement Learning for Personalization – Creation of an application to be used in pain therapy

Jana Kolly

Master Thesis in Mathematics

In the context of this thesis, one was interested in exploring the opportunities that reinforcement learning can offer to support the treatment of pain, as well as to alleviate accompanying symptoms. It has therefore been approached to develop a self-learning application, which uses techniques from psychological pain management, and thereby sends various stimulations to the user of the application. The main advantages such an application offers, are therapist-independency, full personalization, and availability at any given time. With a combination of reinforcement learning techniques and projection of stimulations triggering neural effects, the application would offer new potential in the area of pain management.

For this purpose, the task was analyzed and mapped into the setting of a Markov decision process (MDP). It had to be defined what was considered as an “agent” and “environment”, how we represent the “actions”, “states” and “rewards”, and what requirements our setting imposes on a learning algorithm. In a further step, it was also addressed what data would be used and how to represent it optimally. Thereby, several encoding methods were introduced. On these grounds, a suitable learning algorithm could be chosen: “SARSA(λ) with function approximation”. With this algorithm the following main criteria can be fulfilled: model-free learning, online learning, on-policy learning, handling of large action and state spaces, handling possibility of non-existence of global optimum, and handling sequential knowledge.

Based on the detailed analysis and modelling of the task, a prototype of the application was developed. The prototype provides an illustration of the core idea: it achieves the desired flow of events; the main functionalities are ensured; and there is already a functional underlying reinforcement learning algorithm, which is capable of learning a certain strategy and responding to user feedback. This prototype is considered as a starting point for further developments.

Supervisor: Prof. Dr. Christian Mazza