Courses in English
Course Description

Department	07 Computer Science and Mathematics
Course title	**Reinforcement Learning**
Course number	
Hours per week (SWS)	4
Number of ECTS credits	5
Course objective	Learning objectives 1 The students should be enabled to use models and methods from the field of reinforcement learning as a solution strategy in a wide variety of application scenarios in a methodologically correct and safe manner. The focus is on basic training. 2 In addition, students should be enabled to quickly familiarize themselves with new and current reinforcement learning models and methods. Technical & methodological competencies 1 To become familiar with and understand the theoretical foundations of reinforcement learning as well as the practical application of algorithms to a wide variety of problems. 2 Learning skills to implement and apply the algorithms in Python (one of the leading programming languages in the field of machine learning). 3 Skills to understand, implement, apply and evaluate different reinforcement learning models. Interdisciplinary competencies 1 Group work: Students work on subareas independently and in small groups. 2 Exercises: Students learn practical work within the Data Science process using examples
Prerequisites	Basic programming skills (preferably in Python), basic knowledge of probability theory as acquired, for example, in the DC Bachelor module of the same name; basic knowledge from the area of stochastic processes (Markov chains, Markov decision processes) and Deep Learning would be advantageous
Recommended reading	* Sutton, Richard S. & Barto, Andrew G. (2020). Reinforcement Learning. The MIT Press. * Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. The MIT Press. * Lapan, M. (2020). Deep Reinforcement Learning. Mitp Verlag. * Foster, D. (2019). Generative Deep Learning [Kapitel 8]. O'Reilly
Teaching methods	projector, whiteboard, Jupyter notebooks
Assessment methods	module work, written exam
Language of instruction	English
Name of lecturer	Prof. David Spieler
Email	david.spieler@hm.edu
Link	
Course content	Reinforcement learning is the third pillar of machine learning, alongside supervised and unsupervised learning. In contrast to the other two sub-areas, learning here is not based on a previously collected data set but mostly online in the form of a virtual agent that has to prove itself in an environment. The application areas are numerous and include autonomous driving, practical optimization and control problems, games, and also specialized domains such as protein folding in bioinformatics. In this course, we first clarify the theoretical foundations such as Markov chains, Markovian decision processes, and planning using dynamic programming. This is followed by Temporal-Difference Learning, Policy Gradient methods, and an introduction to the use of function approximators. Furthermore, the interplay between learning and planning as well as the dilemma between exploration and exploitation will be discussed. Optionally, more advanced topics, such as imitation learning, will follow.
Remarks	