Master thesis topic (Masterproefonderwerp)	Deep reinforcement learning: with applications to robotic control
Promoters (Promotoren)	prof. dr. ir. Aleksandra Pižurica – Department of Telecommunications and Information Processing
Mentors (Begeleiders)	Srđan Lazendić – Department of Electronics and Information Systems
Contact person (Contactpersoon)	Srđan Lazendić
Number of master theses (Aantal masterproeven)	1
Number of students (Aantal studenten)	1
Academic year (Academiejaar)	2021-2022
Keywords (Trefwoorden)	Deep learning, Reinforcement learning, Deep reinforcement learning, Robotics
Location – (Locatie)	Technicum, De Sterre, thuis

References (Referenties):

1. K. Arulkumaran et al. Deep Reinforcement Learning: A brief Survey, 2017.
2. Gym: A toolkit for developing and comparing reinforcement learning algorithms. https://gym.openai.com/docs; https://github.com/openai/gym

Description of the problem (Probleemstelling)	Deep reinforcement learning (DRL) is a rapidly evolving research track in the field of machine learning, often indicated as the technology that is expected to revolutionize the field of artificial intelligence (AI). One of the primary goals of the field of AI is to develop fully autonomous agents that learn optimal behaviour by interacting with their environments, and improve over time through trial and error. A mathematical framework for such experience-driven autonomous learning is known as reinforcement learning (RL). Although RL had some successes in the past, previous approaches lacked scalability (due to complexity issues), which was limiting their applicability to relatively low-dimensional problems that are rarely capturing complex real-life scenarios. The rise of deep learning provided new tools to overcome these problems resulting in a powerful new framework – deep reinforcement learning. DRL has already demonstrated remarkable results in various applications, ranging from playing video games to indoor navigation and with various agents including robots and softbots. For example, control policies for robots can be learned directly from camera inputs in the real world.
Goal of the thesis (Doelstelling)	The first goal of this thesis is to study and understand well the principles involved in deep reinforcement learning. Secondly, student should implement and empirically evaluate deep RL architectures making use of the OpenAI Gym toolkit for developing and comparing reinforcement learning algorithms. As the main task the student should study multiple variants of the recently proposed Deep Q-Network algorithm. Different ways to prioritize the awards could be introduced and then used with the existing algorithms to solve different tasks. These tasks will be chosen in the agreement with the student. An optional direction one could go is to use recently proposed Double Deep Q-Network algorithm which can be used to control an industrial robotic arm in a simulated environment. The goal here is to control the movements of the robotic arm to pick up items in front of it. One should try to develop a framework such that different robots perform different tasks with different algorithms. Finally, the learned strategy should allow a physical robot to perform the same tasks. In this way, the introduced RL algorithm could potentially be used to control industrial robots in e-supermarkets in the future. At the beginning of the semester, a minicourse will be organized for a student to put him/her on the right track and to familiarize him/her with the topic. Furthermore, the existing code and literature will be made available to the student.

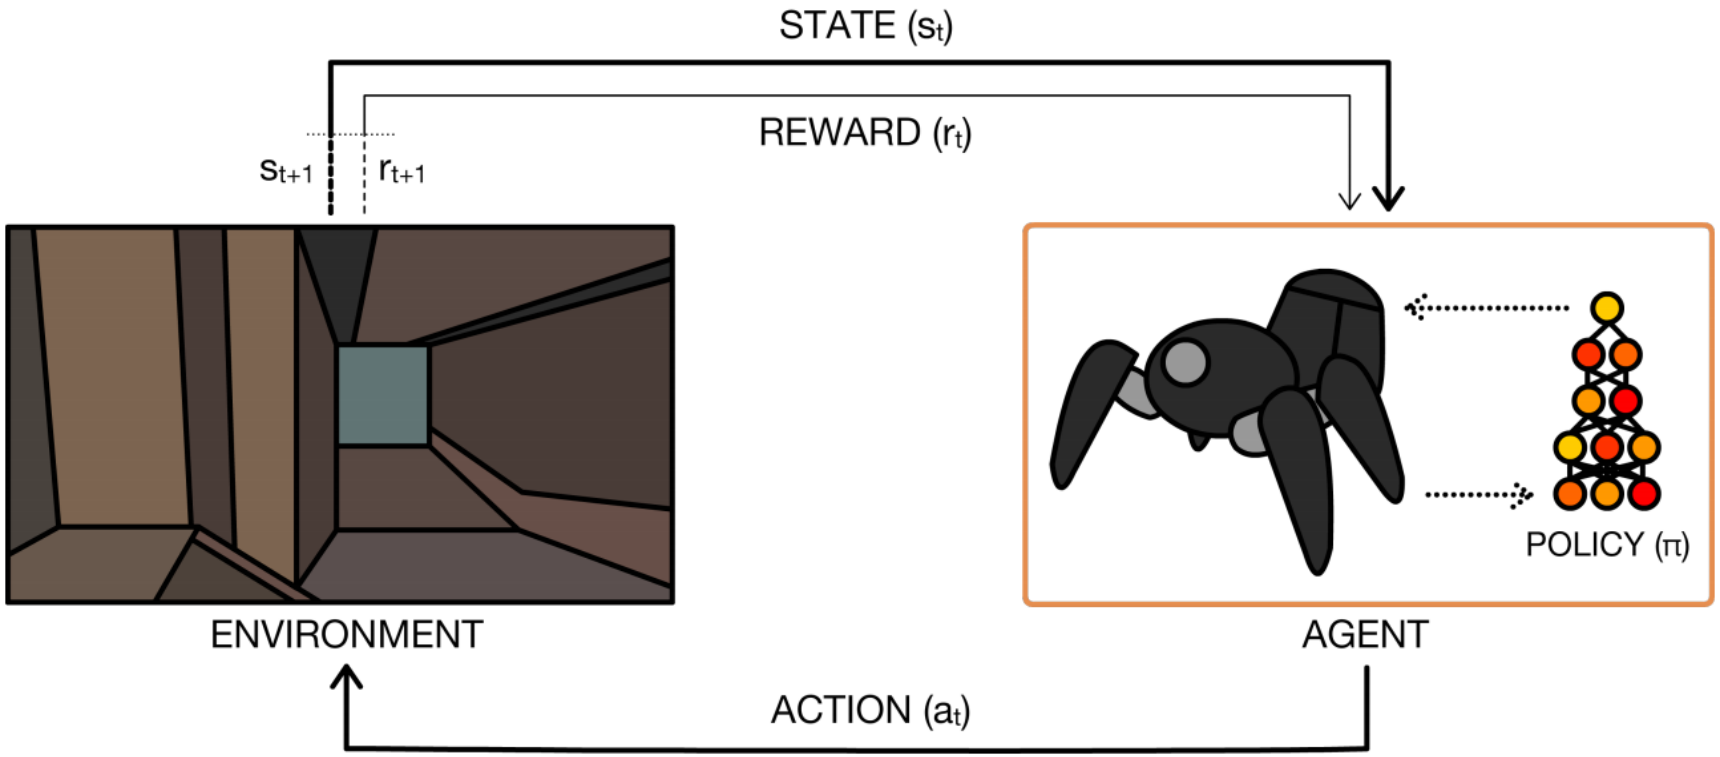

Figure 1: The perception-action-learning loop. Source: Kai Arulkumaran et al. Deep Reinforcement Learning, Signal Processing Magazine, November 2017.