3 ECTS

15D013

Topics in Big Data Analytics II

Reinforcement learning

Overview and Objectives

Reinforcement learning (RL) is the model-based theory of sequential decision-making under uncertainty. The field of reinforcement learning research looks back to a long history with several periods of high and moderate success; recent years saw an unprecedented increase in attention to the field both in the mainstream media and within the machine learning/AI research community. Indeed, RL techniques underly many of the breakthrough successes of AI research that made headlines in recent years (e.g., the ATARI and Go players developed at Google DeepMind). Within the machine learning community, several successful and influential researchers now cite the task of general reinforcement learning as the ultimate challenge for all of machine learning research. Very recently, RL has been highlighted as one of the "Top 10 Breakthrough Technologies of 2017" by the MIT Technology Review.

This short course gives an overview of the fundamental techniques of reinforcement learning starting from the classic temporal-difference methods through approximate dynamic programming all the way to recent developments in deep reinforcement learning. The goal is to provide a strong understanding of the most common methods and provide a basic algorithmic toolbox for building RL systems. Besides familiarizing students with these tools, the course puts a strong emphasis on highlighting the crucial challenges that set RL problems apart from other machine learning problems. Students taking the course are expected to gain the capability to identify and tackle such challenges in various application domains.

Required Activities

Attendance at the theory classes and handing in a report of the group project.

Materials

- R. S. Sutton and A. G. Barto (2018): Reinforcement Learning: An Introduction (second edition). MIT Press.
- Csaba Szepesvári (2010): Algorithms for Reinforcement Learning. Morgan & Claypool.

3 ECTS

15D013

Topics in Big Data Analytics II

Conversational Artificial Intelligence

Overview and Objectives

This course will introduce students to the concepts, techniques and challenges involved in the design and development of conversational agents. Following a brief overview of the dialogue system development, the course will analyse the most recent research for building intelligent conversational agents. Existing commercial bot development platforms and open source alternatives will be introduced and compared. Finally, the last part of the course will be dedicated to hands-on development of a chatbot using the reviewed tools and techniques.

Required Activities

Attendance to classes and group project (evaluation: 100% group project).

The students will work in small teams to develop a fully functional conversational agent for the selected use case. They will use the tools and techniques introduced in the course to develop the final project. The quality of the dialogue systems will be evaluated by every team according to different subjective evaluation metrics. Final grades will take into account the technical contribution, the quality of the generated dialogues and business value.

Materials

- Celikyilmaz A., Deng L., Hakkani-Tür D. (2018) Deep Learning in Spoken and Text-Based Dialog Systems. In: Deng L., Liu Y. (eds) Deep Learning in Natural Language Processing. Springer, Singapore. https://link.springer.com/chapter/10.1007/978-981-10-5209-5_3
- RASA: Open source machine learning tools for developers to create contextual AI assistants and chatbots that go beyond answering simple questions. https://rasa.com/
- ParlAI: A unified platform for sharing, training and evaluating dialogue models across many tasks. http://parl.ai/

3 ECTS

15D013

Topics in Big Data Analytics II

Internet of Things (IoT)

Overview and Objectives

Internet of Things (IoT) is the next ground for Analytics. IoT is spread across the length and breadth of the industry. From consumer electronics, automobiles, aviation, energy, oil and gas, manufacturing, banking, and so on, almost every industry is benefiting from IoT. Data can come from different devices or sensors, data can have different characteristics, diverse latencies, different importance or be plagued with missing values or be exposed through vulnerabilities in the pipeline security.

We, data scientists, are used to developing our solutions on cloud computing architectures, where a server takes care of the entire computation. However, with the IoT ecosystem, this is not always efficient neither cost-effective so we need to rethink our paradigms. Here is where the “edge computing” (the devices/sensor layer) presents itself as a solution. Edge computing is an architecture where the process of data, applications, and services are pushed away from the centralized cloud to the logical extremes of the network. These devices are equipped with the enough computing power and data storage facilities to fulfill the task. After computing, only the rich and condensed yet reusable data is transmitted back to the cloud. Low cost alternative storage can also be considered for the remaining data. The objective of the course is to provide you with a hands-on experience on IoT Analytics that will leverage your already obtained knowledge throughout the Master and equip you with basic skills to deploy your first IoT Analytics solutions.

Required Activities

Final project between 2 or 3 people. The final project grade is based on the IoT solution built and the way students present their concept.

Course evaluation

The final overall course grade will be the simple average of the grades of each part.