VUB SCIENCES & BIOENGINEERING SCIENCES

The Research Group

Artificial Intelligence Lab

has the honour to invite you to the public defence of the PhD thesis of

Tim BRYS

to obtain the degree of Doctor of Sciences

Title of the PhD thesis:

Reinforcement Learning with Heuristic Information

Promotor:

Prof. dr. Ann NOWÉ

The defence will take place on

Friday 2 December 2016 at 17:30 h

in Auditorium D.2.01 at the Campus Etterbeek of the Vrije Universiteit Brussel, Pleinlaan 2 - 1050 Elsene, and will be followed by a reception.

Members of the jury:

Prof. Dr. Wolfgang DE MEUTER (chairman)
Prof. Dr. Bernard MANDERICK (secretary)
Prof.Matthew E. TAYLOR (co-promotor, IRLL, Washington State Univ.)
Dr. Bart JANSEN
Prof. Dr. Enda HOWLEY (NUI Galway, Ireland)
Prof. Dr. Peter VANPLEW (Univ. Australia)

Curriculum vitae

Tim's research has spanned multiple topics, but he currently focuses on the incorporation of prior knowledge in reinforcement learning, looking at ways to efficiently incorporate demonstrations, prior experiences, domain knowledge, etc. He has presented his work at top venues such as AAAI, IJCAI and AAMAS.

Abstract of the PhD research

Artificial intelligence and machine learning specifically are (at least in part) about automating human tasks -- piecing together cars in a factory, recommending books on a webshop, making decisions on the stock market, and, more futuristically but very realistic, delivering your packages, driving your car and even cooking your meals. Experimental successes in the past few years have hinted at the potential of reinforcement learning at tackling such complex tasks. These learning systems are able to learn to solve a task from scratch, but quite often external information is used to guide the learner. A great example is the reinforcement learning system that made headlines as being the first AI system to beat a professional player of the infamously hard board game Go. The system used reinforcement learning guided by a large database of human moves. This thesis focusses on which kinds of information can be available to a reinforcement learner, and how it can use this information.

We investigate how information learned in a previous task can be used in a new task that is similar, analysing the effect of high and low quality information. We investigate how human demonstrations of the task can help the reinforcement learner get the hang of it faster, and how the number of these demonstrations affects learning. And finally, we develop a way of integrating the information from various sources so that the learner can combine this information as good as possible.