SCIENCES &
BIOENGINEERING SCIENCES

The Research Group

Artificial Intelligence lab

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

Pieter LIBIN

to obtain the degree of Doctor of Sciences

Title of the PhD thesis:

Guiding the mitigation of epidemics with reinforcement learning

Promotors:

Prof. dr. Ann Nowé (VUB)
Prof. dr. Philippe Lemey (KU Leuven)

The defense will take place online on

Wednesday, April 15, 2020 at 13:00

and will be broadcasted live via YouTube as specified on [1], a reception will be held after the COVID-19 lockdown.

[1] http://ai.vub.ac.be/team/pieter-libin/

Members of the jury:

Prof. Dr. W. DE MEUTER (VUB, chairman)
Prof. Dr. T. LENAERTS (VUB, secretary)
Prof. Dr. A. NOWE (VUB, promotor)
Prof. Dr. P. LEMEY (KUL, promotor)
Prof. Dr. D. MAES (VUB)
Prof. Dr. B. DE BAETS (UGent)
Prof. Dr. E. HOWLEY (NUI Galway, Ireland)

Curriculum vitae

Pieter Libin (born 02/12/1982) obtained a Master in Informatics at the Vrije Universiteit Brussel in 2014. He started his PhD in the Artificial Intelligence Lab at the VUB under the supervision of Prof. Dr. Ann Nowé and Prof. Dr. Philippe Lemey (KU Leuven), supported by an FWO PhD grant. His research focuses on developing a new methodology to automatically learn prevention strategies to mitigate epidemics through the use of reinforcement learning. His work resulted in 27 peer-reviewed manuscripts, of which many were published in high-ranked journals (e.g., Bioinformatics, PLOS Pathogens, PLOS NTD) and top computer science conferences (e.g., AAMAS, ECML).

Abstract of the PhD research

Epidemics of infectious diseases are an important threat to public health and global economies. To properly understand epidemic processes, and to study emergency scenarios, epidemiological models are necessary. Yet, the development of prevention strategies, which need to fulfil distinct criteria, remains challenging. For this reason, it is important to study how optimization techniques can be used to support decision makers. In this thesis we explored an approach based on reinforcement learning, a field within artificial intelligence, to automatically learn prevention strategies.

We investigated two main lines of research. Firstly, we studied the decision problem where a number of prevention strategies has been defined, and the decision makers need to determine which of the strategies is most efficient, by evaluating the models in a complex and computationally demanding epidemiological model. To perform this evaluation efficiently, we investigate the use of algorithms in the field of reinforcement learning that are grounded in the Bayesian uncertainty framework. This approach enables us to learn faster *and* to quantify the uncertainty of the decisions. To make this possible, we adapt existing algorithms and created new algorithms. Furthermore, we provide theoretical insights in how these algorithms operate. Secondly, we extend this approach such that we can learn adaptive strategies in an epidemiological model. This means that, rather than comparing preventive strategies, we attempt to learn which subsequent steps are necessary to act optimally over time, i.e., by considering the state of the epidemic. Since the state space of the epidemiological models that are necessary to investigate versatile prevention strategies is huge, we need to represent this space such that the reinforcement learner can use this in its learning process. To do this, we use deep reinforcement learning.

We evaluate both research trajectories in the context of pandemic influenza, a pathogen that has made many victims in the past. Our experiments show that our first research trajectory is very useful to evaluate prevention strategies. Furthermore, we show that these techniques are also useful to support other complex decision problems that involve computationally demanding models. In the experiments to validate the second research trajectory, we created a specific model to investigate school closure policies in case an influenza pandemic emerges. Through experiments, we show that our learning technique approximates the optimal strategy. Finally, we investigate whether there is a collaborative advantage when designing school closure policies. We formulate this research question as a multi-agent problem and solve it using deep multi-agent reinforcement learning techniques.