ISI WSC 2019 Short Course Programme

COURSE TITLE :	**SC2 - Deep Learning for the Environmental Sciences**
DURATION :	1.5 Days
DATE :	16 – 17 August 2019
VENUE :	Sasana Kijang
REGISTRATION FEES :	Developed Country MYR 1,760 (Approximately EUR 370) Developing Country / Student* MYR 1,140 (Approximately EUR 240) ** For student, proof of enrolment is required*

INSTRUCTOR 1

Dr. Nathaniel K. Newlands
Agriculture and Agri-Food Canada (Government of Canada)
Canada

Nathaniel's work addresses public-good food-water-energy nexus issues and tackles broad, integrated, complex global problems (e.g., climate change) to help support and advance global sustainable development. He completed a B.Sc. (Hon.) in mathematics and physics, a M.Sc. in astrophysics, and a Ph.D. in resource management and environmental studies, with post-doctoral training at the Department of Mathematics in Applied Mathematics (IAM) and the Pacific Institute of the Mathematical Sciences (PIMS) based at University of British Columbia (UBC) in mathematical ecology. He currently works as a Research Scientist in Environmental Health within the Science and Technology Branch of Agriculture and Agri-Food Canada.

INSTRUCTOR 2

Dr. Etienne Lord
Agriculture and Agri-Food Canada (Government of Canada)
Canada

Etienne's work addresses microbiome identification, classification and evolution as well as the identification of biological network hubs. He also focus on multi-omics and drone imaging analysis to tackle complex environments evolution and sustainability. He completed a B.Sc. in biology, a M.Sc. in veterinary sciences, and a Ph.D. in computer sciences. He did post-doctoral training in machine learning and biology at the University of Montreal (UdeM) and later at Laval University in Quebec city (UL). He currently works as a Study Leader in bioinformatics within the Science and Technology Branch of Agriculture and Agri-Food Canada.

COURSE DESCRIPTION

Modern machine learning and artificial intelligence procedures provide fast and automatic learning of hidden dependencies and structures. Challenges with modelling the dynamics of complex ecosystems also brings to the fore statistical problems of analyzing massive, multi-resolution, multi-source data with a nonstationary space-time structure – where deep learning may help provide an integrated solution. The course will start with an introduction to artificial intelligence, machine-learning and neural networks. An overview of network model architecture, learning structures (i.e., multilayer perceptrons (MLPs), convolutional, recurrent, deep) and supervised/unsupervised optimization and learning algorithms will be discussed. State-of-the-art toolsets and code libraries for tackling deep learning tasks will be surveyed and applied (e.g., h20, Deepnet, Keras, TensorFlow). A set of environmental science examples will be used to illustrate core concepts, features, and benefits. These case examples will include: predicting wine quality, plant phenomics, crop yield, extreme events, and air pollution.

SYLLABUS

Relevance

Modern machine learning and artificial intelligence procedures provide fast and automatic learning of hidden dependencies and structures. It offers many advantages for operational predictive processes and systems by providing higher automation/speed, accuracy, scalability, and strategic foresight. The benefits of deep learning is being realized across a wide variety of real-world applications from speech and image recognition, smart application communications, communications, medicine/drug discovery, manufacturing, equipment failure and flaw detection, transportation to threat/fraud/risk management, and environmental science. Challenges with modelling the dynamics of complex ecosystems also brings to the fore statistical problems of analysing massive, multi-resolution, multi-source data with a nonstationary space-time structure – where deep learning may help provide an integrated solution. In particular, reinforcement algorithms that incorporate deep learning help to predict within the so-called 'delayed return' environment where it can be difficult to understand which action leads to which outcome over many time steps. Here, deep learning may enable near-real-time forecasting and strategic foresight in the more ambiguous 'real-world' environment, whereby one chooses from an arbitrary number of possible actions to achieving goals.

Organization

The course will include a presentation and practical component. The presentation will provide students/attendees with required background, fundamentals, jargon in Deep learning/AI from a multidisciplinary perspective. Following the morning presentation, students/attendees will be presented with a selected set of practical deep learning examples with a description, software and coding guidance and tips provided. Examples will include showcased/code snippets running in both R and python using the RStudio and Spyder Integrated Development Environments (IDE's). The instructor will use their computer to first showcase examples, and utilize a 'docker container' and/or 'cloud-based' link to access a computer to run examples with larger data sets via their class laptop. A set of software installation instructions will also be prepared and handed out.

Course Objectives

1. To introduce students/attendees to the field of Deep learning and artificial

intelligence by providing a state-of-the-art, concise summary of its background, fundamentals, and definitions from a multidisciplinary perspective.

2. To orient students/attendees using prepared coding examples in applying deep learning to environmental science across a range of different applications.
3. To provide students/attendees with the required theory, knowledge of benefits and drawbacks, and ability to build their own applications in their research, development or technology-transfer activities.

Course topics:

Background, fundamentals, definitions

Practical tips, guidance, software install and coding examples (R and python)

Application examples that will include: weather forecasting, fruit image recognition, wine quality, crop yield/actuarial, plant phenomics

Software:

Participants should install the latest version of R Statistical Software and R Studio Desktop, along with Anaconda Python 3.6.6. Keras will be used for deep learning. The scikit-learn, panda and numpy python libraries will also be used.

Additional details and information, including required R libraries and example codes will be provided to participants either in advance of the workshop (if attendees names and emails can be provided in advance to the instructor), or alternatively, at the workshop.

TARGET AUDIENCE

This short course is relevant to anyone with an interest in Environmental Science and Sustainability, Data Science, and Predictive Analytics. Previous training and experience in predictive modelling, machine-learning and R Statistical Language are an asset.