KU LEUVEN

Edge Intelligence Unleashed: Mastering Embedded AI for Industry Advancement
Summer School

16, 17 and 18 September 2024
KU Leuven - Brugge

PUC | CONTINUE

Artificial intelligence is gradually becoming an integral part of production chains and business processes. In recent years, more and more AI tasks are being performed at the device level rather than on powerful cloud servers. Embedded AI offers numerous benefits and can be used for a multitude of tasks (such as process automation, improving customer service, advanced analytics, etc.), resulting in an exponential growth in the number of use cases.

During this exclusive summer school, we will examine in detail how machine and deep learning can be applied at the device level. Which components and frameworks are needed for edge processing? Both hardware and software will be covered, as well as techniques for optimizing network architectures. In addition, there is plenty of room to experiment under the guidance of experts and to bridge the gap to your own business case. After the three-day program, you will have a solid theoretical foundation regarding embedded AI, supplemented with practical tools to use in your own work context through the hands-on approach of the training.

Objectives

1. Gain a comprehensive understanding of machine and deep learning and their applications in embedded systems.
2. Learn about the necessary components and frameworks for edge processing and how to optimize network architectures.
3. Understand the hardware and software required for embedded AI, and the different techniques available for optimizing embedded machine learning algorithms.
4. Develop practical skills using tools, such as Edge Impulse and TensorFlow Lite for implementing embedded AI.
5. Learn how to apply embedded AI to real-world business cases, and how to bridge the gap between theory and practice.
6. Understand the potential benefits and limitations of embedded AI, and how to evaluate the suitability of different techniques for specific use cases.

16 September 2024
Session 1: Introduction

Prof. Mathias Verbeke

Setting the context for the rest of the summer school:
- Introduction to machine learning & deep learning
- Embedded machine learning
- Use cases and company pitches

Session 2: Workshop "Edge Impulse"

Dr. Jonas Lannoo

In this workshop, we use "Edge Impulse", the online edge AI platform: a low-code environment to develop, train and deploy AI models on edge devices. It is widely applicable to quickly train a neural network and place the result on a microcontroller or any other edge device. With a simple but intuitive example, we discover the broad possibilities of this powerful software tool and create our first functional edge-deployed AI model.
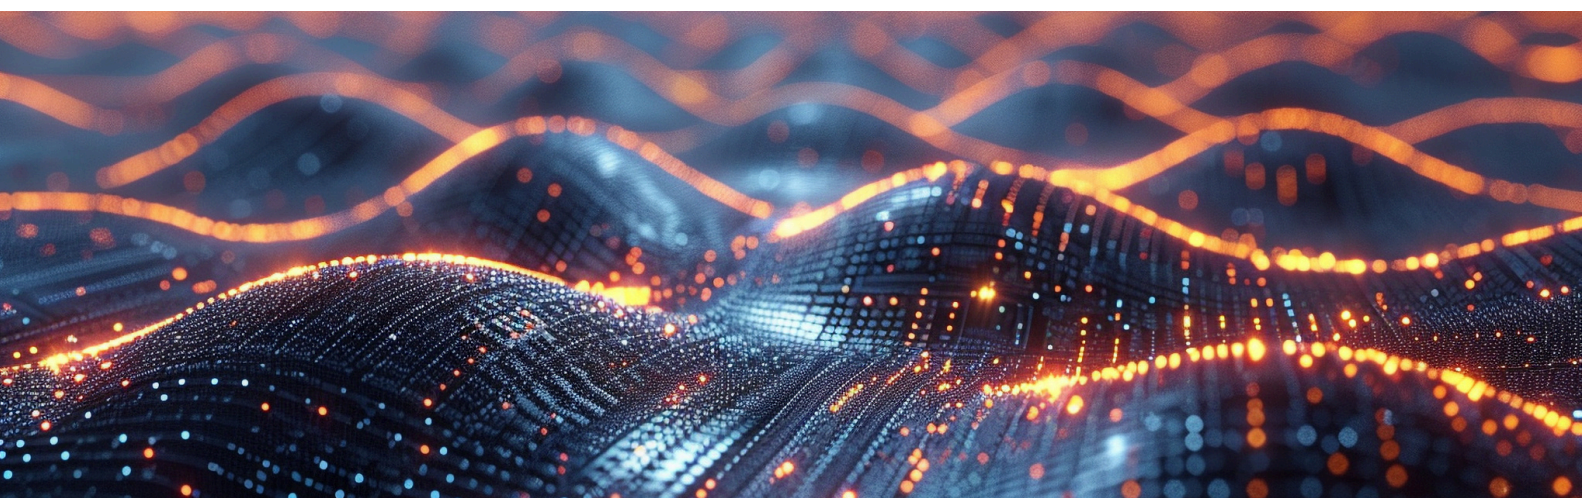
17 September 2024
Session 3 and 4: TinyML (running ML models on microcontrollers)
Prof. Toon Goedemé & Prof. Peter Karsmakers

Theory: TinyML design flow for inference:
- Model compression: hardware-friendly model architectures, quantisation, pruning, knowledge distillation
- Model optimization for edge hardware
- Model deployment

Workshop: Compress and optimize a computer vision application for deployment on a microcontroller. Design validation using an ESP-EYE microcontroller.

18 September 2024
Session 5: Workshop FPGA

Wouter Hellemans

In this session, we will establish the landscape of ML on Field Programmable Gate Arrays (FPGAs) and investigate how CNNs can be accelerated on FPGAs for high-throughput applications. Specifically, we will explore a network security case by developing a CNN-based intrusion detection system. During the session, we will make use of the open-source libraries Brevitas and FINN to quantize our model and convert it to specialized hardware.

Session 6: Distributed AI

Prof. dr. Hans Hallez

Focus on distributed working with artificial intelligence:
- How can AI algorithms be scaled and distributed over a network of systems?
- Deep dive into concepts like Edge Computing, Big Data, Machine Learning

There will be ample time to wrap up the summer school with lessons learned and a closing drink.

Target audience

The target audience for this training is professionals and experts interested in applying machine learning and deep learning to embedded systems in their respective industries. This may include hardware and software engineers, data scientists, machine learning engineers, embedded system designers, functional analysts, R&D managers and researchers working in fields such as computer vision, natural language processing, robotics, and IoT. The training is designed for individuals who already have a basic understanding of machine learning and deep learning and are looking to expand their knowledge and skills in applying these techniques to embedded systems.

Basic notions of data structures and algorithms, signal processing, convolutional neural networks and languages C++, Python & Linux (command line instructions) are needed to follow the programme adequately.

Date and location

- KU Leuven - Brugge (Spoorwegstraat 12, 8200 Brugge)
- 16, 17 and 18 September 2024

Registration

- Register online before 9/09/2024.
- Because of the hands-on character of the training, the number of participants is limited to 20.
- The fee for the series is 1200 euro. For the second participant of the same organisation, a discount of 20% is applied.
- After registration, you will automatically receive the invoice for your participation.

KMO

Save on your participation costs via the kmo-portfolio.
Click here for more information. Our approval number is DV.O102270. When submitting your application, you choose the theme Innovation and the Artificial Intelligence advice.

This programme is organised by PUC - KU Leuven Continue with the support of VAIA - Vlaamse AI Academie.

PUC - KU Leuven Continue
KU Leuven Kulak
E. Sabbelaan 53 bus 7643 - 8500 Kortrijk
+32 56 24 61 84 - puc@kuleuven.be
puc.kuleuven.be

Door mij in te schrijven, verklaar ik mij ermee akkoord dat de gegevens die ik verstrek gebruikt worden om mij te contacteren in de context van deze opleiding en voor alle nuttige opvolging ervan.