

FPGAs for AI: Build & Deploy Your Own Edge AI Accelerator

Half-Day Hands-On Workshop · 3 June 2026 · School of Architecture, Computing and Engineering, University of East London, Docklands Campus · 15–30 participants ·

ABOUT THIS WORKSHOP

AI is transforming every industry but running AI on GPUs is expensive, power-hungry, and often too slow for real-world deployment. **FPGAs offer a better way.** Configurable to match any AI model precisely, they deliver faster inference, lower power, and greater flexibility without writing a single line of hardware code.

In this intensive half-day workshop you will use **Intel Altera Agilex FPGAs** and the **Altera FPGA AI Suite**, a free push-button toolchain that converts any pre-trained model into a hardware accelerator in hours. You will deploy a real image-classification model, estimate its live power, and compare it head-to-head against CPU and GPU in speed, efficiency, and energy.

WHAT YOU WILL BE ABLE TO DO

Understand FPGA AI

Explain why FPGAs excel at inference: flexible datapaths, deterministic latency, and low power.

Navigate the Toolchain

Use FPGA AI Suite end-to-end: OpenVINO conversion, architecture optimisation, IP generation, and Quartus synthesis.

Apply Quantisation

Select INT8 / BFP / FP11–FP16 precision and understand the accuracy vs. resource trade-off.

Run FPGA Emulation

Execute a compiled FPGA AI model in software emulation and validate inference via OpenVINO runtime.

Estimate & Compare Power

Use Quartus Power Analyser; benchmark CPU and GPU; compute FPS/W for each platform.

Make the Sustainability Case

Translate measured power figures into annual energy and argue for FPGA in sustainable AI.

WHO SHOULD ATTEND

- PhD & Postdoc Researchers in AI / CE / Sustainable Computing
- ML Engineers & Data Scientists curious about hardware
- Hardware & Embedded Engineers adding AI acceleration
- MSc Students in EE, CS, or AI
- Industry Practitioners (IoT / Robotics / Smart Energy / Automotive)

PREREQUISITES

Minimal: A basic understanding of neural networks and familiarity with Python. *No prior FPGA, HDL, or hardware-design experience required.* All tools run in a pre-built Docker image or on Intel DevCloud, no installation needed on the day.

HALF-DAY SCHEDULE · 4.5 HOURS

09:00–09:20	Welcome & The Big Picture The AI power problem, the FPGA opportunity, and today's plan.
09:20–10:00	How Agilex FPGAs Work Agilex 5 & 7 architecture, AI Tensor Blocks, and neural-network mapping. Talk + Q&A.
10:00–11:00	Lab A · From Model to Hardware IP ResNet-50 → OpenVINO → FPGA IP core. Architecture Optimizer & INT8 vs BFP precision comparison.
11:00–12:00	Lab B · Synthesis & Power Analysis Quartus Prime: integrate IP, synthesise bitstream, run Power Analyser. No physical board needed.
12:00–13:00	Lab C · Emulate, Estimate & Compare FPGA software emulation vs. CPU vs. GPU. Build a three-way FPS/W comparison table.
13:00–13:30	Sustainability Exercise & Wrap-Up Translate power figures. Key takeaways & open Q&A.

WORKSHOP ORGANISERS

Prof. Amirali Baniasadi Workshop Organiser

Full Professor, ECE, University of Victoria, Canada

Leads a research group in computer architecture, ML hardware, and embedded systems. Extensive publication record spanning processor design, network-on-chip, and applied ML. Provides overall academic leadership for this workshop.

Dr. Meisam Abdollahi Presenter

Embedded System Designer & Postdoctoral Fellow,
Barco Visual Solutions / UVic, Canada

PhD in Computer Architecture (University of Tehran). Industry expertise in AI-accelerated video pipelines on Intel Altera Agilex 7 FPGAs (8K streaming, HDMI 2.1, PCIe, Jetson AGX Orin). 40+ journal & conference publications; 10+ years university teaching across VLSI, Computer Architecture, and Embedded Systems.

CONDITIONAL DELIVERY

Please note that the workshop will run **subject to sufficient registrations**. Promotion will initially be **conditional**, and final confirmation will depend on reaching the required number of participants.

