Joint ICTP-IAEA School on Detector Signal Processing and Machine Learning for Scientific Instrumentation and Reconfigurable Computing

Introduction to Machine Learning and Edge AI

smr4110 | Trieste, Italy - 2025

Romina Soledad Molina, Ph.D.

IAEA
International Atomic Energy Agency
Atoms for Peace and Development

The Abdus Salam International Centre for Theoretical Physics

Outline

- Introduction.
- AI beyond 2025.
- What is Edge AI?
- What the Literature Says.
- Integrating Every Layer of Edge AI Design.
- Case Study for FPGA-based Edge AI: MNIST Binary Classification.

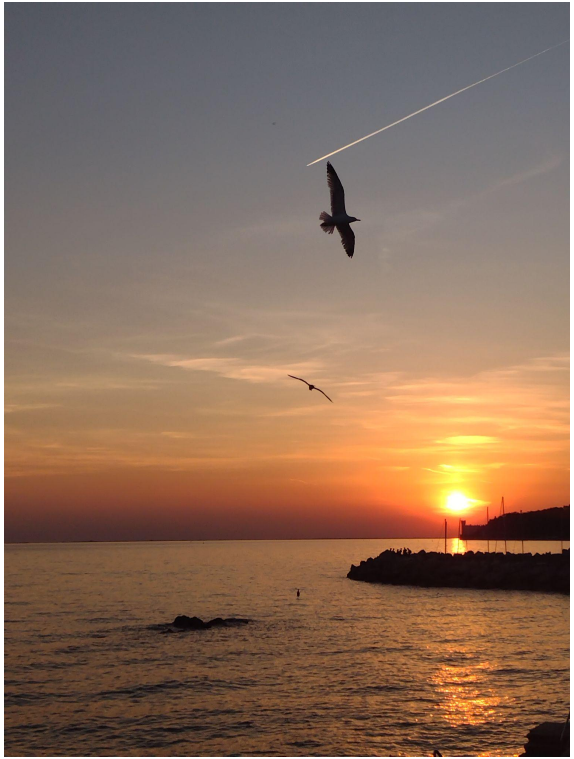

Introduction

Introduction

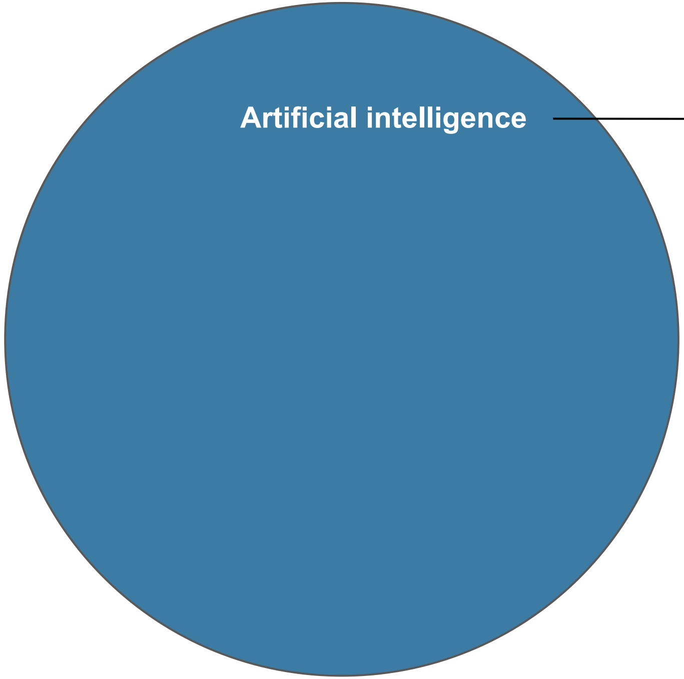

A machine/system capable of imitating human behavior

Introduction

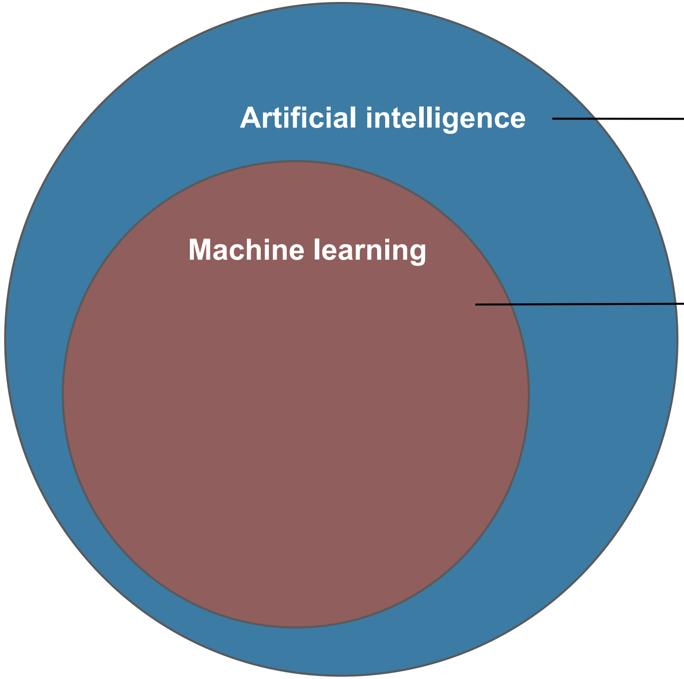

A machine/system capable of imitating human behavior

Mathematics and statistics to learn from data.

Introduction

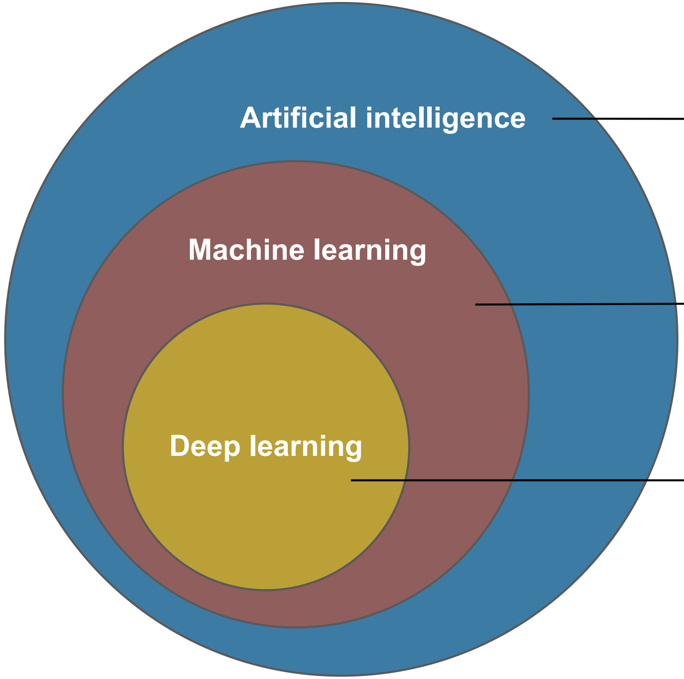

A machine/system capable of imitating human behavior

Mathematics and statistics to learn from data.

Neural networks multi-layer

Introduction

- **1950** – Alan Turing test: if a machine tricks human into thinking it's a human, then it has intelligence.
- **1956** – Arthur Samuel, working at IBM, introduces the term "machine learning."
- **1964** – Joseph Weizenbaum created the first chatbot ELIZA.
- **1997** – Deep Blue, an AI developed by IBM, defeated the world chess champion, Garry Kasparov.
- **1999** – Sony released the AIBO, an entertainment robot designed to mimic the behavior of a pet dog.
- **2002** – iRobot released the Roomba, a revolutionary autonomous robotic vacuum cleaner.
- **2011** – IBM's AI system Watson competed on the television game show Jeopardy!.
- **2012** – Deep learning achieved a major breakthrough with the AlexNet model.
- **2017** – The Transformer model was introduced in the paper "Attention is All You Need".
- **2020** – OpenAI released GPT-3 (Generative Pretrained Transformer 3).
- **2022** – ChatGPT was released by OpenAI.
- **2024** – John J. Hopfield and Geoffrey E. Hinton receive the Nobel Prize in Physics.
- **2025** – Launch of DeepSeek-V3.

Introduction

Why has the explosion of AI happened in recent years? Three main components

Big data

Introduction

Why has the explosion of AI happened in recent years? Three main components

Big data

Software

Introduction

Why has the explosion of AI happened in recent years? Three main components

Big data

Software

Hardware

Introduction

Why has the explosion of AI happened in recent years? Three main components

Big data

Software

Hardware

The massive amount of data generated every day—from social media, sensors, and online activity—provides the fuel that AI systems need to learn and improve

Introduction

Why has the explosion of AI happened in recent years? Three main components

Big data

Software

Hardware

The massive amount of data generated every day—from social media, sensors, and online activity—provides the fuel that AI systems need to learn and improve

New algorithms, neural network architectures, and open-source frameworks have dramatically improved AI's capabilities and made development more accessible.

Introduction

Why has the explosion of AI happened in recent years? Three main components

Big data

The massive amount of data generated every day—from social media, sensors, and online activity—provides the fuel that AI systems need to learn and improve

Software

New algorithms, neural network architectures, and open-source frameworks have dramatically improved AI's capabilities and made development more accessible.

Hardware

Advances in GPUs, TPUs, and cloud computing have made it possible to train large AI models faster and more efficiently than ever before.

AI beyond 2025

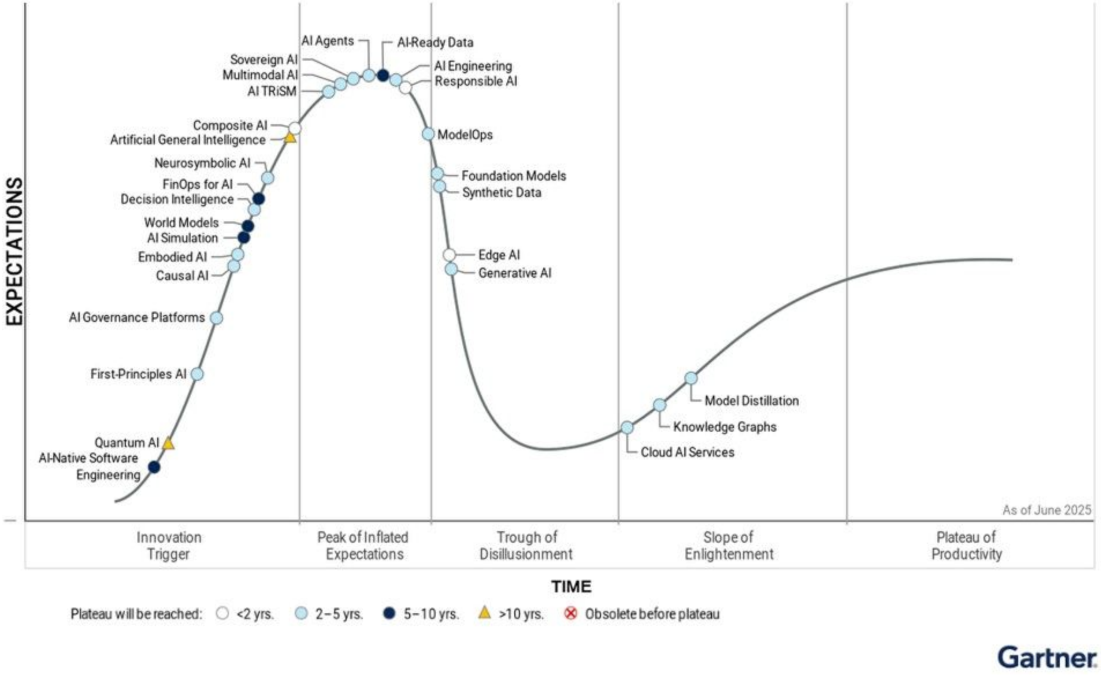

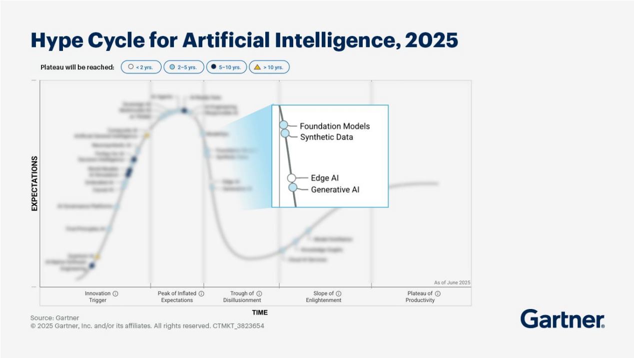

Graphically shows the maturity and adoption of technologies, helping to solve real problems and take advantage of opportunities.

Source: https://www.gartner.com/en/newsroom/press-releases/2025-08-05-gartner-hype-cycle-identifies-top-ai-innovations-in-2025

AI beyond 2025

AI beyond 2025

AI Democratization: Making AI tools and knowledge accessible to everyone, lowering barriers for innovation and participation worldwide.

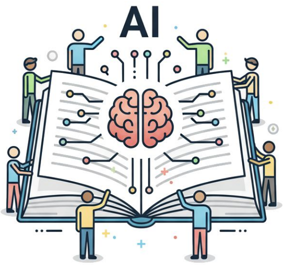

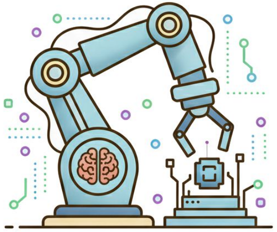

AI beyond 2025

AI Democratization: Making AI tools and knowledge accessible to everyone, lowering barriers for innovation and participation worldwide.

Ethical AI and Regulation: Promoting transparency, fairness, and accountability through responsible AI development and global policy frameworks.

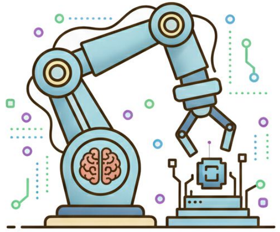

AI beyond 2025

AI Democratization: Making AI tools and knowledge accessible to everyone, lowering barriers for innovation and participation worldwide.

Ethical AI and Regulation: Promoting transparency, fairness, and accountability through responsible AI development and global policy frameworks.

AI in Healthcare: Powering precision medicine, early diagnostics, and improved patient care through data-driven insights.

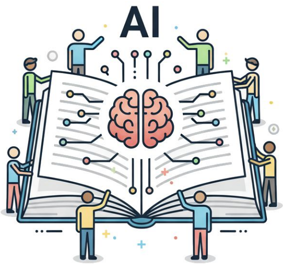

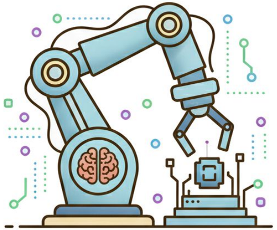

AI beyond 2025

AI Democratization: Making AI tools and knowledge accessible to everyone, lowering barriers for innovation and participation worldwide.

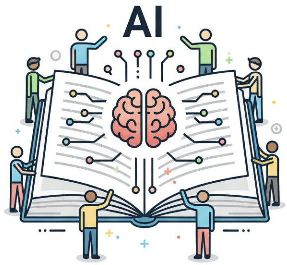

Ethical AI and Regulation: Promoting transparency, fairness, and accountability through responsible AI development and global policy frameworks.

AI in Healthcare: Powering precision medicine, early diagnostics, and improved patient care through data-driven insights.

AI and Human Collaboration: Enhancing creativity and productivity by combining human intuition with machine intelligence.

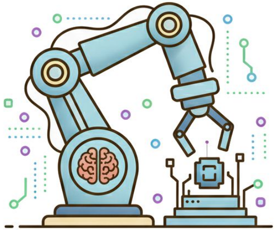

AI beyond 2025

AI Democratization: Making AI tools and knowledge accessible to everyone, lowering barriers for innovation and participation worldwide.

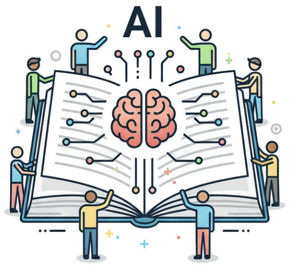

Ethical AI and Regulation: Promoting transparency, fairness, and accountability through responsible AI development and global policy frameworks.

AI in Healthcare: Powering precision medicine, early diagnostics, and improved patient care through data-driven insights.

AI and Human Collaboration: Enhancing creativity and productivity by combining human intuition with machine intelligence.

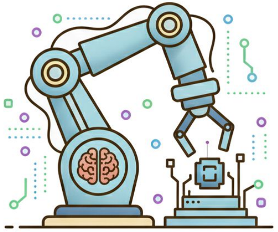

Autonomous Systems and AI-Powered Robotics: Enabling intelligent automation in transportation, manufacturing, and logistics.

AI beyond 2025

AI in Education: Delivering personalized and adaptive learning experiences for students across all levels.

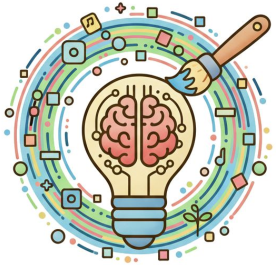

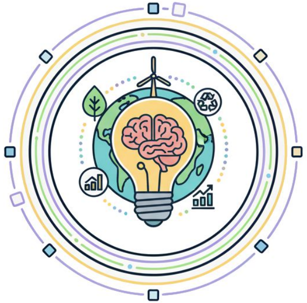

The Abdus Salam International Centre for Theoretical Physics

AI beyond 2025

AI in Education: Delivering personalized and adaptive learning experiences for students across all levels.

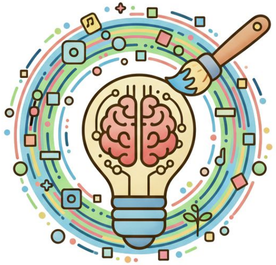

Natural Language Processing (NLP) and Conversational AI: Advancing communication between humans and machines through natural, context-aware interactions.

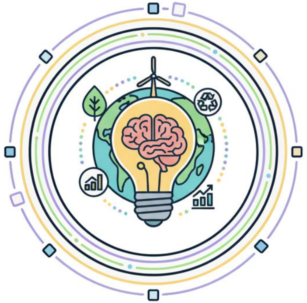

AI beyond 2025

AI in Education: Delivering personalized and adaptive learning experiences for students across all levels.

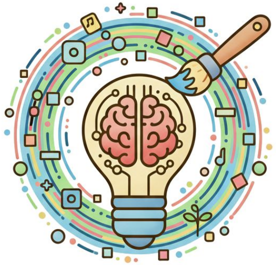

Natural Language Processing (NLP) and Conversational AI: Advancing communication between humans and machines through natural, context-aware interactions.

Quantum AI: Leveraging quantum computing to accelerate learning and solve complex problems beyond classical limits.

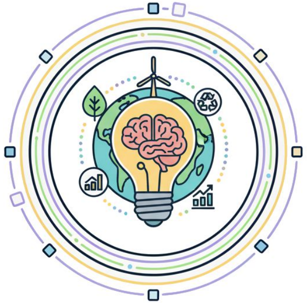

AI beyond 2025

AI in Education: Delivering personalized and adaptive learning experiences for students across all levels.

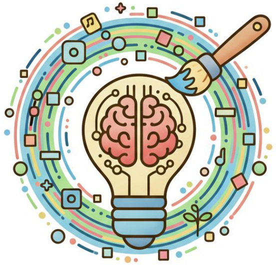

Natural Language Processing (NLP) and Conversational AI: Advancing communication between humans and machines through natural, context-aware interactions.

Quantum AI: Leveraging quantum computing to accelerate learning and solve complex problems beyond classical limits.

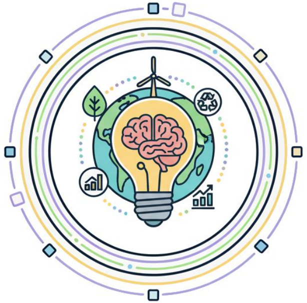

AI and Sustainability: Applying AI to address climate change, optimize energy use, and build greener technologies.

AI beyond 2025

AI in Education: Delivering personalized and adaptive learning experiences for students across all levels.

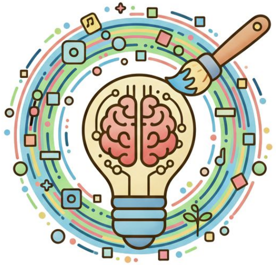

Natural Language Processing (NLP) and Conversational AI: Advancing communication between humans and machines through natural, context-aware interactions.

Quantum AI: Leveraging quantum computing to accelerate learning and solve complex problems beyond classical limits.

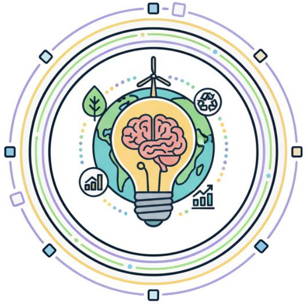

AI and Sustainability: Applying AI to address climate change, optimize energy use, and build greener technologies.

AI-Driven Creativity: Unlocking new forms of art, design, and innovation through generative models.

AI beyond 2025

5G/6G and Edge AI: Combining ultra-fast connectivity with local AI processing to enable real-time, low-latency intelligence at the edge.

AI beyond 2025

5G/6G and Edge AI: Combining ultra-fast connectivity with local AI processing to enable real-time, low-latency intelligence at the edge.

AI Infrastructure and MLOps: Building scalable pipelines for data, model training, and deployment to support enterprise AI adoption.

AI beyond 2025

5G/6G and Edge AI: Combining ultra-fast connectivity with local AI processing to enable real-time, low-latency intelligence at the edge.

AI Infrastructure and MLOps: Building scalable pipelines for data, model training, and deployment to support enterprise AI adoption.

AI and Cybersecurity: Strengthening digital defense systems and detecting threats proactively using intelligent algorithms.

AI beyond 2025

5G/6G and Edge AI: Combining ultra-fast connectivity with local AI processing to enable real-time, low-latency intelligence at the edge.

AI Infrastructure and MLOps: Building scalable pipelines for data, model training, and deployment to support enterprise AI adoption.

AI and Cybersecurity: Strengthening digital defense systems and detecting threats proactively using intelligent algorithms.

Neuromorphic Computing: Designing brain-inspired chips for efficient, adaptive AI.

The Abdus Salam International Centre for Theoretical Physics

AI beyond 2025

- *Question: What do you think about the future of AI?*

AI beyond 2025

AI Market Growth and Impact

Market Growth:

- The global AI market is currently valued at several hundred billion USD and is projected to reach between US $2–4 trillion by the early-2030s, with annual growth rates (CAGR) in the ≈ 20% to 31% range — depending on the forecast.

AI beyond 2025

AI Market Growth and Impact

Market Growth:

- The global AI market is currently valued at several hundred billion USD and is projected to reach between US $2–4 trillion by the early-2030s, with annual growth rates (CAGR) in the ≈ 20% to 31% range — depending on the forecast.

Business Applications:

- According to Gartner, Inc., by 2026 up to 40% of enterprise applications will contain task-specific AI agents (up from less than 5% today).
- They further predict that by 2029, 80% of common customer-service issues will be handled autonomously by what they call “agentic AI” (with little or no human intervention).

AI beyond 2025

AI Market Growth and Impact

Market Growth:

- The global AI market is currently valued at several hundred billion USD and is projected to reach between US $2–4 trillion by the early-2030s, with annual growth rates (CAGR) in the ≈ 20% to 31% range — depending on the forecast.

Business Applications:

- According to Gartner, Inc., by 2026 up to 40% of enterprise applications will contain task-specific AI agents (up from less than 5% today).
- They further predict that by 2029, 80% of common customer-service issues will be handled autonomously by what they call “agentic AI” (with little or no human intervention).

Impact on Businesses & Jobs:

- Business models and job roles are poised to undergo rapid and sometimes unpredictable transformations—driven by autonomous systems, embedded AI in workflows and changing expectations of productivity and capabilities.

AI-Specific Hardware Platforms

- **AI Chips & Processors**
  - **GPUs** (e.g., NVIDIA H100, AMD Instinct MI300) – Still the backbone of AI acceleration.
  - **TPUs** (Google Tensor Processing Units) – Optimized for deep learning workloads.
  - **NPUs** (Neural Processing Units) – Specialized for on-device AI in smartphones and edge devices.
  - **Analog & Optical AI Chips** – Emerging technology for ultra-low-power AI inference.

AI-Specific Hardware Platforms

- **AI Hardware Trends**
  - **Edge AI** – Custom low-power AI chips for real-time processing on IoT devices.
  - **RISC-V AI Processors** – Open-source architecture gaining traction for AI acceleration.
  - **Neuromorphic Computing** – Brain-inspired hardware and algorithms designed for parallel, adaptive, and energy-efficient processing.

AI-Specific Hardware Platforms

- *Question: What do you think about FPGA and AI?*

AI-Specific Hardware Platforms

- **FPGA+AI Trends**
  - **AI-Specific FPGA Architectures** – New FPGA models optimized for deep learning (Xilinx Versal AI Core, Intel Stratix 10 NX, and Lattice Avant AI).
  - **Quantum + FPGA Hybrid Computing** – Research into integrating FPGA with quantum acceleration.
  - **FPGA as-a-Service (FaaS)** – Cloud-based FPGA solutions for scalable AI workloads.
  - **Embedded AI & Edge Computing** – FPGAs in IoT & industrial AI, enabling real-time decision-making with ultra-low power.

Edge AI

Edge AI

"**Edge artificial intelligence** (Edge AI) refers to the deployment of AI algorithms and AI models directly on local edge devices such as sensors or Internet of Things (IoT) devices, which enables real-time data processing and analysis without constant reliance on cloud infrastructure." [IBM]

[IBM] https://www.ibm.com/think/topics/edge-ai

Edge AI

“**Edge artificial intelligence** (Edge AI) refers to the deployment of AI algorithms and AI models directly on local edge devices such as sensors or Internet of Things (IoT) devices, which enables real-time data processing and analysis without constant reliance on cloud infrastructure.” [IBM]

On the edge processing

[IBM] https://www.ibm.com/think/topics/edge-ai

Edge AI

"**Edge artificial intelligence** (Edge AI) refers to the deployment of AI algorithms and AI models directly on local edge devices such as sensors or Internet of Things (IoT) devices, which enables real-time data processing and analysis without constant reliance on cloud infrastructure." [IBM]

On the edge processing

Low latency

[IBM] https://www.ibm.com/think/topics/edge-ai

Edge AI

“**Edge artificial intelligence** (Edge AI) refers to the deployment of AI algorithms and AI models directly on local edge devices such as sensors or Internet of Things (IoT) devices, which enables real-time data processing and analysis without constant reliance on cloud infrastructure.” [IBM]

On the edge processing

Privacy and security

Low latency

[IBM] https://www.ibm.com/think/topics/edge-ai

Edge AI

"**Edge artificial intelligence** (Edge AI) refers to the deployment of AI algorithms and AI models directly on local edge devices such as sensors or Internet of Things (IoT) devices, which enables real-time data processing and analysis without constant reliance on cloud infrastructure." [IBM]

- On the edge processing
- Privacy and security
- Low latency
- Reduced Bandwidth Usage

[IBM] https://www.ibm.com/think/topics/edge-ai

Edge AI

"**Edge artificial intelligence** (Edge AI) refers to the deployment of AI algorithms and AI models directly on local edge devices such as sensors or Internet of Things (IoT) devices, which enables real-time data processing and analysis without constant reliance on cloud infrastructure." [IBM]

- On the edge processing
- Privacy and security
- Energy efficiency
- Low latency
- Reduced Bandwidth Usage

[IBM] https://www.ibm.com/think/topics/edge-ai

Edge AI

Intelligence closer to data

- **Low latency:** Real-time decisions, reducing cloud dependence.
- **Privacy & security:** Data stays on the device.
- **Lower bandwidth use:** Less data sent to the cloud.
- **Energy efficiency:** Optimized local processing.
- **Scalability & reliability:** Devices run autonomously, even offline

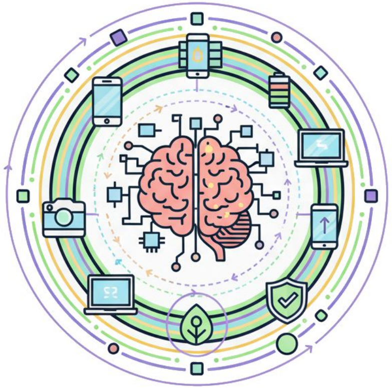

Bringing intelligence closer to the data, enabling low-latency, privacy-preserving, and energy-aware AI.

Edge AI
Challenges

Hardware & Energy	Computation & Performance	Data & Deployment
Limited power	Latency & real-time	Privacy & security
Memory/storage limits	Limited compute capacity	Network issues
	Hardware heterogeneity	

Edge AI

Challenges

- *Question:*

  Given all these challenges… how can we make AI models run efficiently on tiny, power-limited devices?

Edge AI

Challenges

- *Question:*

  What if we could combine the strengths of different processors CPUs, GPUs, and FPGAs — to overcome these limits?

Edge AI

Heterogeneous Computing: A Key Enabler for Edge AI.

- By incorporating CPUs, GPUs, and FPGAs into a single system, **heterogeneous computing** becomes possible.

Edge AI

Heterogeneous Computing: A Key Enabler for Edge AI.

- By incorporating CPUs, GPUs, and FPGAs into a single system, **heterogeneous computing** becomes possible.

- This approach maximizes the strengths of each component, efficiently distributing edge workloads to improve both performance and energy consumption.

Edge AI

Heterogeneous Computing: A Key Enabler for Edge AI.

- **CPU:** handles general-purpose tasks like system management, control logic, and lightweight AI inference.

Edge AI

Heterogeneous Computing: A Key Enabler for Edge AI.

- **CPU:** handles general-purpose tasks like system management, control logic, and lightweight AI inference.

- **GPU:** accelerates parallelizable tasks such as deep learning inference, image processing, and high-performance computing.

Edge AI

Heterogeneous Computing: A Key Enabler for Edge AI.

- **CPU:** handles general-purpose tasks like system management, control logic, and lightweight AI inference.
- **GPU:** accelerates parallelizable tasks such as deep learning inference, image processing, and high-performance computing.
- **FPGA:** provides real-time, ultra-low-latency processing for tasks like sensor fusion, encryption, and custom AI accelerators.

Edge AI based on FPGA

FPGA / SoC-based on FPGA

Edge AI based on FPGA

Low latency

FPGA / SoC-based on FPGA

Edge AI based on FPGA

Low latency

Energy Efficiency

FPGA / SoC-based on FPGA

Edge AI based on FPGA

- Low latency
- Energy Efficiency
- High parallelism

FPGA / SoC-based on FPGA

Edge AI based on FPGA

Low latency

Energy Efficiency

High parallelism

Scalability

FPGA / SoC-based on FPGA

Edge AI based on FPGA

- Low latency
- Energy Efficiency
- High parallelism
- Scalability
- Customizable AI Acceleration

FPGA / SoC-based on FPGA

Edge AI based on FPGA

- Low latency
- Energy Efficiency
- High parallelism
- Scalability
- Customizable AI Acceleration

FPGA / SoC-based on FPGA

Resource-constrained devices

Edge AI based on FPGA

Original
ML-based model

Edge AI based on FPGA

Original
ML-based model

Edge AI based on FPGA

Edge AI based on FPGA

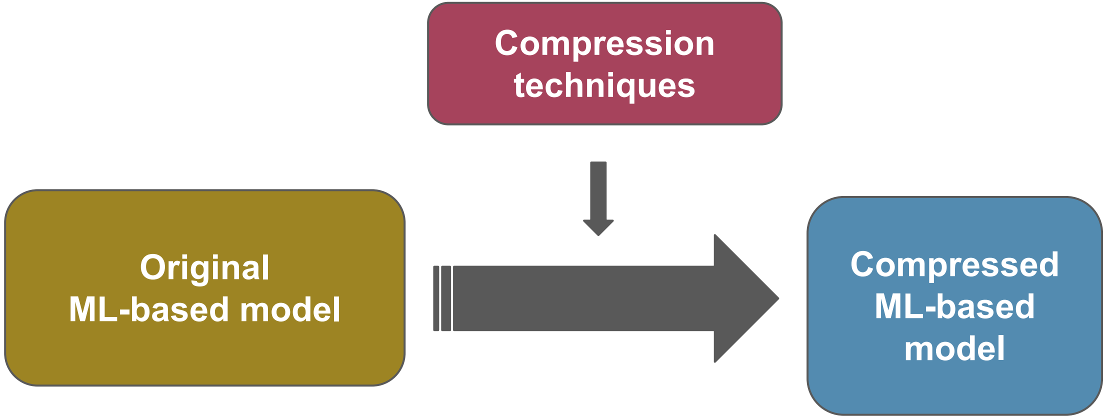

Edge AI based on FPGA

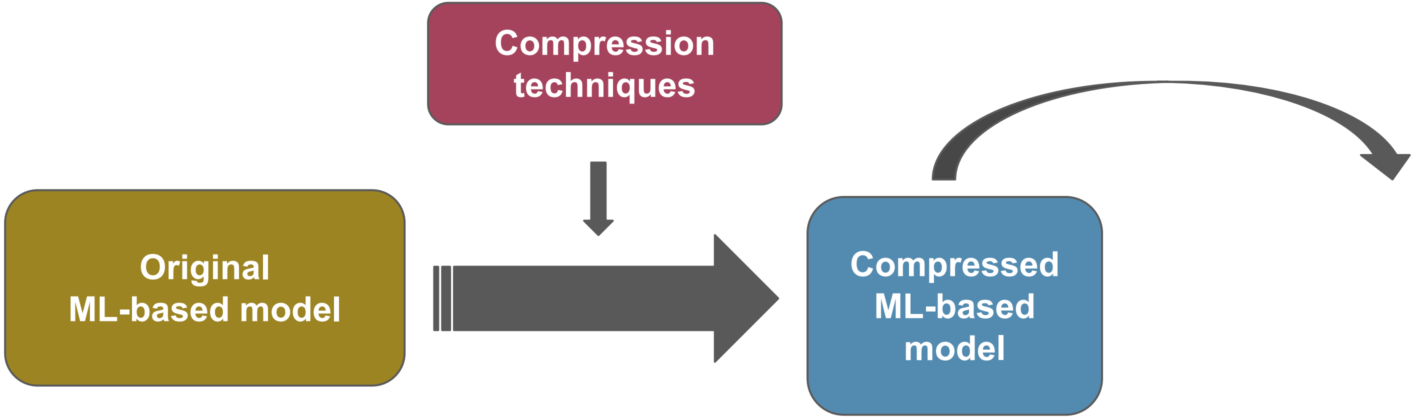

Edge AI based on FPGA

Compression techniques

Original ML-based model

Compressed ML-based model

Embedded system (SoC-based FPGA)

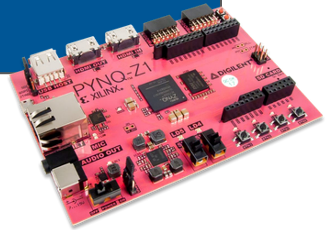

What the Literature Says

What the Literature Says

- **ML on constrained devices is now a major focus**
  - TinyML and on-device inference are growing research areas.
- **Classical state-of-the-art models are being adapted for the edge**
  - MobileNetV2, TinyBERT, and YOLO-Lite show that efficiency is now as important as accuracy.
- **Compression through pruning & quantization dominates optimization**
  - These remain the most used techniques to fit large models into edge hardware.
- **Off-chip memory access is a key bottleneck**
  - Memory transfers can consume more energy than computation.
- **Workflows are still fragmented across tools**
  - Full end-to-end edge pipelines are rare; integration remains an open challenge.

What the Literature Says

Memory footprint and latency

What the Literature Says

Memory footprint and latency

Ensemble of compression techniques

What the Literature Says

Memory footprint and latency

Ensemble of compression techniques

On-chip memory deployment

What the Literature Says

Memory footprint and latency

Ensemble of compression techniques

On-chip memory deployment

End-to-end workflow

What the Literature Says

Memory footprint and latency

Ensemble of compression techniques

On-chip memory deployment

End-to-end workflow

Productivity

Integrating Every Layer of Edge AI Design

Integrating Every Layer of Edge AI Design

Optimization is holistic — every phase, from model to hardware, impacts overall performance at the edge.

Software development

Model training

Model optimization & compression

Integrating Every Layer of Edge AI Design

Optimization is holistic — every phase, from model to hardware, impacts overall performance at the edge.

Software development	Hardware platform selection
Model training	Edge devices
Model optimization & compression	Cloud/GPU/CPU
	Other computing platforms

Integrating Every Layer of Edge AI Design

Optimization is holistic — every phase, from model to hardware, impacts overall performance at the edge.

Software development	Hardware platform selection	Firmware creation
Model training	Edge devices	C/C++ application
Model optimization & compression	Cloud/GPU/CPU	Python application
	Other computing platforms	HDL design

Integrating Every Layer of Edge AI Design

Optimization is holistic — every phase, from model to hardware, impacts overall performance at the edge.

Software development	Hardware platform selection	Firmware creation	PCB design
Model training	Edge devices	C/C++ application	Key components
Model optimization & compression	Cloud/GPU/CPU	Python application	Signal and power integrity
	Other computing platforms	HDL design	Methodological design
			EMC and EMI

Case Study for FPGA-based Edge AI: MNIST Binary Classification.

ML and model compression techniques for SoC/FPGA

Demo: MNIST-based binary classification

MNIST-based
binary classification

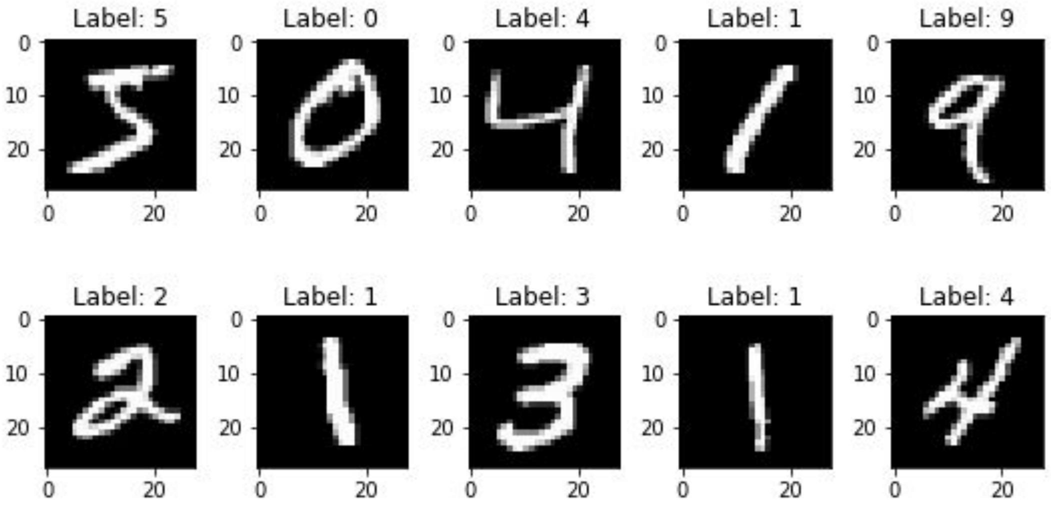

ML and model compression techniques for SoC/FPGA

Demo: MNIST-based binary classification

MNIST-based binary classification

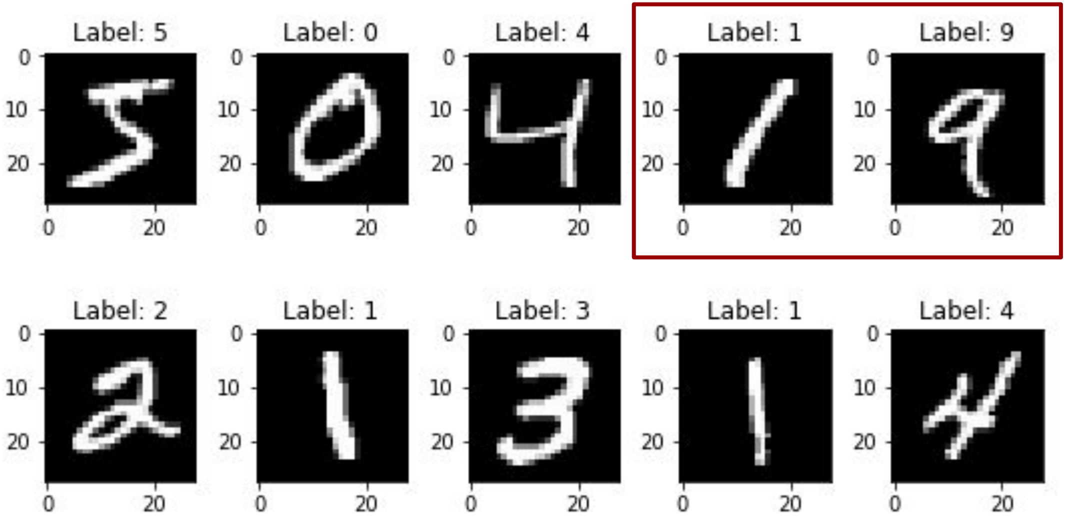

ML and model compression techniques for SoC/FPGA

Demo: MNIST-based binary classification

MNIST-based binary classification

- Quantization-Aware Pruning
  - 8-bits fixed point precision
  - 20% target sparsity
  - QKeras for model definition

```
_________________________________________________________________
Layer (type)                 Output Shape              Param #
=================================================================
fc1_input (QDense)           (None, 5)                 3925
_________________________________________________________________
relu_input (QActivation)     (None, 5)                 0
_________________________________________________________________
fc1 (QDense)                 (None, 7)                 42
_________________________________________________________________
relu1 (QActivation)          (None, 7)                 0
_________________________________________________________________
fc2 (QDense)                 (None, 10)                80
_________________________________________________________________
relu2 (QActivation)          (None, 10)                0
_________________________________________________________________
output (QDense)              (None, 2)                 22
_________________________________________________________________
sigmoid (Activation)         (None, 2)                 0
=================================================================
Total params: 4,069
Trainable params: 4,069
Non-trainable params: 0
_________________________________________________________________
```

ML and model compression techniques for SoC/FPGA

Demo: MNIST-based binary classification

MNIST-based binary classification

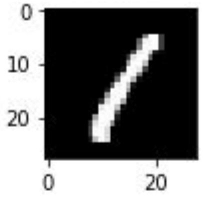

Class 0

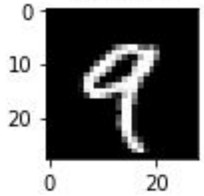

Class 1

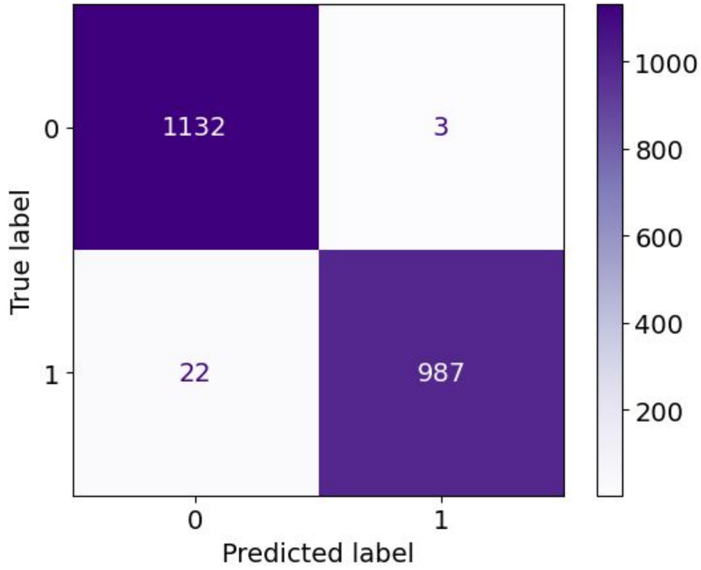

ML and model compression techniques for SoC/FPGA

Demo: MNIST-based binary classification

MNIST-based
binary classification

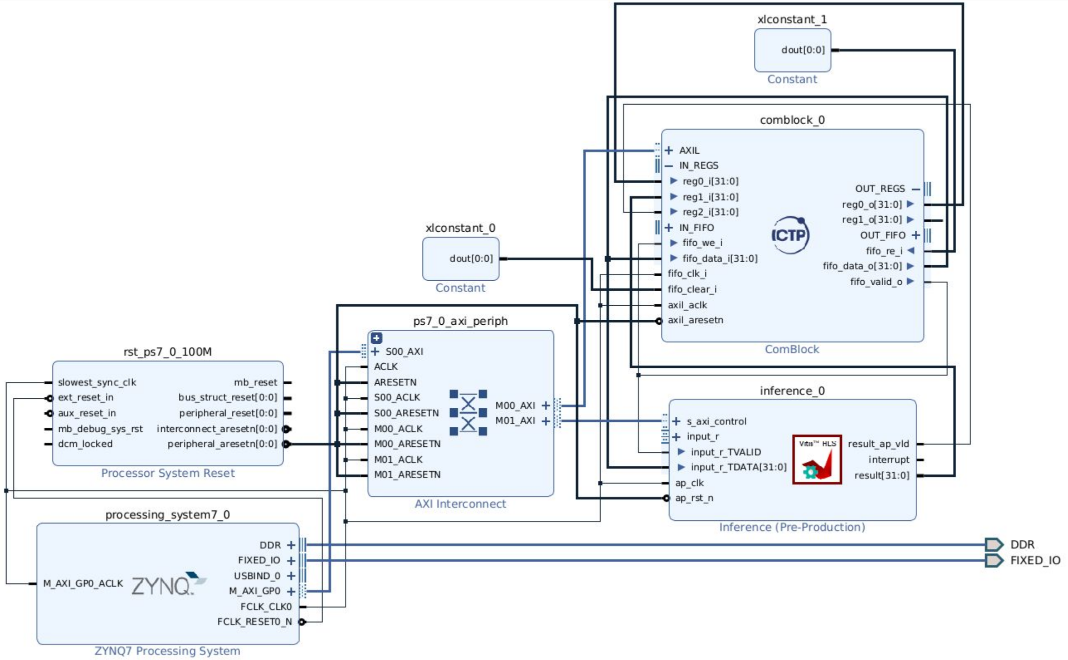

ML and model compression techniques for SoC/FPGA

Demo: MNIST-based binary classification

MNIST-based
binary classification

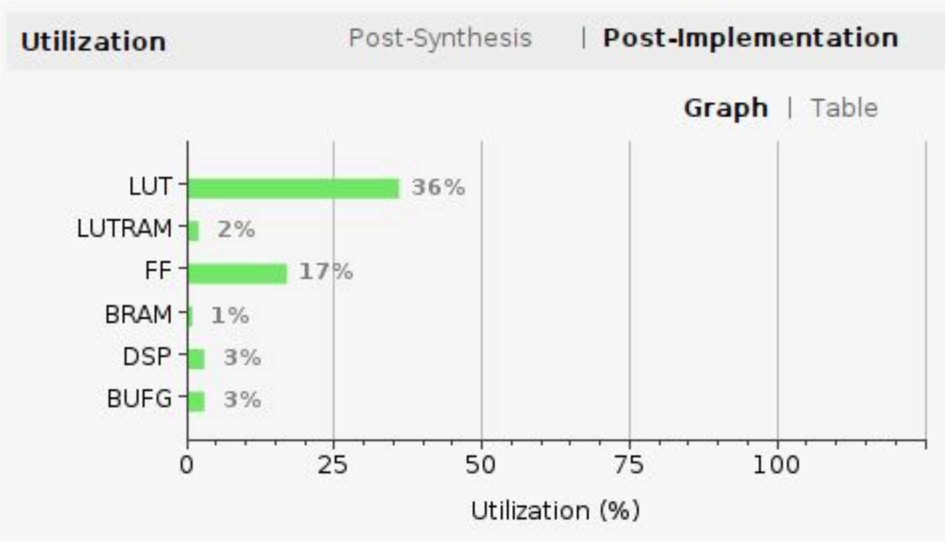

ML and model compression techniques for SoC/FPGA

Demo: MNIST-based binary classification

MNIST-based
binary classification

IP core based on ML integrated with PYNQ framework

In [ ]:
```
from pynq import Overlay
from pynq import MMIO
import comblock as cbc

import numpy as np
import matplotlib.pyplot as plt
```

Load Overlay

In [ ]:
```
# Load the overlay (bitstream) onto the FPGA.

ol = Overlay("design_1_wrapper.xsa")
```

The information from the **comblock_0** block is read to verify everything that is obtained. Since the object is mapped to AXI Lite, it is noted that the AXI Full address is omitted.

In [ ]:
```
## Overlay information

ol.ip_dict
```

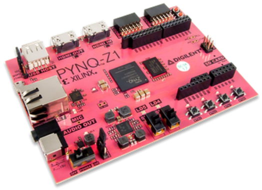

PYNQ™

ML and model compression techniques for SoC/FPGA

Demo: MNIST-based binary classification

MNIST-based
binary classification

ComBlock information

ComBlock for PYNQ: https://github.com/Mballina42/PynQ_ComBlock

For convenience, the *comblock.py* Python script is established which contains useful constants for interacting with the ComBlock.

In [ ]:
```
ol.ip_dict['comblock_0']
```

In [ ]:
```
# The object is created based on the comblock_0 IP
cb = ol.comblock_0
```

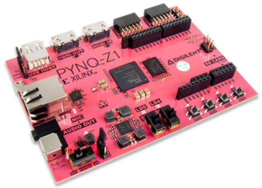

ML and model compression techniques for SoC/FPGA

Demo: MNIST-based binary classification

MNIST-based binary classification

Data preparation

```
In [ ]: signal_1 = [
            0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
          ]
```

```
In [ ]: imageArray = np.array(signal_1)

        image_2d = imageArray.reshape((28, 28))

        # Display as an image
        plt.imshow(image_2d, cmap='gray', interpolation='nearest')
        plt.colorbar()  # Optional: Show color scale
        plt.show()
```

```
In [ ]: signal_2 = [
            0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
        ]
```

```
In [ ]: imageArray = np.array(signal_2)

        image_2d = imageArray.reshape((28, 28))

        # Display as an image
        plt.imshow(image_2d, cmap='gray', interpolation='nearest')
        plt.colorbar()  # Optional: Show color scale
        plt.show()
```

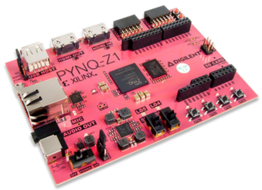

PYNQ™

ML and model compression techniques for SoC/FPGA

Demo: MNIST-based binary classification

MNIST-based
binary classification

Interacting with ComBlock

Signal 1

Write FIFO - Send image to the FPGA

In [ ]:
```python
cb.write(cbc.OREG1, 1)

# Send data to the ComBlock's FIFO
data_size = 28*28
for i in range(data_size):
        cb.write(cbc.OFIFO_VALUE, signal_1[i])
```

Read registers - Read inference result from the FPGA

In [ ]:
```python
# Read IREG1 to obtain the result of the inferece process
cb.read(cbc.IREG1)
```

Signal 2

Write FIFO - Send image to the FPGA

In [ ]:
```python
cb.write(cbc.OREG1, 1)

# Send data to the ComBlock's FIFO
data_size = 28*28
for i in range(data_size):
        cb.write(cbc.OFIFO_VALUE, signal_2[i])
```

Read registers - Read inference result from the FPGA

In [ ]:
```python
cb.read(cbc.IREG1)
```

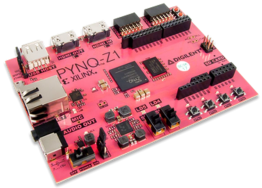

PYNQ™

Joint ICTP-IAEA School on Detector Signal Processing and Machine Learning for Scientific Instrumentation and Reconfigurable Computing

Introduction to Machine Learning and Edge AI

smr4110 | Trieste, Italy - 2025

Romina Soledad Molina, Ph.D.

IAEA
International Atomic Energy Agency
Atoms for Peace and Development

The Abdus Salam International Centre for Theoretical Physics