Benchmarking Neuromorphic and Conventional AI for Energy-Efficient Inference

Supervisor: *Sifat Rezwan, Juan Cabrera*

Neuromorphic computing, particularly through spiking neural networks (SNNs), promises significant energy savings compared to conventional deep learning. However, systematic benchmarks of SNNs versus traditional neural networks on real hardware accelerators are scarce. With access to heterogeneous hardware (Axelera, Hailo, Jetson Nano, Asus, GPUs), this project investigates whether neuromorphic models provide measurable benefits in accuracy, energy, latency, and robustness, contributing to the evidence base for Green AI.

What you will do

- Implement spiking neural networks using frameworks (e.g., Nengo, snnTorch, Brian2, spikingjelly).
- Train and benchmark SNNs on standard classification/detection datasets.
- Compare SNNs against equivalent ANN baselines (e.g., LeNet, ResNet variants, YOLO-lite).
- Deploy models across hardware (Axelera, Hailo, Jetson Nano, Asus, GPU).
- Evaluate and compare:
    - Accuracy: Top-N classification accuracy, mAP for detection.
    - Latency: average inference time per sample.
    - Throughput: effective frames per second (FPS).
    - Energy: Joules per inference.
    - Energy–accuracy Pareto front analysis, if possible.
    - Comparative analysis: ANN vs. SNN on each hardware platform.
- Perform an engineering analysis of neuromorphic vs. conventional approaches in real-world edge AI deployment.

Why these matters

Spiking neural networks are often touted as the future of sustainable AI, but their advantages need rigorous evaluation. This thesis clarifies whether neuromorphic approaches outperform conventional models in practical edge AI scenarios, providing evidence for or against their role in Green AI.

Required skills

- Python (must), ML frameworks (PyTorch/TensorFlow).
- Familiarity with computer vision architectures.
- Familiarity with spiking neural network libraries (optional).
- Basic knowledge of Linux and hardware benchmarking.