

Comparing Energy Consumption of Local and Cloud ML Processing on Embedded Systems

Emil Njor

December 21, 2022

Supervisor: Xenofon Fafoutis <xefa@dtu.dk>

Co-Supervisor: Emil Njor <emjn@dtu.dk>

Background: tinyML is a promising research area concerned with running machine learning models on ultra-low power devices, typically in the range of milliwatts or below. The state of the art of tinyML involves training machine learning models on larger computers (e.g., a laptop computer), and subsequently applying optimizations and deploying the models on ultra-low power devices.

It is often claimed that tinyML will improve the energy efficiency of embedded devices by running machine learning inference locally. The reasoning behind this claim is that, when processing machine learning locally, the device does not need to send and receive data from a remote server. As the networking module on a device is typically the most power-hungry component, this should result in a significant reduction in energy consumption.

However, according to my knowledge, this claim has not yet been tested in an empirical study.

Project Description: The goal of this project is to compare the energy consumption of local and cloud-based machine learning processing on embedded systems. The project will involve implementing tinyML systems on an embedded device, and equivalent cloud based machine learning systems where the embedded device acts as a data capture device. A methodology for measuring the energy consumption of each system should be developed, and the results should be compared.

Recommended Background Knowledge: Embedded Systems, Cloud Applications, Machine Learning, C++ & Python.