netidee STIPENDIEN

Dynamic Power Management for Edge AI: A Sustainable Self-Adaptive Approach

Interim Report | Call 19 | Scholarship ID 7383

License: CC BY

Contents

1 Introduction

This thesis investigates how a reinforcement learning (RL) agent can dynamically select object detection modes on a Raspberry Pi to balance accuracy and energy use. The project milestones included literature review, proposal, and framework selection by November 2024; RL prototype development by December 2024; testbed automation by January 2025; parameter studies through May 2025; RL training and experiments in summer 2025; and results evaluation, thesis writing, and defence planned for autumn and late 2025. So far, key objectives have been achieved, and new insights have helped refine the RL design. This report summarizes the work on the RL prototype and measurement pipeline, reviews parameter studies and action space updates, compares progress to the original plan, and outlines the updated timeline.

---

2 2. Status

2.1 2.1 Milestone 1 – Initial RL Agent Setup and Framework Selection

The first milestone concentrated on establishing the theoretical and practical foundations of the RL agent. A critical early step involved defining the state space, action space, and reward function to be used in the RL model. These components needed to be carefully designed to accurately reflect the system's operational parameters and goals. After reviewing multiple Python-based RL frameworks, Tensorforce[1] was selected due to its ability handling multivariate action spaces and compatibility with the project's computational environment. This choice enabled the implementation of a simple RL agent capable of simultaneously adjusting several parameters, such as frame rate, model variant, and image resolution.

Initial development involved constructing dummy functions to simulate energy consumption and energy production by the solar panel, which allowed preliminary testing without relying on real hardware measurements. These simulations enabled the creation of a first power consumption model, parameterized by frame rate, model variant, and

[1] "Tensorforce: A TensorFlow Library for Applied Reinforcement Learning — Tensorforce 0.6.5 Documentation."

resolution. Alongside, a preliminary rudimentary reward function to balance the competing objectives of energy efficiency and detection confidence.

Although the prototype RL agent operated within these defined spaces and executed actions accordingly, it was unable to learn meaningful policies. The failure was attributed to the oversimplification inherent in the dummy functions, which did not capture the complex interactions affecting power consumption on the Raspberry Pi. This outcome underscored the necessity for empirical data and accurate power profiling to guide further development. Profiling the power consumption wasn’t included in the original plan but soon proved necessary. As a result, experiments were set up to measure the Raspberry Pi’s energy use across different configurations.

2.2 2.2 Milestone 2 – Experimental Setup and Automation Pipeline

Milestone two involved designing and setting up an automated experimental system to measure the Raspberry Pi’s energy consumption under different workloads and settings. This was essential for gathering accurate data to train the RL agent realistically. The setup was built in the laboratory of TU Wien’s HPC Research Group[2] and included a Raspberry Pi, a digital power meter for precise energy monitoring, and a compatible camera module.

A workload generation script was created and run on the Raspberry Pi to perform object detection using selected models (different variants of YOLO11[3]), frame rates, and image resolutions. Samples from the COCO2017[4] dataset were scaled to three different resolutions to ensure consistent and repeatable tests.

Alongside, automation scripts were developed to remotely start the workload, collect power meter readings, and save data, including power consumption and prediction confidence, into CSV files. This automation allowed for systematic and reliable experiments.

Setting up the experiment came with challenges, such as gaining remote access to the Raspberry Pi, configuring the programming environment, and ensuring stable communication between the controlling laptop and the device. It was also found that the workload needed about 20 seconds to warm up before power readings stabilized; measurements taken without this warm-up were unreliable, so the protocol was adjusted accordingly.

[2] HPC Research Group
[3] Khanam and Hussain, “YOLOv11.”
[4] Lin et al., “Microsoft COCO.”

Despite these challenges, the automated measurement system was successfully completed. Although it took more time than initially expected, the approach stayed aligned with project goals and produced valuable data.

2.3 Milestone 3 – Parameter Study and Action Space Refinement

Between March and April 2025, multiple experiments were conducted to collect data on the Raspberry Pi's power consumption under various configurations. The analysis showed that the choice of model variant consistently had the most noticeable impact on both power consumption and detection accuracy. However, the influence of frame rate and image resolution was highly context-dependent. In some cases, configurations exhibiting significant differences may not be practical or beneficial as part of the RL agent's action space. Should this be the case, the project plans to shift focus toward alternative action dimensions.

A simple example can be shown using the figures below. The resolution parameter influences the power consumption of the smaller model (yolo11n) more than it influences the power consumption of the larger model (yolo11s). This is due to the fact that the larger model reaches maximum throughput at lower resolutions than the smaller model does. However, when we consider the second figure showing the distribution of average detection confidence per frame, it reveals that even though lower resolutions can lower the energy consumption they might not be suitable for realistic scenarios, since they significantly lower detection confidence.

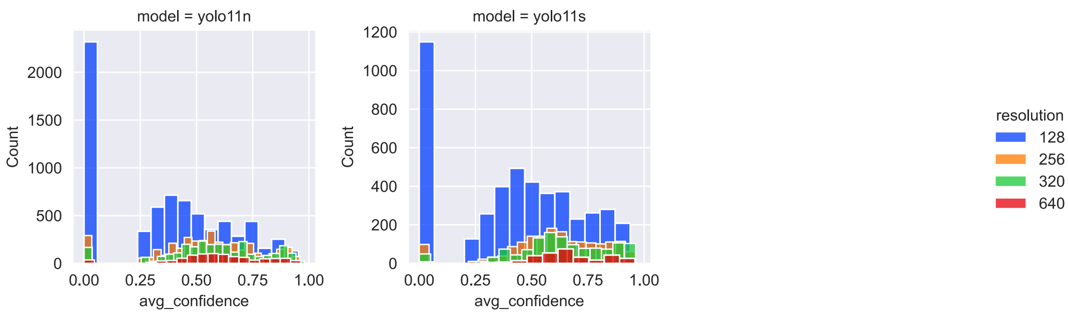

This realization challenged the initial assumptions and prompted a reassessment of the RL agent's action space. While focusing solely on the model variant could overlap with existing research, additional variables are now being considered. One key candidate is toggling data transmission over the network, which is commonly used in real-world applications for remote monitoring or cloud processing and has a well-documented impact[5] on energy consumption. If time allows, activating or deactivating the network module itself may also be incorporated. These additions introduce important new factors into the decision-making process and emphasize the project's integration with solar panel and battery management systems.

Detailed experiments on the power consumption effects of network activity will not be conducted within this thesis. Instead, well-established models from the literature will be used to estimate these impacts. This approach balances practical limitations while maintaining focus on the core RL development and evaluation.

3 Summary and Plan Update

The original project timeline planned for the completion of literature research, proposal submission, and selection of frameworks and datasets by November 2024. Energy consumption simulation on the Raspberry Pi was scheduled for January 2025, with the training of the learning engine and prototype design set for March to May. The intermediate report was due in May 2025. The following months were designated for programming the prototype, developing the workload generation module, conducting experiments, analyzing results, writing the final thesis, and defending it after November 2025.

Most of the early stages proceeded as planned. However, the initial RL environment developed in December 2024, which relied on simplified assumptions, did not produce meaningful learning results. This led to a shift toward empirical experimentation to accurately measure the device's power consumption, requiring the creation of an automated experimental setup. This setup phase extended until February 2025, followed by ongoing experiments through April. By May, thorough data analysis revealed that the choice of model variant was the primary factor affecting power consumption. This finding led to a revision of the RL agent's action space to include additional variables, ensuring the project's novelty and relevance.

Despite this necessary adjustment, the overall project timeline remains feasible. Upcoming tasks, such as training an advanced RL agent with baseline comparisons, conducting further

[5] Kaup, Gottschling, and Hausheer, "PowerPi."

experiments, evaluating results, and writing the thesis, are well aligned with the original schedule. These changes represent a strategic and important evolution of the research, enhancing its foundation and expected impact.

4 References

Kaup, Fabian, Philip Gottschling, and David Hausheer. “PowerPi: Measuring and Modeling the Power Consumption of the Raspberry Pi.” In *39th Annual IEEE Conference on Local Computer Networks*, 236–43, 2014. https://doi.org/10.1109/LCN.2014.6925777.

Khanam, Rahima, and Muhammad Hussain. “YOLOv11: An Overview of the Key Architectural Enhancements.” arXiv, October 23, 2024. https://doi.org/10.48550/arXiv.2410.17725.

Lin, Tsung-Yi, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. “Microsoft COCO: Common Objects in Context.” arXiv, February 21, 2015. https://doi.org/10.48550/arXiv.1405.0312.

“Tensorforce: A TensorFlow Library for Applied Reinforcement Learning — Tensorforce 0.6.5 Documentation.” Accessed May 26, 2025. https://tensorforce.readthedocs.io/en/latest/.