Edge Intelligence Resource Consumption by UAV-based IR Object Detection

Andrii Polukhin[1], Yuri Gordienko[1], Mairo Leier[2], Gert Jervan[2]
Oleksandr Rokovyi[1], Oleg Alienin[1], and Sergii Stirenko[1]

[1] National Technical University of Ukraine "Igor Sikorsky Kyiv Polytechnic Institute", Kyiv, Ukraine
[2] Tallinn University of Technology, Tallinn, Estonia

Abstract

The investigation of the feasibility of using the YOLO (You Only Look Once) architecture for object detection in infrared images from unmanned aerial vehicles (UAVs) on low-power devices, specifically the Raspberry Pi, and Orange Pi, is conducted. The study measures the consumption of computing resources for each device, such as inference time (ms), peak power consumption (W), memory consumption (MB), inference energy (J), memory consumption (MB), and storage consumption (MB). It also investigates the correlation between number of model parameters and resource consumption of the different YOLO model sizes. Finally, the study draws conclusions about the expediency and realism of using YOLO on low-power devices for Edge Intelligence and proposes methods of speeding up work. The results show that YOLO can be used effectively on low-power devices with some optimizations to increase performance and energy efficiency.

Introduction

In our study, we address the challenge of deploying convolutional neural networks (CNNs) like YOLO on low-power devices, including Raspberry Pi and Orange Pi, which have limited computational resources. We investigate the performance of YOLO in the context of object detection, especially in the domain of UAVs. Our findings offer valuable insights for both academic researchers and industry professionals working on computer vision applications for low-power platforms, shedding light on the feasibility and efficiency of deploying advanced CNN models in resource-constrained environments.

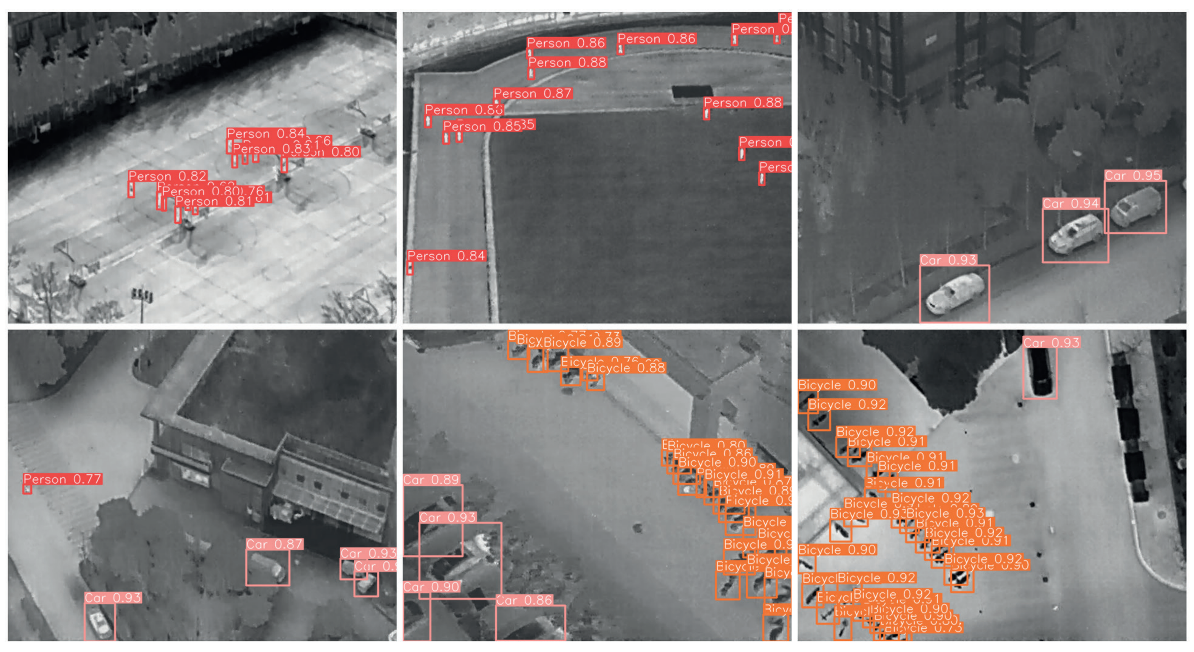

Figure 4: Sample predictions of the YOLO v5 (x).

Methodology

We aim to measure the performance of low-power devices in executing the YOLO architecture for object detection in infrared images from UAVs. We trained the model on the HIT-UAV dataset, which is comprehensively described in our previous work. To achieve this, we use a set of metrics related to the consumption of computing resources, such as inference time (s), peak power consumption (W), inference energy (J), memory consumption (MB), and storage consumption (MB).

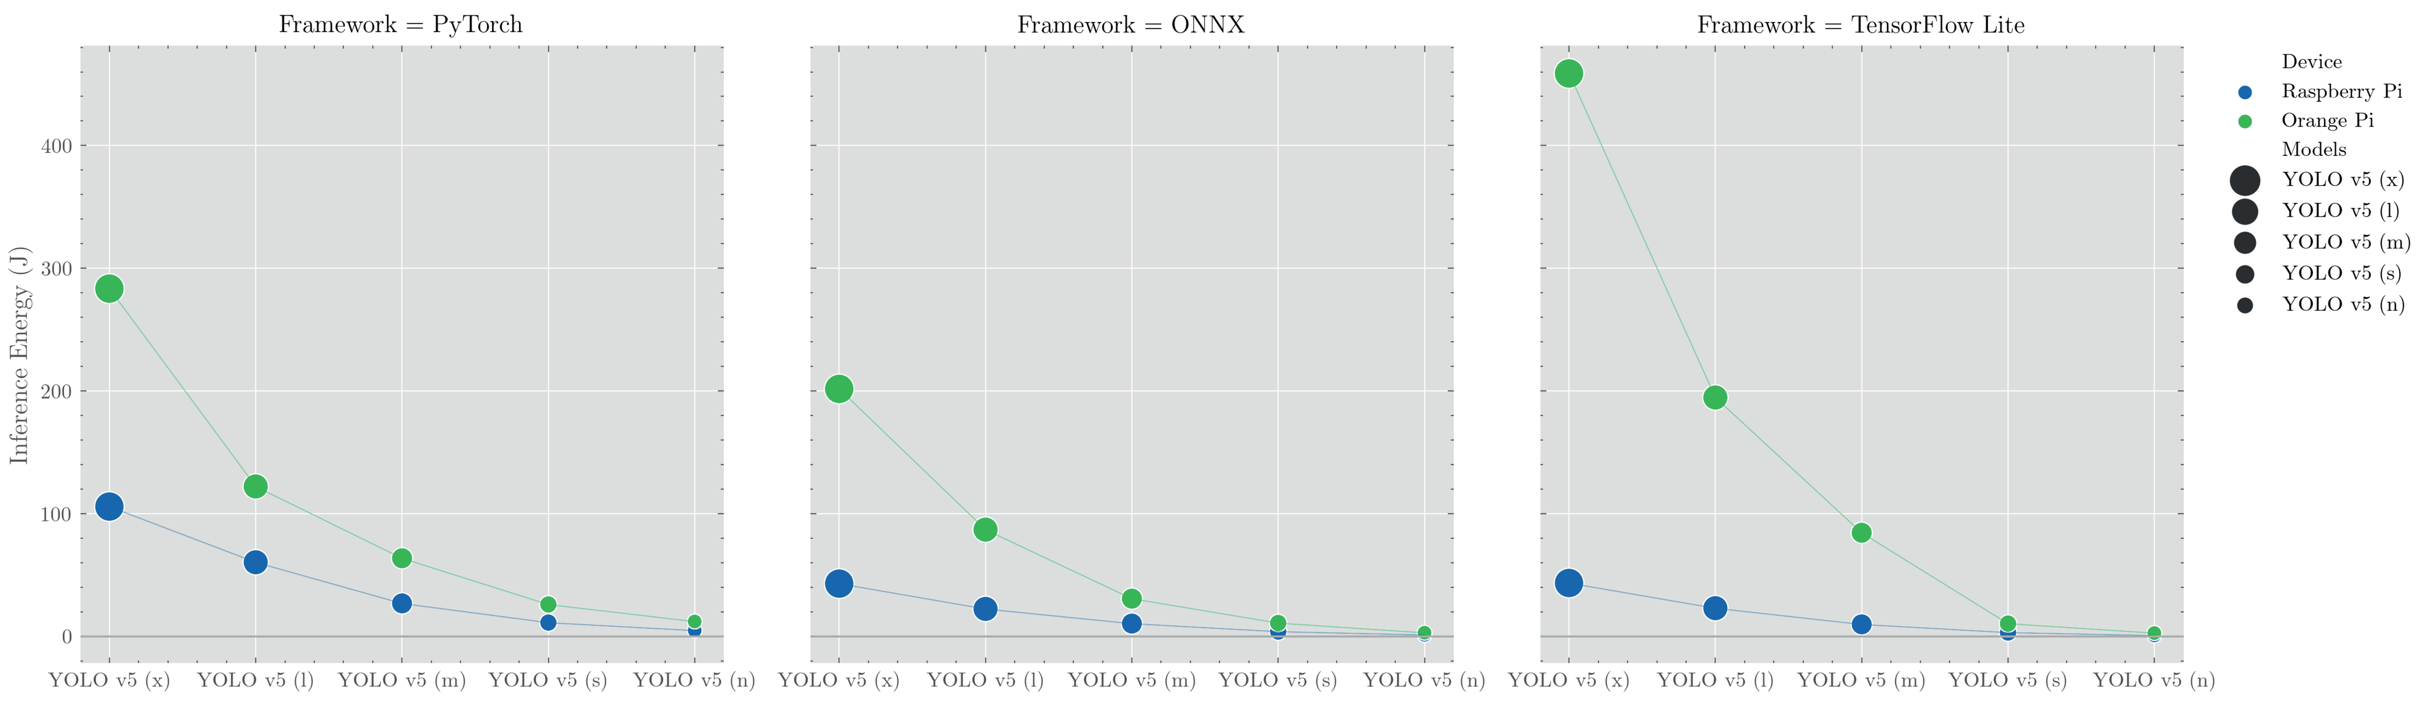

Figure 1: Comparison of Inference Energy (J) vs different model sizes using fp32 data type on Orange Pi and Raspberry Pi.

Results

We present the results and discussions of our experiments with YOLO v5 models on low-power devices, namely Raspberry Pi and Orange Pi, using various machine learning frameworks and data types. We highlight the resilience of these models when applied to infrared (IR) images, especially on coarser data. Notably, Raspberry Pi outperforms Orange Pi in terms of inference time and memory consumption, making it a superior choice for object detection tasks with UAV-based IR imaging. Additionally, we analyze the impact of model size, data types, and machine learning frameworks on inference performance and power consumption, offering insights into their trade-offs and implications for practical deployments.

Table 1: Inference Time (s) of YOLO v5 model sizes with float32 data type and different ML frameworks and devices.

		YOLO v5 Model				
	Framework	(n)	(s)	(m)	(l)	(x)
Orange Pi	PyTorch	6.8	13.0	**29.0**	**55.6**	**123.2**
	ONNX	**3.4**	**11.0**	30.8	79.1	168.1
	TF Lite	4.0	13.1	76.8	177.0	327.7
Raspberry Pi	PyTorch	1.9	3.8	7.9	16.8	27.9
	ONNX	**0.7**	**1.8**	**4.8**	**10.2**	**18.7**
	TF Lite	**0.7**	2.2	6.2	13.6	25.6

Our study also underscores the significant influence of the machine learning framework on both inference time and power consumption. ONNX shines in memory efficiency, TensorFlow Lite emerges as the energy-efficient choice for smaller models, and PyTorch delivers consistent performance, particularly when memory and power consumption need a balanced approach.

Conclusion

We showcase the practicality of deploying YOLO architecture for object detection in infrared images captured by UAVs on resource-constrained Edge Intelligence devices. Notably, our findings reveal the Raspberry Pi as a more energy-efficient choice compared to the Orange Pi. We emphasize the significant impact of framework selection, model size, and precision levels on power consumption and inference performance, underscoring the need for careful optimization to align with specific use case requirements.

Table 2: Storage and Memory Consumption (MB) of YOLO v5 model sizes with float32 data type and different ML frameworks.

		YOLO v5 Model				
	Framework	(n)	(s)	(m)	(l)	(x)
Storage Consumption	PyTorch	7.2	**26.9**	80.4	177.0	330.0
	ONNX	7.1	27.2	80.1	**176.0**	**329.0**
	TF Lite	**6.8**	**26.9**	**79.8**	**176.0**	**329.0**
Initialization Memory Consumption	PyTorch	**10.8**	**30.9**	**84.3**	**181.5**	**335.0**
	ONNX	27.2	81.9	205.7	431.0	711.1
	TF Lite	17.1	57.1	163.2	365.5	664.5
Inference Memory Consumption (MB)	PyTorch	35.3	72.9	91.2	101.3	116.1
	ONNX	**25.3**	**47.0**	**58.3**	**87.5**	**93.8**
	TF Lite	75.9	106.4	140.2	183.2	209.1

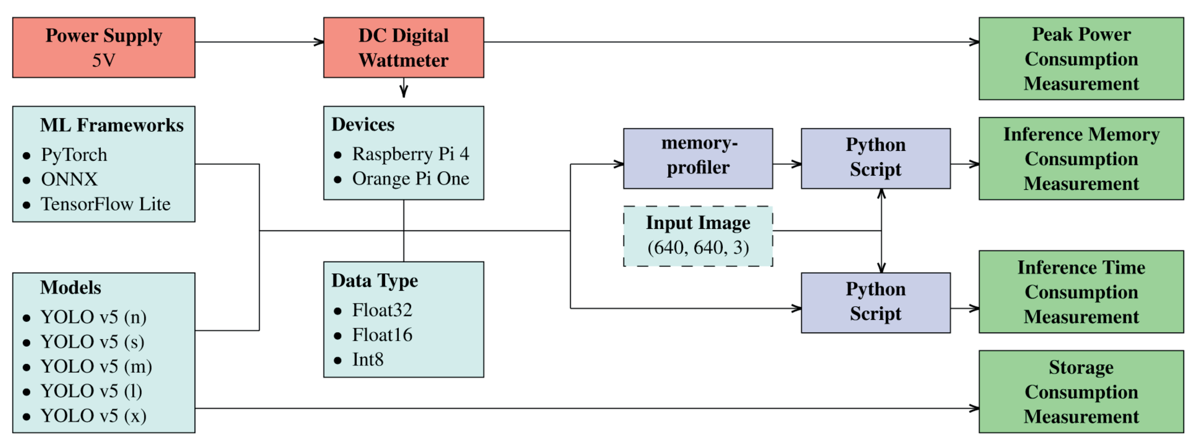

Figure 3: Resource Consumption Measurement Process Visualization.

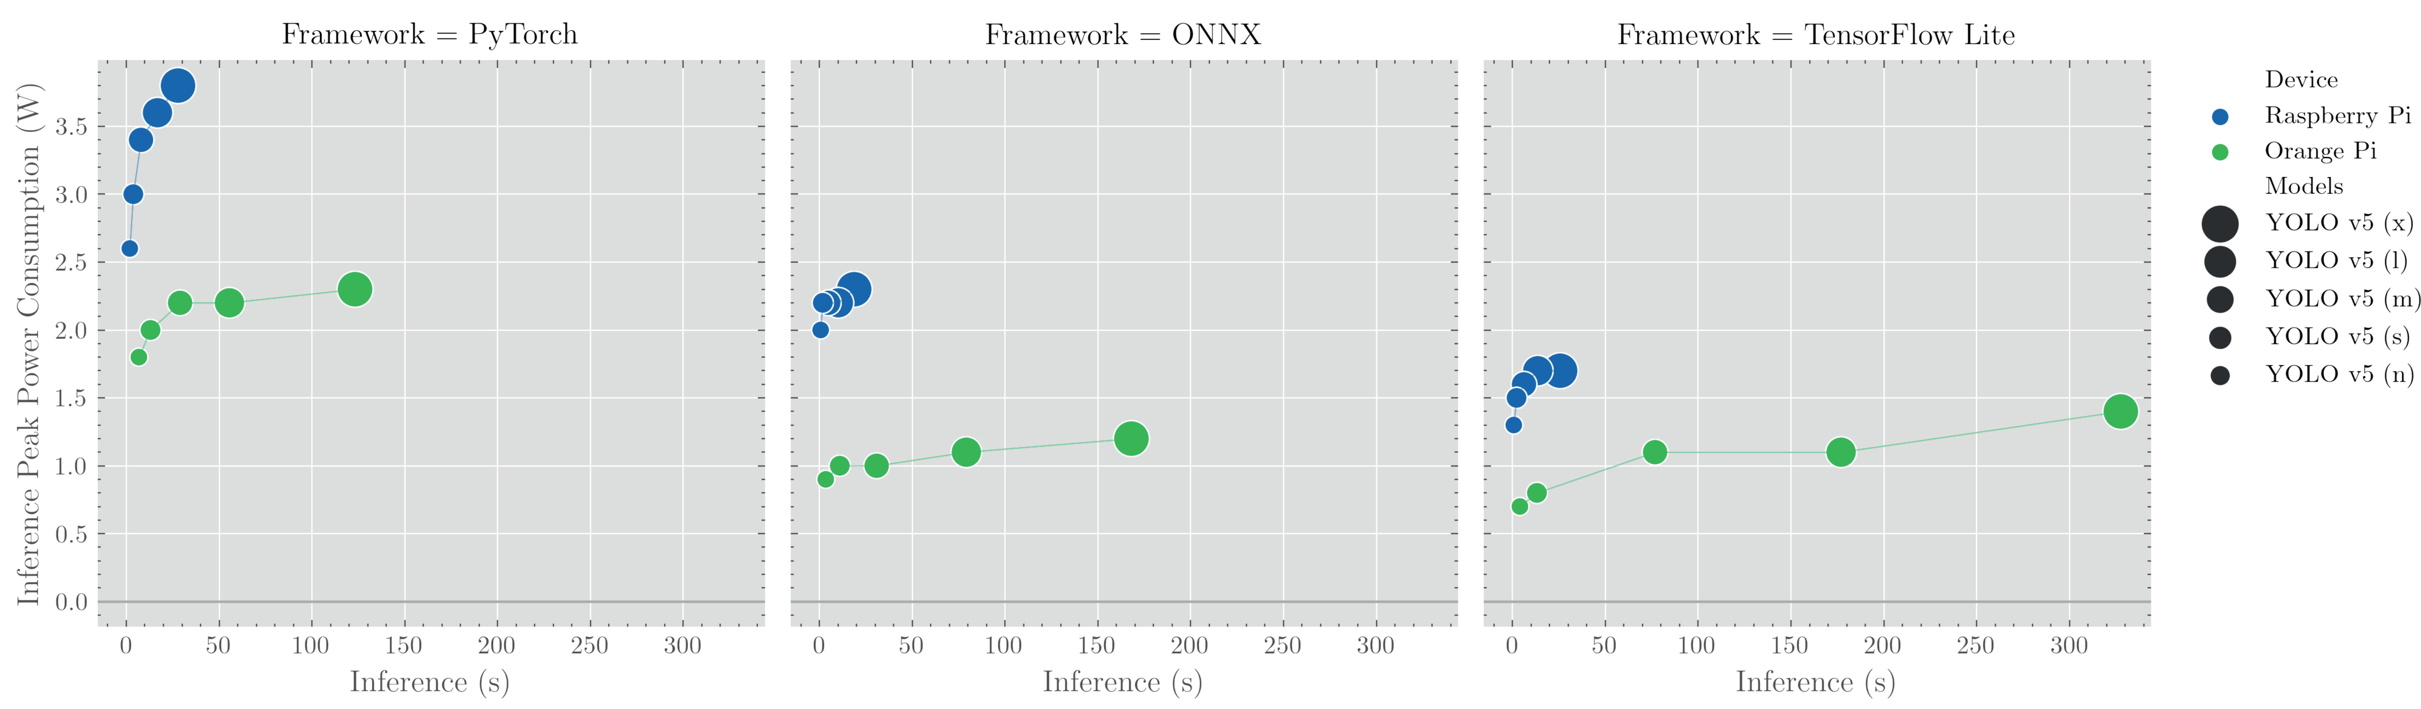

Figure 2: Comparison of Inference Peak Power Consumption (W) vs different Machine Learning Frameworks on Orange Pi and Raspberry Pi.

NATO OTAN
This publication is supported by:
The NATO Science for Peace and Security Programme (grant G6032)