



Detecting Soccer Balls with Reduced Neural Networks

A Comparison of Multiple Architectures Under Constrained Hardware Scenarios

Douglas De Rizzo Meneghetti¹ · Thiago Pedro Donadon Homem² · Jonas Henrique Renolfi de Oliveira¹ · Isaac Jesus da Silva¹ · Danilo Hernani Perico¹ · Reinaldo Augusto da Costa Bianchi¹

Received: 14 October 2020 / Accepted: 2 February 2021 / Published online: 27 February 2021
© Springer Nature B.V. 2021

Abstract

Object detection techniques that achieve state-of-the-art detection accuracy employ convolutional neural networks, implemented to have lower latency in graphics processing units. Some hardware systems, such as mobile robots, operate under constrained hardware situations, but still benefit from object detection capabilities. Multiple network models have been proposed, achieving comparable accuracy with reduced architectures and leaner operations. Motivated by the need to create a near real-time object detection system for a soccer team of mobile robots operating with x86 CPU-only embedded computers, this work analyses the average precision and inference time of multiple object detection systems in a constrained hardware setting. We train open implementations of MobileNetV2 and MobileNetV3 models with different underlying architectures, achieved by changing their input and width multipliers, as well as YOLOv3, TinyYOLOv3, YOLOv4 and TinyYOLOv4 in an annotated image dataset captured using a mobile robot. We emphasize the speed/accuracy trade-off in the models by reporting their average precision on a test data set and their inference time in videos at different resolutions, under constrained and unconstrained hardware configurations. Results show that MobileNetV3 models have a good trade-off between average precision and inference time in constrained scenarios only, while MobileNetV2 with high width multipliers are appropriate for server-side inference. YOLO models in their official implementations are not suitable for inference in CPUs.

Keywords Object detection · Convolutional neural networks · Mobile robots · Constrained hardware

1 Introduction

The recent successes in the field of object detection are mostly due to the use of deep neural networks, more specifically convolutional neural networks (CNN). The underlying operations that compose CNNs are highly optimized for fast execution in graphics processing units (GPU). However, in some domains, GPUs may be unavailable and these processes must be executed in CPUs. One such domain is

mobile robotics, comprised of systems such as autonomous terrestrial, aerial or underwater vehicles, as well as specialized units, such as humanoid robots. These systems may be subject to limitations regarding space, weight and energy consumption, constraining the robot's hardware to contain only CPUs, thus hindering the performance of systems based on deep learning techniques.

With these limitations in mind, this work presents an empirical and comparative analysis of the performance of recent CNN architectures proposed for constrained hardware settings and applied to object detection tasks, with the goal of emphasizing the speed/accuracy trade-off present in current models. Motivated by the creation of a near real-time soccer ball detection system to be executed in mobile robots equipped with embedded computers, that contain only CPUs under an x86 architecture and no graphics processing units, we train multiple models in the same annotated image dataset and compare their average precision (AP) in a test data set, as well as their inference times in both constrained and unconstrained hardware settings.

The authors acknowledge the São Paulo Research Foundation (FAPESP Grant 2019/07665-4) for supporting this project. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

✉ Douglas De Rizzo Meneghetti
douglasrizzo@fei.edu.br

Extended author information available on the last page of the article.

In preceding work [11], we presented a similar analysis of MobileNetV1, a CNN whose architecture was optimized for low latency inference in mobile phones, adapted for the task of object detection in an Intel NUC mini-PC equipped with a Core-i7 CPU. Here, we expand our work by analyzing MobileNetV2 [39] and MobileNetV3 [17], models that incorporate more recent advances in deep learning research to their architectures in order to achieve even lower latency in CPU-only mobile systems. We focus on benchmarking the performance of both architectures under different combinations of their width and input multipliers, two hyperparameters that control the underlying architecture of the MobileNets. In this work, we also include the YOLOv3 [36] and YOLOv4 [5] object detection models and their “tiny” counterparts, due to also being one-stage CNN-based object detectors.

With this study, we aim to fill an information gap related to the performance of the selected CNN models and object detection techniques, which are aimed at fast and accurate inference in mobile systems, when executed in a hardware platform typically used in mobile robots. In the case of the MobileNets, we investigate how a change in their width and input multipliers, two hyperparameters that control the overall topology of the model, affect their final accuracy and latency in the proposed task and hardware. For YOLO, we showcase their performance in the same task, as well as corroborate their high latency when executed in CPUs, especially less powerful ones.

We also expect, given our use case, to provide useful information to teams who participate in humanoid soccer competitions, regarding the capabilities of the selected models to perform near real-time inference with acceptable precision in a dataset with relevant applications to the competition. More broadly, this study is also aimed towards guiding the choice of a low latency, high accuracy object detection model to be executed in embedded computers with an x86 architecture, with a focus on recent MobileNet models that perform object detection using the Single Shot MultiBox Detector (SSD) [29] method, and multiple YOLO and TinyYOLO implementations.

This work also investigates the latency of the selected implementations of the models by gathering their inference times when processing videos in multiple input resolutions. These measurements are done in both constrained and unconstrained environments, providing readers with more relevant results to make decisions regarding the choice of neural network model to use in a single-object detection system under local, mobile, constrained hardware scenarios, but also accounting for situations in which remote and/or unconstrained hardware configurations may be available.

The text is organized as follows: Section 2 lists the recent advances in the state-of-the-art in object detection, as well as techniques to create smaller network topologies while

still maintaining high detection accuracy. We also describe the network architectures utilized in this work and their detection mechanisms. Section 3 depicts related works. The experimental methodology is presented in Section 4. Results are presented and discussed in Section 5. Lastly, Section 6 provides the conclusions and future work.

2 Research Background

In recent years, object detection techniques have advanced at great pace due to the equally fast advances of deep learning and convolutional neural networks applied to computer vision [49]. Two-stage detectors such as Faster R-CNN [37] first generate region proposals and then detect objects only in the selected regions, while one-stage detectors such as YOLO [34] and SSD [18] generate bounding box coordinates and class predictions at the same time, with YOLO using only convolutional layers for this task.

More recently, effort has centered around building strategies for efficiently scaling network models, reaching a trade-off between FLOPS, number of trainable parameters and accuracy. The MobileNetV3 architecture [17] has been partially achieved via hardware-aware neural architecture search techniques [42, 47], while AmoebaNet’s architecture [33], which achieved state-of-the-art classification accuracy on ImageNet [12], was evolved using evolutionary algorithms.

Other techniques attempt to shrink or expand the dimensions of convolutional layers using hyperparameters. MobileNetV1 [18] introduced the width multiplier and input resolution parameters, discussed later in the text, while EfficientNet [43] and EfficientDet [44] use a compound coefficient to scale all three dimensions of convolutional layers in order to maximize the network’s accuracy. This, allied with neural architecture search, introduced the current state-of-the-art in image classification and object detection using CNNs.

2.1 MobileNets

MobileNets [18] are convolutional neural network architectures whose number of trainable parameters can be controlled by two hyperparameters. The first is the *width multiplier* $\alpha \in (0, 1]$, which controls the number of channels in each layer of the network. Smaller values of α reduce the number of parameters in each layer of the network uniformly, also reducing computational cost. The second parameter is the *resolution multiplier* $\rho \in (0, 1]$, which is used to reduce the resolution of the input images and, consequently, the number of operations throughout all layers of the network.

Additional features introduced in MobileNetV1 are batch normalization [21] for learning stabilization, as well as depthwise-separable convolutions [40], a convolution operation that uses fewer parameters to achieve comparable results to regular convolutions.

MobileNetV2 [39] advanced the state-of-the-art by introducing linear bottleneck layers in the network, reducing the size of the inputs in subsequent layers while preventing information from being lost by non-linear activation functions. The ReLU6 non-linearity [26] was chosen instead of regular ReLU to prevent loss of information when calculations with low-precision data types are performed.

Finally, the MobileNetV3 [17] builds upon MobileNetV2 and MnasNet [42] by using Mobile Neural Architecture Search (MNAS) [42], an algorithm that manipulates blocks of layers in the network, with the goal of maximizing accuracy while ensuring inference is conducted under a latency budget. The architecture is further simplified by the use of the NetAdapt algorithm [47], which reduces the number of filters in each layer of each block created by MNAS, further reducing network latency while trying to minimally affect network accuracy. For each iteration that either NAS algorithm is applied to the current network, its inference latency is measured in a mobile phone to ensure that the resulting model architecture is optimized for that type of hardware [17].

Both MobileNetV2 as well as the gradual enhancements performed in MobileNetV3 have their latency measured in the ARM-based processor of the Google Pixel 1 [17, 39]. In this work, we expect the performance gain achieved by these enhancements to also be visible in our target hardware, which differs from the original hardware the networks were tested on.

2.2 Single-Shot MultiBox Detector

The Single-Shot MultiBox Detector (SSD) [29] is a technique that utilizes a convolutional neural network, called the base network, combined with multiple subsequent convolutional filters of different sizes, to perform detection under different scales and aspect ratios in multiple regions of an input image. The feature maps of the base network may be pretrained in a classification or detection class. When training the network for a detection task, SSD employs techniques such as data augmentation and hard negative mining for faster training, as well as a loss function that is a weighted sum of both localization and classification losses.

2.3 You Only Look Once

You Only Look Once (YOLO) [34] simplified the object detection problem, which was then composed of a

region proposal step followed by an image classification step, to a single regression step composed of bounding box coordinates and class probabilities. YOLOv2 [35] introduced the use of anchor boxes to the algorithm, a technique that allowed the detection of multiple objects with different aspect ratios in the same quadrant of an input image, using only convolutional layers.

YOLOv3 [36] introduces the prediction of bounding box coordinates across multiple scales and the use of residual layers [15] to speed up training, while YOLOv4 [5] adapts multiple data augmentation and feature extraction techniques to allow efficient training and inference of a model on a single GPU with 8 to 16 GB of VRAM.

3 Related Work

This work lies at the intersection of two problems, which may be considered complementary to each other. The first is the execution of computer vision tasks in embedded computers, which introduces complications regarding energy efficiency, storage and memory capacity and processing power. The second relates to compressing, reducing or simplifying a CNN model to achieve lower latency while sacrificing as little model accuracy as possible.

3.1 Neural Network Reduction Strategies

One way to make neural network inference feasible in embedded computers is to use techniques that simplify existing network architectures before deployment in the constrained hardware. The goal of this approach is to either reduce the amount of memory necessary to store the model; the number of multiply–accumulate operations (MAC) in a forward pass of the network, or other more direct metrics, such as the wall-clock inference time. Given the possible loss in representation power of the model resulting from its reduction, a loss of accuracy in the target task is usually expected as trade-off.

One model reduction technique is network pruning [14], which consists in the detection and removal of low-valued weights from a network model, followed by the retraining of the network, promoting the re-purposing of the remaining weights. There have been works that have applied pruning in order to make network models more fitting to be executed on embedded systems [1, 27].

Another approach to reduce the memory necessary to store a neural network is the quantization of the network's parameters [20], a technique which consists in representing numbers from a range of values by a less precise subset of values, consuming less memory and possibly making operations simpler. A more extreme

approach to quantization are the binary neural networks [10, 32], whose weights are represented by sets of binary or ternary values, enabling straightforward implementation of the models direct in the hardware level, at the cost of unstable training and reduced precision. Other works have already covered the applications of quantization [23, 25, 45] and binarization [25] of networks for use in embedded systems.

Model compression [2, 6], also known as knowledge distillation [16] is a process in which one shallow network (called the student) learns to imitate the output of one or more networks with more complex architectures (the teachers), achieving better results than if trained directly in the source dataset. This is usually done by using the predictions of the teacher models to classify or generate synthetic dataset, which is then used to train the student model. Thus, the student learns to approximate the function modeled by teacher ensemble instead of the function that represents the original data, which may be more accurate.

There have also been works that focus on representing neural network weights in alternative ways, *e.g.* using low-rank approximations of the weight matrices [38]. However, these techniques may require domain knowledge to be implemented without side-effects. Another option is to rethink the operations that compose the neural networks. One example of relevance to this work are depthwise separable convolutions [9, 40], which decompose the standard convolution operation applied to multiple channels into multiple single-channel convolutions, followed by a single pointwise convolution that joins the results from multiple channels. Depthwise separable convolutions are used in MobileNetV2 and V3, which are analyzed in this work.

Neural architecture search is yet another family of techniques aimed towards adapting a network's architecture according to an optimization objective and algorithm. A relevant example to this work is MNAS [42], a reinforcement learning-based technique which aims to maximize model accuracy and minimize wall-clock latency in mobile devices.

For individual works that build upon the aforementioned network reduction techniques, the reader is pointed to surveys on the topic [8, 38, 41].

3.2 CNNs in Embedded Systems

Due to the relative novelty of the use of deep convolutional neural networks in computer vision tasks, their transposition to more specialized and energy-efficient hardware is an ongoing topic of research. One category of such specialized hardware are the portable NVIDIA Jetson development boards, equipped with an ARM CPU and CUDA-capable GPU. There have been MobileNetV1 [31] and V2 [48]

implementations aimed towards object detection tasks in Jetson boards, as well as Fast R-CNN [30] and Faster R-CNN [22].

Previous work [7] has analyzed the inference time of multiple CNNs in image classification tasks under the Jetson TX1. A linear upper bound was observed when comparing model accuracy with inference time, *i.e.* models that took longer to perform a forward pass on an image were more accurate and the relation between both quantities was linear. In order to minimize memory usage and accelerate convolutions in the Jetson TX1, another work [27] explores the representation of CNN weights as sparse matrices and the use of different sparse matrix multiplication algorithms as alternatives to the convolution operation.

There have also been efforts to implement CNNs in field-programmable gate arrays (FPGA), such as the proposal of toolsets and methods to convert, accelerate and analyze the latency and precision of CNN operations in the new hardware [13, 46]. A common trend in these works is the observation of a negligible loss of precision due to the quantization of the parameters of the network [13, 25], which is greatly offset by a considerable gain in terms of energy efficiency and latency. Some works report a speed up of 10 to 16 times in inference when compared to executing the same architectures in the aforementioned Jetson boards [3, 13].

Another work [24] similar to this one has analyzed YOLOv3, YOLOv4, TinyYOLOv3, TinyYOLOv4 and a custom network both in a server equipped with an Intel Xeon E5-2678 v3 and an NVIDIA GeForce GTX 1080Ti, as well as a Raspberry Pi 3B. Like this work, they have reported prohibitively high inference times in their constrained hardware, ranging from 3.225 seconds to 5.555 seconds when doing a forward pass in a single image. Unlike them, we provide an alternative with faster inference times for our constrained hardware.

3.3 Discussion

This section has presented related works in two main areas of which this work is an intersection of: neural network reduction techniques and the execution of neural networks in embedded devices.

This work differs from its related works by focusing on the analysis of convolutional network architectures applied to object detection tasks in mobile computers with an x86 architecture processor. More specifically, we analyse the MobileNetV2 and V3 neural networks in a different target hardware than their original one and how their precision and inference times differ when a set of hyperparameters that dictate their architectures (input and width multipliers) change. We also provide the same analysis for the YOLOv3, YOLOv4, TinyYOLOv3 and TinyYOLOv4 networks, due

to their prominence in object detection tasks with high frame rate needs.

4 Experimental Methodology

This study is motivated by the development of a computer vision system for an autonomous humanoid robot, which provides the robot with the capabilities of detecting soccer balls in a time that is compatible with the dynamics of the game. The target hardware of the vision system is CPU-only embedded computer of the x86 architecture, such as a mini-PC. We identify two properties that are desirable in an object detection system, given the aforementioned requirements. The first is high detection accuracy, a desirable feature in all object detection systems. The second is low latency, providing the robot with updated object locations under actionable intervals of time.

These two properties can be considered opposite, since higher accuracy usually indicates a more complex model, with more layers and operations. Also, when working with mobile robots and embedded computers, there is a tendency towards employing simpler machine learning models, in order to achieve acceptable operating thresholds with regards to time and power consumption, while sacrificing as little precision as possible.

To compose a detection system with low latency and high accuracy, we have selected multiple recent neural network architectures developed for computer vision tasks which attempt to reach a balance between both properties. Models were selected according to their performance in object detection tasks, their availability for the x86 processor architecture and homogeneity in pretraining in the MSCOCO [28] dataset.

The remainder of this section describes the selected neural network models; the image dataset used to train the models and analyze their performance; the humanoid robot our experiments are geared towards (with no loss of generality over other x86 embedded systems); and the hardware used for training, as well as the different hardware used for testing the models under constrained and unconstrained scenarios, providing even more broad and informative results.

4.1 Network Architectures and Training

Twenty MobileNetV2 configurations were tested by modifying the values for the width and resolution multipliers. For the width multiplier, the values 1, 0.75, 0.5 and 0.35 were used, resulting in networks with 3.47, 2.61, 1.95 and 1.66 million trainable parameters, respectively. The values used for the resolution multiplier were chosen so that the input resolution of the network is equal to 224, 192, 160, 128 and

96. The combined values of both hyperparameters resulted in a total of twenty models that were trained using the soccer ball dataset, described on Section 4.2.

MobileNetV3 [17] models are composed of the “Large” and “Small” variants, both with width multipliers of 1 and 0.75, possessing 5.4 (Large, $\alpha = 1$), 4 (Large, $\alpha = 0.75$), 2.9 (Small, $\alpha = 1$) and 2.4 million (Small, $\alpha = 0.75$) trainable parameters, as well as minimalistic versions of both variants with $\alpha = 1$, possessing 3.9 and 2 million parameters. Minimalistic models do not contain the more advanced squeeze-and-excite units, hard-swish, and 5×5 convolutions operations from the non-minimalistic counterparts. YOLO models are composed of the v3 [36] and v4 [5] versions of the neural networks, as well as their “tiny” counterparts.

The MobileNet implementations selected for this work are provided in the TensorFlow Object Detection API [19], while YOLO is provided in the original paper implementation [5]. All models used pretrained weights learned in the COCO dataset [28].¹

4.2 Image Dataset

The dataset used in this work [4]² consists of 4364 images in 1920×1080 resolution, collected from the point-of-view of the humanoid robot described in Section 4.4. A fish-eye lens is used to maximize the field of view of the robot and so all images in the dataset also inherit this feature.

Of the 4364 images, 4014 compose the annotated training set and 250, the annotated test set. In these sets, the soccer balls visualized by the robot have been marked with bounding boxes. The training and test sets were collected from different sets of videos, with the purpose of minimizing data correlation.

Each image contains a single soccer ball, captured under multiples lighting conditions, as well as at different angles and distances from the camera. There are pictures of both stationary and moving soccer balls. Figure 1b presents examples of the dataset.


4.3 Training Procedure

The models were trained in a server with Intel Xeon Gold 5118@2.3 GHz processors totaling 48 CPUs, 192 GB of RAM and an NVIDIA Tesla V100-PCIE with 16 GB of memory, running CentOS 7.6.1810. The MobileNet models were trained for 50000 training steps. The YOLO

¹In order to facilitate the replication of these results and encourage the development of similar approaches by other researchers, the software used in this paper is available for download at: <https://github.com/douglasrizzo/JINT2020-ball-detection>.

²The dataset was also made available at <http://iee-dataport.org/open-access/open-soccer-ball-dataset>

Fig. 1 Image dataset and humanoid robot used to collect it. **a** One of the teen-sized robots of the RoboFEI team. **b** Examples of images from the soccer ball dataset



and TinyYOLO models were trained for a total of 6000 training steps, following recommendations from the original developers of the model, given the number of classes to be detected.

All MobileNetV2 models were trained using batches of 32 images and the RMSProp optimizer with initial learning rate of $4 \cdot 10^{-3}$, an exponential decay schedule with a decay factor of 0.95 and a momentum coefficient of 0.9. All MobileNetV3 models were trained using batches of 32 images, stochastic gradient descent with initial learning rate of 0.4, a cosine decay schedule and a momentum coefficient of 0.9.

YOLO and TinyYOLO models were trained using batches of 64 images, stochastic gradient descent with a momentum coefficient of 0.9. YOLOv3 and TinyYOLOv3 models used a learning rate of 10^{-3} , while YOLOv4 and TinyYOLOv4 used a learning rate of $2 \cdot 10^{-3}$. Two step decays at 80% and 90% of the training were applied to these learning rates.

4.4 Humanoid Robot

Our use case for this work is a mobile humanoid robot built for the task of playing soccer. One of its capabilities is the detection of soccer balls through visual inputs. The robot weighs about 5.9 kg and measures 81 cm in height. It is composed of 19 Dynamixel servomotors (a combination of MX-64, MX-106 and XM430 models), totaling 19 degrees of freedom. The humanoid robot uses a Genius WideCam F100 (Full HD) camera for image capture and a CH Robotics UM7 orientation sensor. The center of mass has a height of 36.1 cm and the robot has a foot area of 174 cm^2 . Other measurements include 39.5 cm of shoulder length, 38.5 cm of leg height, 18.8 cm of neck height and 38.5 cm of arm length. The robot is equipped with an Intel NUC Core i7 mini-PC. A picture of the robot is presented in Fig. 1a.

4.5 Performance Metrics

To evaluate the performance of the selected models in both time and correctness, we gathered their average precision and inference time during the experiments. The meaning of these two measures is explained in this section.

4.5.1 Average Precision

In this work, both ground-truth annotations and predictions are represented as rectangular bounding boxes. Given a predicted rectangle A and a ground truth rectangle B , it is possible to calculate a similarity measure between both rectangles, called *intersection over union* (IoU),

$$IoU = \frac{(A \cap B)}{(A \cup B)}$$

When the IoU among both rectangles surpasses a certain threshold t , the detection represented by A is considered a true positive (TP). Otherwise, if no other predicted rectangle has an acceptable $IoU \geq t$ wrt. B , the failed detection is considered a false negative (FN). Lastly, all predicted boxes which have $IoU < t$ for all ground truth boxes in an image are considered false positives (FP). These quantities allow us to calculate a detector's precision and recall,

$$Precision = \frac{TP}{TP + FN}, \quad Recall = \frac{TP}{TP + FP}.$$

By ordering all predictions in order of decreasing confidence and calculating precision and recall at each new prediction, we arrive at a monotonically decreasing precision/recall (PR) curve. Average precision (AP) is the area under the PR curve. It is a single-valued metric which summarizes model performance when retrieving objects of a single class.

4.5.2 Inference Time

In this work, inference time is defined as the time it takes a detection model to generate predictions for a given input image, in milliseconds. To get a more accurate measure, all models operated over a 30-second video taken from the robot’s point-of-view and the inference time of the model is taken as the average time over all frames of the video. Experiments were conducted in the same video at the native resolution of 1920×1080 pixels, as well as other versions scaled down to 1280×720 , 640×480 and 480×360 pixels. All networks then processed these four versions of the same video and the mean inference time over all frames of each video for each network was recorded.

4.6 Constrained Hardware for Inference

The constrained hardware configuration in which the inference time of the models was captured is equipped with an i5-4210U CPU @ 1.70GHz and 8 GB of RAM and no GPU, in line with the hardware typically used by an autonomous mobile robot. For comparison purposes, the same experiments were performed in the training computer, under GPU and CPU-only settings. In all cases, when CPU-only experiments were executed, all CPU cores were allowed to be utilized.

5 Results

Table 1 displays the AP of all trained models in the test set, as well as the inference time in milliseconds for different hardware configurations. In this table, the inference time was calculated with the videos in their native 1920×1080 resolution. These results allow us to conclude that the canonical YOLOv3 and YOLOv4 implementations, as well as their tiny counterparts, are not optimized for inference on CPUs, achieving the highest inference times of all models in the Intel Core i5-4210U. In fact, it is stated by the author of the model [5] that their implementations are optimized for inference in single GPUs. This can be seen by the comparatively low times achieved in the Tesla V100 GPU, especially by the TinyYOLO models, which achieved the lowest inference times of all models.

As for the MobileNet models, we can see that MobileNetV3 and MobileNetV2 with $\alpha = 1$ achieved the highest AP in the test data set. However, with the high variation in inference times between all combinations of hyperparameters, the results from Table 1 alone do not provide enough information to compare model performances. To remedy that, we calculate a performance score for each neural network in each hardware setting

$p_{m,h} = \frac{AP_m}{t_{m,h}}$, where AP_m represents the AP of network m (a value that is hardware-independent) and h , the hardware setting the inference time t of model m was gathered from. Then, the performance scores of all models in the same hardware are normalized by the highest performance score in that hardware, leading to the normalized score $s_{m,h} = \frac{p_{m,h}}{\max_{\eta} p_{\eta,h}}$. Achieving a normalized score $s_{m,h} = 1$ means that model m had the highest AP/inference time ratio off all models in hardware h .³

Table 2 presents the normalized scores of all models. Overall, MobileNetV2 models with width multipliers $\alpha \in \{0.75, 1\}$ had the best scores of all MobileNetV2 models in all hardware settings. We can also see that the five best models in unconstrained hardware settings are the same for both CPU and GPU. However, when operating under the Intel Core i5 4210U processor, both MobileNetV3 models (large and small) with $\alpha = 0.75$ achieved the highest scores, making MobileNetV3 models a viable option for an object detection system that operates under constrained hardware settings, whereas they did not exhibit the same performance in GPUs.

5.1 Performance with Different Input Video Resolutions

This section presents the inference times in milliseconds of all model implementations when processing input videos of several resolutions (1920×1080 , 1280×720 , 640×480 , 480×360). Table 3 presents the mean and standard deviation of the inference time of all YOLO models for the tested hardware configurations. Overall, all YOLO and TinyYOLO models achieved low standard deviation in this test, implying that their implementation [5] is indifferent to the input resolution of images.

As for the MobileNet results, we first discuss the distributions of results in the three hardware settings. In total, 104 values were collected in each setting (twenty V2 and six V3 models applied to videos in four resolutions), with the following means and standard deviations: $\mu_{i5} = 68.292$, $\sigma_{i5} = 17.374$, $\mu_{V100} = 47.24$, $\sigma_{V100} = 8.756$, $\mu_{Xeon} = 47.009$ and $\sigma_{Xeon} = 8.981$.

Due to the similarity in the inference times collected from the NVIDIA Tesla V100 GPU and the Intel Xeon Gold 5118, we executed a two-sample Kolmogorov-Smirnov test between the two samples, with a result of $p = 0.97371$, indicating that the measurements collected in the 48 quad-core CPUs and the single GPU are similar with high

³This score may easily break or be less informative if one neural network in the sample has disproportionately low inference time or high AP. However, given the well-behaved values presented in Table 1, we consider the use of the proposed score appropriate for the purposes of our analysis.

Table 1 Average precision (higher is better) and inference time in milliseconds (lower is better) of the trained models

Network	Width mult.	Input res.	AP	Inference time (ms)		
				Core i5	V100	Xeon
MobileNetV2	0.35	96	0.4065	99.964	65.37	64.959
		128	0.7095	101.166	51.062	49.642
		160	0.6304	88.651	52.642	50.97
		192	0.6756	87.006	53.613	52.232
		224	0.4984	78.553	42.852	41.703
	0.5	96	0.4065	91.403	58.919	57.626
		128	0.6986	81.08	44.529	43.388
		160	0.3361	91.775	58.288	58.616
		192	0.0944	116.076	65.629	64.828
		224	0.3253	78.624	42.759	43.18
	0.75	96	0.7284	86.569	51.065	51.528
		128	0.6954	84.159	42.097	41.866
		160	0.6679	81.351	41.883	42.309
		192	0.6952	78.699	42.347	41.776
		224	0.7874	85.186	48.343	47.854
1	96	0.8133	122.853	56.992	57.83	
	128	0.7672	82.277	46.921	47.799	
	160	0.8597	88.886	52.278	52.569	
	192	0.3632	110.75	61.263	60.052	
	224	0.8177	79.547	42.438	42.183	
MobileNetV3 (large min.)	1	224	0.6007	85.808	58.706	59.581
MobileNetV3 (large)	0.75	224	0.8847	89.362	63.515	63.703
MobileNetV3 (large)	1	224	0.6875	120.045	88.017	91.369
MobileNetV3 (small min.)	1	224	0.6024	79.142	48.68	49.236
MobileNetV3 (small)	0.75	224	0.7067	60.654	49.328	47.741
MobileNetV3 (small)	1	224	0.8651	96.975	70.689	70.017
TinyYOLOv3			0.3381	588.235	33.557	85.47
TinyYOLOv4			0.3504	714.286	29.851	119.048
YOLOv3			0.1355	5000	44.248	588.235
YOLOv4			0.1419	5000	50	833.333

The five best in each category are marked in bold

statistical relevance. Because of that, in this section we only report results for the MobileNets in the NVIDIA Tesla V100 GPU and the Intel i5-4210U processor.

The distance between μ_{i5} and μ_{V100} is indicative of the performance lost by executing deep learning models in constrained CPUs, while a larger standard deviation on the CPU (σ_{i5}) indicate that there is a larger variation in network performance, given the resolution of the input video.

Figures 2 and 3 present a series of boxplots containing the inference time by frame in milliseconds for the MobileNets when processing the same video under multiple resolutions in the Intel i5-4210U CPU and the NVIDIA Tesla V100 GPU, respectively. The central line in each box represents the median value in the group; boxes represent the first

and third quartiles of the data; and whiskers represent the minimum and maximum acceptable range, whereas dots outside the whiskers are considered outliers.

In all of the figures, it is possible to observe a gradual decrease in inference time as the resolution of the input videos decreases. This information may be relevant, as the largest input resolution used by a MobileNet model is 224×224 . Furthermore, the implementations used in this work [19] already operate in downscaled images, with a resolution of 300×300 pixels. Both of these observations indicate that using a low-resolution input feed for object detection is a valid strategy to achieve lower inference times. This speedup is visualized in both constrained CPU and GPU settings.

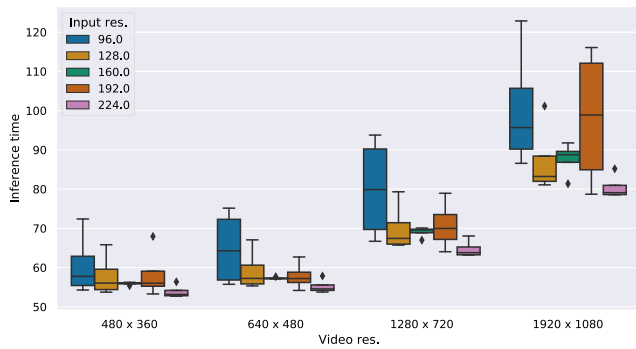
Table 2 Normalized scores of the trained models under different hardware

Network	Width mult.	Input res.	Normalized score		
			Core i5	V100	Xeon
MobileNetV2	0.35	96	0.349	0.323	0.323
		128	0.602	0.721	0.737
		160	0.61	0.622	0.638
		192	0.666	0.654	0.667
		224	0.545	0.604	0.617
	0.5	96	0.382	0.358	0.364
		128	0.74	0.814	0.831
		160	0.314	0.299	0.296
		192	0.07	0.075	0.075
		224	0.355	0.395	0.389
	0.75	96	0.722	0.74	0.729
		128	0.709	0.857	0.857
		160	0.705	0.828	0.814
		192	0.758	0.852	0.858
		224	0.793	0.845	0.849
	1	96	0.568	0.741	0.726
		128	0.8	0.849	0.828
		160	0.83	0.853	0.844
		192	0.281	0.308	0.312
		224	0.882	1	1
MobileNetV3 (large min.)	1	224	0.601	0.531	0.52
MobileNetV3 (large)	0.75	224	0.85	0.723	0.716
MobileNetV3 (large)	1	224	0.492	0.405	0.388
MobileNetV3 (small min.)	1	224	0.653	0.642	0.631
MobileNetV3 (small)	0.75	224	1	0.744	0.764
MobileNetV3 (small)	1	224	0.766	0.635	0.637
TinyYOLOv3			0.049	0.523	0.204
TinyYOLOv4			0.042	0.609	0.152
YOLOv3			0.002	0.159	0.012
YOLOv4			0.002	0.147	0.009


Higher is better. A normalized score of 1 indicates that the model performed the best under that hardware, compared with the others

Table 3 Inference time of YOLO and TinyYOLO models when processing videos of multiple input resolutions


Network	Hardware	Inference time	
		mean	std. dev.
TinyYOLOv3	Tesla V100	41.957	10.022
	Xeon Gold 5118	88.983	3.369
	i5-4210U	588.235	0.000
TinyYOLOv4	Tesla V100	38.808	9.136
	Xeon Gold 5118	114.529	6.094
	i5-4210U	714.286	0.000
YOLOv3	Tesla V100	46.726	4.081
	Xeon Gold 5118	588.235	0.000
	i5-4210U	5000.000	0.000
YOLOv4	Tesla V100	50.385	0.676
	Xeon Gold 5118	817.308	32.051
	i5-4210U	5000.000	0.000




(a) MobileNetV2 by input resolution on i5-4210U




(a) MobileNetV2 by input resolution on V100



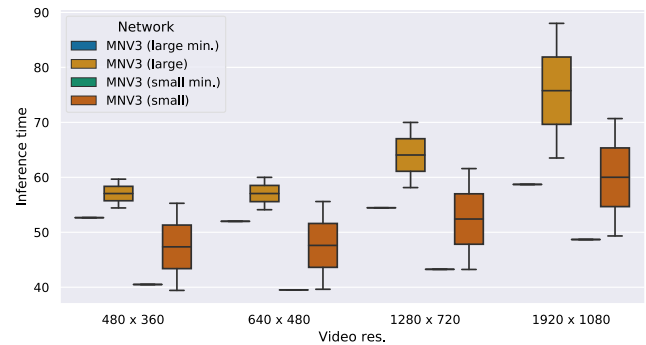
(b) MobileNetV2 by width multiplier on i5-4210U



(b) MobileNetV2 by width multiplier on V100



(c) MobileNetV3 on i5-4210U



(c) MobileNetV3 on V100

Fig. 2 Inference time of MobileNetV2 and V3 models in the Intel i5-4210U

Fig. 3 Inference time of MobileNetV2 and V3 models in the NVIDIA Tesla V100

To discover whether the distributions presented in Figs. 2 and 3 can be considered different with statistical significance, a series of one-way analysis of variance (ANOVA) tests were performed, whose results can be viewed in Table 4. We can see that MobileNetV2 is sensitive to the input resolution hyperparameter only when inference is done on a GPU and not the constrained CPU.

The other results in Table 4 also confirm that there is a difference in feeding video frames in different resolutions to the TensorFlow implementations of the MobileNets. MobileNetV2 demonstrated statistically significant differences both on the constrained CPU and the GPU,

while MobileNetV3 demonstrated statistically significant differences only in the constrained CPU. Lastly, in all cases, a change in the width multiplier of the networks did not result in a significant change in their inference times.

To discover which combinations of the aforementioned parameters actually resulted in lower inference times for each neural network, the ANOVA tests were followed by a series of Tukey's Honestly Significant Difference (HSD) tests, whose results are displayed in Figs. 4 and 5. Whenever there is no vertical overlap between the intervals covered by two pairs of lines, their distributions are considered different with a statistical significance greater than 95%. It can be

Table 4 One-way analysis of variance of the inference times achieved by the MobileNets

Network	Parameter	Hardware	<i>F</i>	<i>p</i>
MobileNetV2	Input res.	Core i5	1.476	0.2176
		V100	6.773	0.0001
	Width mult.	Core i5	0.639	0.5921
		V100	2.606	0.0577
	Video res.	Core i5	63.744	1.0595 · 10⁻²⁰
		V100	14.929	9.8719 · 10⁻⁸
MobileNetV3	Width mult.	Core i5	2.471	0.1302
		V100	1.873	0.1848
	Video res.	Core i5	7.350	0.0016
		V100	1.885	0.1646

Values in bold indicate results that were different with $p \leq 5\%$

seen on Fig. 4a that MobileNets with input resolutions of 128 and 224 pixels achieve lower inference time than the ones with input resolutions of 96 and 192 pixels.

The other subfigures confirm that feeding video frames in a natively low resolution to the MobileNets result in lower inference times. In all cases, the optimal video resolutions are 640×480 and 480×360 , in contrast with 1920×1080 . In the case of MobileNetV2 on CPU, 1280×720 achieves a statistically significant compromise among the two aforementioned groups.

5.2 Discussion

The results presented in Tables 1 and 2 show that MobileNetV2 models with $\alpha \leq 0.5$ do not achieve

the top results with regards to AP. While some of the V2 models do achieve low inference times (in bold in Table 1), the ANOVA results from Table 4 indicate that the differences in latency resulting from a change in the α hyperparameter are not statistically significant, which leads us to the conclusion that, under our experimental settings, MobileNetV2 models with $\alpha > 0.5$ should be prioritized.

Another interesting observation is that MobileNetV3 is the only model whose latency lowered (in comparison with the other models) in the less powerful CPU, in comparison with the server CPU and GPU. This is a strong indication that the optimizations performed in MobileNetV3 [19] on top of MobileNetV2 and MnasNet, which initially targeted





Fig. 4 Tukey's HSD tests applied to the MobileNet groups who had $p \leq 5\%$ in the ANOVA tests on the Intel Core i5-4210U





Fig. 5 Tukey's HSD tests applied to the MobileNet groups who had $p \leq 5\%$ in the ANOVA tests on The NVIDIA

a performance gain in ARM-based mobile phones, are also visible in Intel-based x86 processors.

Lastly, since the MobileNets studied in this work receive images with a maximum input size of 224×224 pixels, the observation that using video feeds with smaller resolutions boost the inference time of all MobileNets is a clear indicator that low latency object detection applications should always prioritize using the lowest video resolution possible to avoid any overhead related to processing or resizing high definition images before each forward pass.

As for YOLO, its high latency in CPUs can be attributed to the optimization of its canonical implementation towards single GPU systems [5]. More specifically, the particularly high latency in low power CPUs achieved in this work for all YOLO models is in line with related work that also performed inference time measurements of the same models, albeit in different hardware [24].

6 Conclusions

This work presented a comparative analysis of the precision and inference time of reduced CNN architectures and single-stage object detectors in the task of single-class object detection under constrained hardware situations. Multiple MobileNetV2 and V3 models, as well as YOLOv3, YOLOv4 and their tiny counterparts were trained to detect soccer balls and tested in a computer with an Intel Core i5-4210U CPU and no graphics capabilities, following the use case of a humanoid robot operating with an embedded computer.

We used open implementations of all networks, provided either from the TensorFlow Object Detection API [19] or from the official YOLO repositories [5] and compared their average precision in a test data set, as well as their inference times in the constrained CPU setting, an unconstrained CPU setting and a server-class GPU, when operating on videos under multiple resolutions.

Results have shown that MobileNetV2 models with high width multipliers have the best trade-off between average precision and inference time in unconstrained hardware settings, being suitable when executing inference in remote servers is an option. However, while MobileNetV3 models did not show remarkable performance when operating in the unconstrained hardware settings, a performance boost was observed when inference was executed under a local, constrained, CPU-only scenario. Lastly, the official implementations of YOLO and TinyYOLO, being optimized for inference in GPUs, displayed poor results in our low-end Intel Core i5-4210U processor.

Author Contributions

- Conceptualization: D. R. Meneghetti; T. P. D. Homem; J. H. R. de Oliveira; R. A. C. Bianchi
- Methodology: D. R. Meneghetti; T. P. D. Homem; D. H. Perico
- Software: D. R. Meneghetti
- Investigation: D. R. Meneghetti
- Formal Analysis: D. R. Meneghetti
- Validation: D. R. Meneghetti
- Data curation: D. R. Meneghetti; T. P. D. Homem; J. H. R. de Oliveira; I. J. da Silva; D. H. Perico; R. A. C. Bianchi
- Writing – original draft: D. R. Meneghetti; T. P. D. Homem
- Writing – review & editing: D. R. Meneghetti; T. P. D. Homem; D. H. Perico; R. A. C. Bianchi
- Visualization: D. R. Meneghetti; T. P. D. Homem; R. A. C. Bianchi
- Resources: D. R. Meneghetti; R. A. C. Bianchi
- Funding acquisition: D. R. Meneghetti; R. A. C. Bianchi
- Project administration: R. A. C. Bianchi
- Supervision: R. A. C. Bianchi

Funding The authors acknowledge the São Paulo Research Foundation (FAPESP Grant 2019/07665-4) for supporting this project. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

Materials Availability The scripts utilized in the experiments presented in this paper are available at <https://github.com/douglasrizzo/JINT2020-ball-detection>. The image dataset used in the experiments was also made available online [4], along with a static copy of the aforementioned scripts.

Declarations

Conflict of Interests The authors declare that they have no conflict of interest.

References

1. Alippi, C., Disabato, S., Roveri, M.: Moving convolutional neural networks to embedded systems: the alexnet and VGG-16 case. In: 2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN), pp. 212–223 (2018). <https://doi.org/10.1109/IPSNS.2018.00049>
2. Ba, L.J., Caruana, R.: Do deep nets really need to be deep? arXiv:1312.6184[cs] (2014)
3. Bettoni, M., Urgese, G., Kobayashi, Y., Macii, E., Acquaviva, A.: A convolutional neural network fully implemented on FPGA for embedded platforms. In: 2017 New Generation of CAS (NGCAS), pp. 49–52 (2017). <https://doi.org/10.1109/ngcas.2017.16>
4. Bianchi, R.A.D.C., Perico, D.H., Homem, T.P.D., da Silva, I.J., Meneghetti, D.D.R.: Open soccer ball dataset. IEEE Dataport. <https://doi.org/10.21227/0vvr-5c61> (2020)
5. Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: YOLOv4: optimal speed and accuracy of object detection. arXiv:2004.10934[cs, eess] (2020)
6. Buciluă, C., Caruana, R., Niculescu-Mizil, A.: Model compression. In: Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Min-

- ing - KDD '06, p. 535. ACM Press, Philadelphia (2006). <https://doi.org/10/fkdh9m>
7. Canziani, A., Culurciello, E., Paszke, A.: Evaluation of neural network architectures for embedded systems. In: 2017 IEEE International Symposium on Circuits and Systems (ISCAS), pp. 1–4 (2017). <https://doi.org/10.1109/ISCAS.2017.8050276>
 8. Cheng, Y., Wang, D., Zhou, P., Zhang, T.: A survey of model compression and acceleration for deep neural networks. *IEEE Signal Process. Mag.* **35**(1), 126–136 (2020). <https://doi.org/10.1109/MSP.2017.2765695>
 9. Chollet, F.: Xception: deep learning with depthwise separable convolutions. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1800–1807 (2017). <https://doi.org/10.1109/CVPR.2017.195>
 10. Courbariaux, M., Hubara, I., Soudry, D., El-Yaniv, R., Bengio, Y.: Binarized neural networks: training deep neural networks with weights and activations constrained to +1 or -1. arXiv:1602.02830[cs] (2016)
 11. de Oliveira, J.H.R., da Silva, I.J., Homem, T.P.D., Meneghetti, D.D.R., Perico, D.H., Bianchi, R.A.D.C.: Object detection under constrained hardware scenarios: a comparative study of reduced convolutional network architectures. In: 2019 XVI Latin American Robotics Symposium and VII Brazilian Robotics Symposium (LARS/SBR). IEEE (2019)
 12. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition, pp. 248–255 (2009). <https://doi.org/10/cvc7xp>
 13. Guo, K., Sui, L., Qiu, J., Yu, J., Wang, J., Yao, S., Han, S., Wang, Y., Yang, H.: Angel-eye: a complete design flow for mapping CNN onto embedded FPGA. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **37**(1), 35–47 (2018). <https://doi.org/10/gf2ntg>
 14. Han, S., Pool, J., Tran, J., Dally, W.: Learning both weights and connections for efficient neural network. *Advances in Neural Information Processing Systems* **28**, 1135–1143 (2015)
 15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016). <https://doi.org/10.1109/cvpr.2016.90>
 16. Hinton, G., Vinyals, O., Dean, J.: Distilling the Knowledge in a Neural Network. arXiv:1503.02531[cs, stat] (2015)
 17. Howard, A., Sandler, M., Chu, G., Chen, L.C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., Le, Q.V., Adam, H.: Searching for MobileNetV3. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 1314–1324 (2019). https://openaccess.thecvf.com/content/ICCV_2019/html/Howard_Searching_for_MobileNetV3_ICCV_2019_paper.html
 18. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. arXiv:1704.04861 (2017)
 19. Huang, J., Rathod, V., Sun, C., Zhu, M., Korattikara, A., Fathi, A., Fischer, I., Wojna, Z., Song, Y., Guadarrama, S., Murphy, K.: Speed/accuracy trade-offs for modern convolutional object detectors. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2017). <https://doi.org/10.1109/cvpr.2017.351>
 20. Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., Bengio, Y.: Quantized Neural Networks: Training Neural Networks with Low Precision Weights and Activations. arXiv:1609.07061[cs] (2016)
 21. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: International Conference on Machine Learning, pp. 448–456 (2015). <http://proceedings.mlr.press/v37/ioffe15.html>
 22. Jaramillo-Avila, U., Anderson, S.R.: Foveated image processing for faster object detection and recognition in embedded systems using deep convolutional neural networks. In: Martinez-Hernandez, U., Vouloutsis, V., Mura, A., Mangan, M., Asada, M., Prescott, T.J., Verschure, P.F. (eds.) *Biomimetic and Biohybrid Systems, Lecture Notes in Computer Science*, pp. 193–204. Springer International Publishing, Cham (2019). https://doi.org/10.1007/978-3-030-24741-6_17
 23. Jian, B., Yu, C., Jinshou, Y.: Neural networks with limited precision weights and its application in embedded systems. In: 2010 Second International Workshop on Education Technology and Computer Science, vol. 1, pp. 86–91 (2010). <https://doi.org/10.1109/ETCS.2010.448>
 24. Jiang, Z., Zhao, L., Li, S., Jia, Y.: Real-time object detection method based on improved YOLOv4-tiny. arXiv:2011.04244[cs] (2020)
 25. Jiao, L., Luo, C., Cao, W., Zhou, X., Wang, L.: Accelerating low bit-width convolutional neural networks with embedded FPGA. In: 2017 27th International Conference on Field Programmable Logic and Applications (FPL), pp. 1–4 (2017). <https://doi.org/10.23919/FPL.2017.8056820>
 26. Krizhevsky, A.: Convolutional deep belief networks on CIFAR-10. Tech. rep. (2010)
 27. Li, Q., Xiao, Q., Liang, Y.: Enabling high performance deep learning networks on embedded systems. In: IECON 2017 - 43rd Annual Conference of the IEEE Industrial Electronics Society, pp. 8405–8410 (2017). <https://doi.org/10/ghpz8h>
 28. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *Computer Vision – ECCV 2014, Lecture Notes in Computer Science*, pp. 740–755. Springer International Publishing, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_48
 29. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: SSD: single shot multibox detector. In: *Computer Vision – ECCV 2016*, vol. 9905, pp. 21–37. Springer International Publishing (2016). https://doi.org/10.1007/978-3-319-46448-0_2
 30. Mao, H., Yao, S., Tang, T., Li, B., Yao, J., Wang, Y.: Towards real-time object detection on embedded systems. *IEEE Transactions on Emerging Topics in Computing* **6**(3), 417–431 (2018). <https://doi.org/10/gd7rvr>
 31. Niazi-Razavi, M., Savadi, A., Noori, H.: Toward real-time object detection on heterogeneous embedded systems. In: 2019 9th International Conference on Computer and Knowledge Engineering (ICCKE), pp. 450–454 (2019). <https://doi.org/10/ghpz8c>
 32. Qin, H., Gong, R., Liu, X., Bai, X., Song, J., Sebe, N.: Binary neural networks: a survey. *Pattern Recogn.* **105**, 107281 (2020). <https://doi.org/10/ggs3g4>
 33. Real, E., Aggarwal, A., Huang, Y., Le, Q.V.: Regularized evolution for image classifier architecture search. In: 33rd AAAI Conference on Artificial Intelligence, AAAI 2019 (2019). arXiv:1802.01548
 34. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: unified, real-time object detection. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2015). <https://doi.org/10.1109/cvpr.2016.91>
 35. Redmon, J., Farhadi, A.: YOLO9000: better, faster, stronger. arXiv:1612.08242[cs] (2016)
 36. Redmon, J., Farhadi, A.: YOLOv3: an incremental improvement. arXiv:1804.02767[cs] (2018)
 37. Ren, S., He, K., Girshick, R.B., Sun, J.: Faster r-CNN: towards real-time object detection with region proposal networks. arXiv:1506.01497. <https://doi.org/10.1109/tpami.2016.2577031> (2015)
 38. Roth, W., Schindler, G., Zöhrer, M., Pfeifenberger, L., Peharz, R., Tschiatschek, S., Fröning, H., Pernkopf, F., Ghahramani,

- Z.: Resource-Efficient Neural Networks for Embedded Systems. arXiv:2001.03048[cs, stat] (2020)
39. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: inverted residuals and linear bottlenecks. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4510–4520 (2018). <https://doi.org/10/gfxgjj>
 40. Sifre, L.: Rigid-Motion Scattering for Image Classification. Ph. D. Thesis, Ecole Polytechnique, CMAP, Palaiseau, France (2014)
 41. Sze, V., Chen, Y., Yang, T., Emer, J.S.: Efficient processing of deep neural networks: a tutorial and survey. *Proc. IEEE* **105**(12), 2295–2329 (2017). [10/gcnp38](https://doi.org/10/gcnp38)
 42. Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., Le, Q.V.: MnasNet: Platform-Aware Neural Architecture Search for Mobile. arXiv:1807.11626[cs] (2019)
 43. Tan, M., Le, Q.V.: EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. arXiv:1905.11946[cs, stat] (2020)
 44. Tan, M., Pang, R., Le, Q.V.: EfficientDet: Scalable and Efficient Object Detection. arXiv:1911.09070[cs, eess] (2020)
 45. Tripathi, S., Dane, G., Kang, B., Bhaskaran, V., Nguyen, T.: LCDet: low-complexity fully-convolutional neural networks for object detection in embedded systems. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 411–420 (2017). <https://doi.org/10/ghpz79>
 46. Venieris, S.I., Kouris, A., Bouganis, C.S.: Deploying Deep Neural Networks in the Embedded Space. arXiv:1806.08616[cs] (2018)
 47. Yang, T.J., Howard, A., Chen, B., Zhang, X., Go, A., Sandler, M., Sze, V., Adam, H.: NetAdapt: platform-aware neural network adaptation for mobile applications. In: European Conference on Computer Vision (ECCV) (2018). https://openaccess.thecvf.com/content/ECCV_2018/papers/Tien-Ju_Yang_NetAdapt_Platform-Aware_Neural_ECCV_2018_paper.pdf
 48. Zhao, Z., Zhang, Z., Xu, X., Xu, Y., Yan, H., Zhang, L.: A lightweight object detection network for real-time detection of driver handheld call on embedded devices. <https://www.hindawi.com/journals/cin/2020/6616584/> <https://doi.org/10.1155/2020/6616584> (2020)
 49. Zhao, Z.Q., Zheng, P., Xu, S.T., Wu, X.: Object detection with deep learning: a review. *IEEE Transactions on Neural Networks and Learning Systems*, pp 1–21. <https://doi.org/10.1109/tnnls.2018.2876865> (2019)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Douglas De Rizzo Meneghetti is a Doctorate student in Electrical Engineering from the University Center FEI, where he does research in artificial intelligence and machine learning, focusing on reinforcement learning and object detection for robotics. He also holds a BTech in Software Development (2012) and a Master's degree in Electrical Engineering. Prior to working with machine learning, he also collaborated on research in the field of electronic educational assessments.

Thiago Pedro Donadon Homem holds a Doctorate in Electrical Engineering from the University Center FEI (2018). He is currently Associate Professor at the Federal Institute of Education, Science and Technology of São Paulo in São Paulo, Brazil. He researches in the following areas: Robotics, Artificial Intelligence and Machine Learning.

Jonas Henrique Renolfi de Oliveira holds a degree in Automation and Control Engineering from University Center FEI (2019). He is currently a student at the State University of Campinas in Campinas, Brazil. He researches in the following areas: Robotics, Artificial Intelligence and Machine Learning.

Danilo Hernani Perico received his Bachelor's degree (2009), Master's degree (2012) and PhD degree (2017) in Electrical Engineering from University Center FEI. He is currently Assistant Professor at the University Center FEI in São Bernardo do Campo, Brazil. He researches in the following areas: Robotics, Artificial Intelligence, Machine Learning, Knowledge Representation, Spatial Reasoning and Computer Vision.

Isaac Jesus da Silva received his Bachelor's degree, Master's degree and PhD degree in Electrical Engineering from University Center FEI. He is currently an Adjunct Professor at University Center FEI. He researches in the following areas: Deep Learning applied to autonomous robots, Robot Soccer, Mobile Robots, Computer Vision and Humanoid Robots Control.

Reinaldo Augusto da Costa Bianchi holds a Doctorate in Electrical Engineering from the University of São Paulo (2004). He held a internship at the Institute of Investigation en Inteligencia Artificial, Barcelona, Catalunya, Spain. He is currently Full Professor at the University Center FEI in São Bernardo do Campo, Brazil. He researches in the following areas: Robotics, Artificial Intelligence, Computer Vision, Machine Learning and Multi-Agent Systems.

Affiliations

Douglas De Rizzo Meneghetti¹  · **Thiago Pedro Donadon Homem²** · **Jonas Henrique Renolfi de Oliveira¹** · **Isaac Jesus da Silva¹** · **Danilo Hernani Perico¹** · **Reinaldo Augusto da Costa Bianchi¹**

Thiago Pedro Donadon Homem
thiagohomem@ifsp.edu.br

Jonas Henrique Renolfi de Oliveira
uniejonoliveira@fei.edu.br

Isaac Jesus da Silva
isaacjesus@fei.edu.br

Danilo Hernani Perico
dperico@fei.edu.br

Reinaldo Augusto da Costa Bianchi
rbianchi@fei.edu.br

- ¹ FEI University Center, 3972-B Humberto de Alencar Castelo Branco Ave., Assunção, São Bernardo do Campo, SP, 09850-901, Brazil
- ² Federal Institute of Education, Science and Technology of São Paulo, 951 Mutinga Ave., Jardim Santo Elias, São Paulo, SP, 05110-000, Brazil