

Object Detection in the DCT Domain: is Luminance the Solution?

Benjamin Deguerre

INSA Rouen Normandie

ACTEMIUM Paris Transport

Email: benjamin.deguerre@insa-rouen.fr

Clement Chatelain

INSA Rouen Normandie

Email: clement.chatelain@insa-rouen.fr

Gilles Gasso

INSA Rouen Normandie

Email: gilles.gasso@insa-rouen.fr

Abstract—Object detection in images has reached unprecedented performances. The state-of-the-art methods rely on deep architectures that extract salient features and predict bounding boxes enclosing the objects of interest. These methods essentially run on RGB images. However, the RGB images are often compressed by the acquisition devices for storage purpose and transfer efficiency. Hence, their decompression is required for object detectors. To gain in efficiency, this paper proposes to take advantage of the compressed representation of images to carry out object detection usable in constrained resources conditions.

Specifically, we focus on JPEG images and propose a thorough analysis of detection architectures newly designed in regard of the peculiarities of the JPEG norm. This leads to a $\times 1.7$ speed up in comparison with a standard RGB-based architecture, while only reducing the detection performance by 5.5%. Additionally, our empirical findings demonstrate that only part of the compressed JPEG information, namely the luminance component, may be required to match detection accuracy of the full input methods. Code is made available at : https://github.com/D3lt4ph4/jpeg_deep.

I. INTRODUCTION

Deep architectures, especially Convolutional Neural Networks (CNN) have become a standard for object detection. Most of the proposed architectures rely on the same two components: a feature extractor (backbone) pre-trained on a classification task and a detection head to output the box predictions. Impressive results were achieved on datasets growing in complexity such as Pascal VOC [1] or MS-COCO [2]. However, even though most of the images are compressed in order to limit the storage and the transfer bandwidth requirements, state-of-the-art detection architectures are designed for processing RGB inputs. Hence the usual procedure for any detection task observes the two following steps:

- The image is uncompressed and possibly pre-processed, namely with normalisation and/or resizing operations.
- Then the detection task is performed, typically using a deep network.

Although effective, this procedure have some drawbacks as the image decoding step induces a computational cost. Moreover, the involved deep architectures may include a tremendous amount of parameters, due to the RGB image resolution, making the network computationally expensive. These facts hinder the large scale deployment of these networks for applications with real-time and computation constraints

such as city surveillance [3] or road traffic monitoring and management [4].

In this article, we propose to directly perform object detection in compressed JPEG images, JPEG being one of the most popular and effective compression algorithms. As the JPEG compression transform images in a tiled frequency space, through Discrete Cosine Transform (DCT), it raises the question whether a detection network is able to efficiently map the frequency domain into a spatial domain in order to output the position of the objects in the original image. Detection in a JPEG compressed domain is also faced to the chrominance (C_r and C_b component) sub-sampling that is usually performed, that may lead to a resolution change when compared to the luminance component (Y component).

Our preliminary work [5] demonstrates that object detection in JPEG images is feasible. By taking advantage of the DCT blocks representation of the image to reduce the overall computation cost, we were able to speed up the inference stage by a factor of 2, at the cost of a reduced accuracy, 45% lower than the detection performances of the RGB-based architectures. Furthermore, the sub-sampled chrominance problem aforementioned was not addressed.

The present work shows that not only object detection in the compressed domain is achievable, but also that very close detection performance to those of RGB-based architectures can be reached, still with a significant speed up gain. Additionally, we investigate the use of the sole Y channel as input of the proposed detection networks, the benefit being the reduction of the required bandwidth. We empirically demonstrate that using only the Y channel is enough to reach an accuracy equivalent to the one of networks relying on YC_bC_r .

To reach these goals, we train several backbone classification networks designed to account for the specificity of the JPEG norm. With regard to the detection task, we incorporate these networks within a SSD [6], and evaluate the overall designed detector on two common detection benchmarks, Pascal VOC [1] and MS-COCO [2]. Finally, as a side contribution, we provide with a thorough analysis of the learned classification networks, in view of usage on restricted computational resources. We evaluate the classification networks on the ImageNet [7] dataset and show that networks that may be highly preferment when evaluated with powerful GPUs may behave poorly in more constrained settings.

To summarise, our contributions are threefold:

- We show that detection in the compressed domain can nearly reach detection performance of the RGB domain, and highlight its interest in embedded conditions. This is achieved by embedding adapted classification networks in a detection architecture such as the SSD [6].
- We provide a thorough evaluation of several detection models (backbone + SSD) in the compressed domain on two public datasets, namely Pascal VOC and MS-COCO.
- And finally, we experimentally show that using only the luminance is enough for detection, effectively reducing the need for bandwidth.

The rest of the paper is as follows: Section II goes over the existing literature and Section III presents the JPEG norm. The proposed approach and architectures are detailed in Section IV. Section V present the obtained results and discusses on the strength and limitation of compressed JPEG inputs.

II. RELATED WORKS

This article focuses on object detection in embedded conditions. Section II-A gives related works to object detection in RGB images, while section II-B details existing literature on the usage of the compressed representation of data coupled with deep learning.

A. Object detection

Object detection aims at detecting multiple objects within an image. This is generally achieved by finding the coordinates of bounding boxes around objects in the images. Methods for object detection are broadly based on RGB inputs and rely on two main deep architectures: two-stage detectors and one-stage detectors. The former networks are based on a region proposal stage followed by a box classification stage while the latter architectures use anchor-boxes and predict a set of pre-defined boxes as object or background.

Two-stage detectors were introduced with R-CNN [8], using Selective Search [9] as region proposal. Further, R-CNN was improved in Fast R-CNN [10] with feature computation sharing and in Faster R-CNN [11] with a Region Proposal Network. Optimising object detection and object segmentation simultaneously, Mask R-CNN [12] yielded another level of performance improvement. Trying to reduce the number of close multiple detection of the same object, Cascade R-CNN [13], proposes a cascaded detection pipeline using increasing Intersection over Union (IoU) thresholds to select the positive examples used for training. CNet [14], the current top-1 ranking approach in the MS-COCO detection challenge [2] is built on Mask R-CNN and uses cascaded classification backbone networks to improve the accuracy of the detector.

One-stage detectors were popularised by YOLO [15] and SSD [6]. While originally suffering a lower accuracy than SSD, YOLO was successively improved [16], [17] through anchor-boxes K-means selection, multi-label object class prediction or prediction across scales. RetinaNet [18] proposed the focal loss to alleviate the background imbalance problem for one-shot detectors. Using a unified module prediction

and connection between low and high prediction layers, [19] further improved the accuracy of single stage detectors. RepPoints [20], gets rid of the rectangular anchor boxes through point prediction. Using a bi-directional feature pyramid network to share information between the different scales, EfficientDet is currently the best performing one-stage detector on the MS-COCO dataset, behind CNet [14].

While two-stage detectors tend to be more accurate, one-stage detectors provide faster inference time. Because we focus on settings with limited computational resource, we opt for one-stage detectors. Of particular note, many methods were proposed to speed up the inference stage of RGB-based architectures without harming the detection performances [21] [22] [23]. The approach we propose rather takes another viewpoint as it aims at improving the network's inference speed by considering light-weight input YC_bC_r or Y . Although out of the scope of this paper, the combination of our approach with the aforementioned methods could be of interest.

B. Deep-Learning on compressed images/videos

Using compressed images has been explored in the past for various computer vision tasks. Many applications estimating flows from videos take advantage of the compression format as the encoded data often include displacement information in order to reduce the size of the videos by exploiting temporal redundancy. Wang *et al.* [24] proposed an architecture for object detection within compressed videos. They use motion vectors and residuals to infer objects through time, while only partially decoding the compressed video flux. Wu *et al.* [25] also proposed to exploit encoded motion vectors and residuals to improve both accuracy and inference speed for a task of action recognition in videos. Taking an other approach, [26] detects action by first generating an optical flow from motion vectors and residuals and then using it to classify the action. We shall mention that, while not coupled with deep learning, motion vectors and/or DCT coded frames were leveraged on to count vehicles [4] or to estimate the vehicles' speed and density on highways [27].

Besides videos, compressed representations were considered for image classification, segmentation or detection tasks. Gueguen *et al.* [28] investigated different network architectures for JPEG image classification. They reached state-of-the-art classification performance while speeding-up the prediction pass, even for their architectures requiring more FLOPs than their RGB counterparts. The gain for such architectures mainly stem from the reduced data transfer between CPU and GPU due to the image compression. In [29], image segmentation using JPEG compressed images is proposed. The methods adapts a RGB network by removing the down-sampling blocks to match the shape of the DCT inputs. The authors show impressive results almost reaching the RGB baseline with similar level of FLOPs. Finally, relying on JPEG2000 compression norm, Lahiru *et al.* [30] showed results matching the RGB baseline as well as speed improvements for classification. To do so, they stack sub bands of half decoded images and feed them to a modified neural network.

III. JPEG NORM

JPEG encoding is a lossy compression algorithm. It relies on the human eye sensibility to chroma components in an image. JPEG norm exploits the sparsity of the DCT (Discrete Cosine Transform) representation of an image to carry the compression.

The whole JPEG compression/decompression pipeline is described in figure 1. The compression procedure is as follows: first the RGB image is converted to the YC_bC_r domain. The human eye being less sensitive to the C_b and C_r components, this allows for easy compression by subsampling (usually $1/2$). Then a blockwise DCT (of size 8×8 pixels) followed by a quantization is applied. The compressed image is attained using the RLE/Huffman entropy coding.

As shown in figure 1, up until the entropy coding, the processed data is similar to an image in shape, and, for two images of the same size, the compressed representation will have the same size. Because quantization and Huffman compression are subject to variations depending on the compression level, we choose to use the de-quantized DCT coefficients as input to the models we propose (see figure 1, bottom).

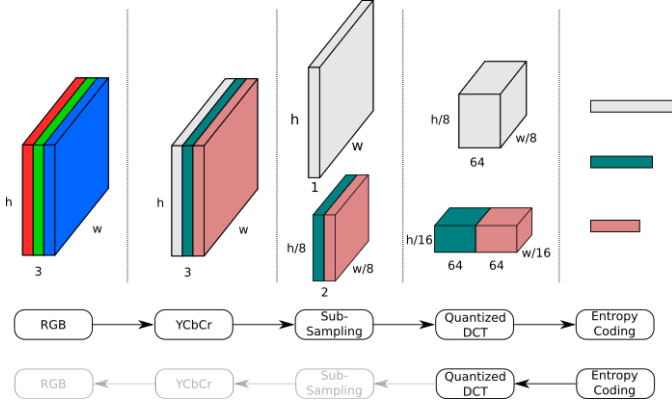


Fig. 1: JPEG CODEC pipeline. Top: obtained representation through the encoding steps, Middle: encoding steps, Bottom: decoding steps; the faded boxes show the unused steps. Y is represented in gray, C_b in dark green and C_r in faded pink.

IV. PROPOSED APPROACH

A. Detection in the frequency domain

Our goal is to design an object detection network starting from compressed JPEG images. The main challenge lies in classifying and localising an object using frequency domain information, namely the DCT block coefficients of the intended image. Indeed, contrary to RGB images, the spatial information relative to the objects may be impeded by the DCT transform. We lay our proposal on SSD [6] which was originally designed to work on RGB inputs. SSD is based on a multi-scale box prediction, each of the successive feature maps are designed to predict boxes of different sizes. In order for the SSD to work with JPEG compressed inputs, we need to modify it. As shown in Figure 2, due to the blockwise DCT,

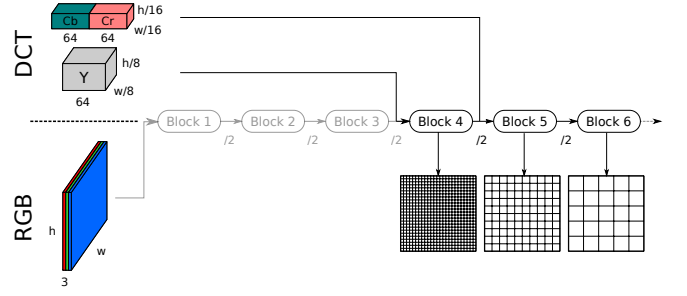


Fig. 2: Principle of the DCT-based object detector. The grids at the bottom right represent the predicted boxes by the network. The further we advance in the network, the bigger the predicted boxes. Depending on the setup, YC_bC_r or only Y inputs are fed to the networks. For clarity, not all the prediction layers are shown.

Y inputs to the network are downsized by 8 and the number of channel is augmented to 64 when compared with the original RGB image. Furthermore, due to the down-sampling on the chrominance components, we can see the C_b and C_r inputs further downsized by a factor of two. Hence, the first blocks of the genuine SSD are skipped (or replaced by blocks without down-sampling).

B. Design of the backbones, YC_bC_r input and Y input

As aforementioned, the JPEG compression algorithm requires to design carefully the detection network. Indeed, the C_b and C_r components are down-sampled, which leads to input matrices smaller by a factor of two for these components. As a result, the smallest predicted boxes are not provided with the chromatic information. This is illustrated in Figure 2. To fix this, we propose to use a deconvolution layer to scale up the C_b and C_r components to use them in combination with the Y component.

Because of the change in resolution induced by the JPEG compression algorithm, we make the hypothesis that all the components may not be useful. Intuitively, as the sub-sampling operation is carried before the blockwise DCT, the C_b and C_r 8×8 blocks are actually a sparse representation of an equivalent four Y 8×8 blocks. Hence, the learning algorithm has to deal with both sparsity and resolution problems when the C_b and C_r inputs are concatenated with the main part of the networks, potentially impeding the training. We thus propose to test a variation of the networks that only uses the Y component as input.

C. Proposed architectures

We now detail all of the tested architectures. Originally, the SSD was built on top of a VGG classification network. Besides modified VGGS for frequency domain, we also investigate ResNet50 based classification backbones as such networks have provided impressive results in terms of accuracy and inference speed for JPEG image classification [28]. The features of these detection architectures are summarised in table

I. Notice that the networks using the sole Y channel are not detailed as they are simplified instances of the YC_bC_r -based networks.

We start by reviewing the RGB architectures in IV-C1. Then in IV-C2 we detail the architectures using the compressed inputs that do not use deconvolution layers. And we finish with the deconvolution architectures in IV-C3.

1) *RGB baselines*: For each of the backbones, we compare our results with the RGB based architecture. For the VGG based SSD, we use the original architecture [6]. In order to use the ResNet50 as backbone, few modifications must be applied to the original SSD. Originally, part of the SSD convolutional layers, the fc6 and fc7 layers (c.f Table I), were designed to use the VGG dense layers' weights. As the ResNet50 does not contains such layers, we need to modify the SSD architecture to correctly incorporate the ResNet50 backbone. We remove all the original SSD layers up to (including) the fc7 layer and replace them with the ResNet50 ones (except for the last classification layer). This way, as for the VGG, only the weights from the classification layer are not pre-loaded into the SSD.

2) *YC_bC_r DCT methods*: We present the DCT architectures that do not account for the discrepancies between the Y and C_b, C_r components. Hence, for all the architectures of this subsection, the smallest boxes only rely on the Y input for predictions (c.f Figure 2). The VGG-based architecture, SSD DCT, is directly re-implemented from [5]. For the ResNet50-based architectures, we select two classification networks, Late-Concat-RFA (LC-RFA) and Late-Concat-RFA-Thinner, which have shown good precision/accuracy ratios as evidenced in [28]. RFA stands for Receptive Field Aware. The LC-RFA architectures imitate the original ResNet50 receptive field by removing the downsizing carried by some of the ResNet50 convolution layers. LC-RFA-Thinner is a lighter, hence faster, variant of LC-RFA.

a) *SSD DCT*: This method is based on the original SSD, the first three convolution blocks are removed and the Y and C_b, C_r inputs are respectively plugged in the fourth and fifth block. Due to the hard pruning of the first blocks, this is one of the fastest detection architecture proposed.

b) *SSD LC-RFA*: We integrate the LC-RFA classification network in the same way as the ResNet50 SSD RGB version. This architecture does not prune away the first convolution blocks, instead, the downsizing operations are removed by changing the stride of the convolutions from 2 to 1 when required. It is worth noticing that we had to modify part the original architecture from [28] as the one given by the authors lead to channel incompatibility on some of the layers. Specifically, we replace the last two CB_4 blocks along the Y and C_b, C_r axis with CB_3 blocks.

c) *SSD LC-RFA-Thinner*: It is a lighter version of SSD LC-RFA with a reduced number of layers [28]. This approach increases detection speed while keeping fairly good accuracy. Modifications in the number of channels is as follows: along the Y axis, channels from the three first CB blocks are set from

$\{1024, 512, 512\}$ to $\{384, 384, 768\}$ and along the C_b, C_r axis, the channels of the CB block set from $\{512\}$ to $\{256\}$.

3) *YC_bC_r DCT Deconvolution methods*: We use deconvolution methods to align the size of the down-sampled C_b, C_r to the one of Y . The related architectures are given below.

a) *SSD DCT-Deconv*: It is based on a VGG where the first three blocks are skipped. The C_b and C_r inputs first go through a deconvolution layer each and are concatenated with the Y component. Then a Batch Normalization is applied and outputs the input of the fourth block. The rest of the network follows the original SSD.

b) *SSD Deconvolution-RFA*: It is based on ResNet50. We test using the deconvolution module proposed in [28]. Contrary to SSD DCT-Deconv, the first blocks are not skipped, instead the stride of the first convolutions is changed from 2 to 1 when required. This architecture is mostly equivalent in speed and accuracy to the LC-RFA network for classification purpose.

V. EXPERIMENTS AND RESULTS

Experiments were conducted to evaluate the investigated detection networks. As a preliminary, we first implement and train the classification networks ResNet50, VGG, LC-RFA and their variants (see Section IV-C) using JPEG images (namely their block DCT parameters). For this, we use the Imagenet2012 training set [31]. The learned DCT-based networks are then evaluated on the validation set.

As for detection, we train and evaluate on two datasets, Pascal VOC and MS-COCO. All reported results on Pascal VOC are the models evaluation on 2007 test set. For the MS-COCO dataset, we use the standard train-validation-test sets.

A. Implementation details

a) *Classification*: We train the classification networks using a distributed environment and follow the recommendations from [32]. The models were trained using horovod on 4 nodes amounting to a total of 8 GPUs. Default training parameters were used for VGG and ResNet as described in [33], [34]. For data-augmentation, we rescale the images so that the smallest side is 256 pixels, we randomly crop a 224x224 patch and then randomly apply horizontal flip. The learning rate is decayed by 10 whenever the validation loss plateaued. While the original articles apply weight decay on the loss, due to framework limitation, we use a per layer one. We found out later that in this setting, the weight decay should actually be reduced by a factor of 2. Given that we get results close to the authors' baselines for the RGB networks, we keep the setting. The trained classification networks served as backbone for the detection networks.

We evaluate the number of Frames Per Seconds (FPS) that can be processed for each of the networks. We use a NVIDIA GTX 1080, set the batch size to 8 and do 10 runs of 200 predictions. The final FPS value is the average over the runs. FPS from other papers (when provided by the authors) are not directly reported as they used different GPU implementations.

TABLE I: Features of the proposed networks. A * indicates that the layer is used for boxes prediction. The lines are arbitrary and do not represent the shape of the layers. Note that LC-RFA-thinner is skipped as it is a variation of the LC-RFA model.

SSD [6]	SSD DCT [5] (reference: SSD)		SSD DCT-deconv (reference: SSD)		SSD ResNet (reference: SSD)	SSD LC-RFA (reference: SSD ResNet)		SSD Deconvolution-RFA (reference SSD ResNet)	
RGB (300,300,3)	Y (38,38,64)	Cb, Cr (19,19,128)	Y (38,38,64)	Cb, Cr (19,19,128)	RGB (300,300,3)	Y (38,38,64)	Cb, Cr (19,19,128)	Y (38,38,64)	Cb, Cr (19,19,128)
C ₁₁ , C ₁₂									
P ₁									
C ₂₁ , C ₂₂									
P ₂									
C ₃₁ , C ₃₂ , C ₃₃									
P ₃									
C ₄₁ , C ₄₂ , C ₄₃									
P ₄									
C ₅₁ , C ₅₂ , C ₅₃									
P ₅									
fc6, fc7									
C ₆₁ , C ₆₂									
C ₇₁									
C ₇₂									
C ₈₁ , C ₈₂									
C ₉₁ , C ₉₂									

Legend			
RGB	RGB pixel input	Deconv	Deconvolution with 64 output channels, filter size 2, stride 2. Separate deconvolution layers are applied to Cb and to Cr, resulting in 128 total output channels
Y	Y channel DCT input	BN	BatchNormalization
Cb, Cr	Cb and Cr channel DCT input	R	Relu
C _{ij}	J-th convolution layer of the I-th block in the original SSD architecture	CB _n	ConvBlock stage n, with number of channels as in original ResNet-50 paper, kernel size = 3 and stride = 2 unless specified otherwise.
P _i	Pooling layer of the I-th block in the original SSD architecture	IB	IdentityBlock, with number of channels matched to preceding CB layer (as in ResNet-50)
C	Convolution(channels, filter size, stride)	M	MaxPooling(pool size, stride)
Concat	Channel wise concatenation	←	Layers after this point are the same as reference
---	Channel is being concatenated	*	Layer is used for boxes prediction

b) Detection: The SSD-based detection networks were trained on a single GPU. The networks were initialised using the weights from the corresponding classification networks.

For the VGG based SSD, we follow [6] and convert the dense classification layers into convolutional layers. When converting the VGG's dense layers weights to fit the convolution layers from the SSD, we use a pre-set sub-sampling of 0:4:4096 to extract 1024 channels from the original 4096.

For the PASCAL VOC, we train on two different sets, the original 2007 training/validation and the 2007+2012 training/validation. They are respectively denoted as 07 and 07+12 in the result tables. For the MS-COCO, we use the provided training/validation datasets.

We evaluate the FPS of each detection network using the same parameters as for classification. When evaluating the speed of the networks we find the Non-Maximum Suppression to be the limiting factor. For some of the architectures, this led to a sub-optimal usage of the GPU's capacities, especially for the VGG-DCT based networks. To account for this, we report two speed evaluations: i) with one model instantiated on the GPU, and ii) with two models instantiated on the GPU. We stopped at two instances as the models were saturating the capacities of the GPU.

B. Evaluation of the classification networks

We rescale the smallest size of the images to 256 and keep the proportions. We feed them to the networks and average the predictions through a Global Average Pooling layer whenever required. For each of the networks, we also retrain on RGB images to set a baseline given our data-augmentation. Results

on the ImageNet validation dataset are reported in Table II and the Accuracy vs FPS is shown in Figure 3.

a) Training with YC_bC_r DCT inputs: We first retrain all the architecture presented in [28], namely LC-RFA, LC-RFA-Thinner and Deconvolution-RFA, and get accuracy results that are about 1~2 % lower than the original ones. We attribute these differences to small changes in the evaluation method as well as some possible differences in hyper-parameters as they were not fully disclosed. The main difference concerns the FPS of the networks, we do not reproduce the same speed improvements between the RGB and DCT architectures. The main gains are obtained for the LC-RFA-Thinner architecture when compared with RGB with a $\times 1.2$ ($\times 1.77$ in [28]) speed improvement, at equivalent accuracy. We believe these differences are due to the different GPU used for the evaluation. Our testing GPU does not process the images at a rate sufficient to take advantage of the reduce data transfer between CPU and GPU entailed by the compressed inputs.

We retrain the VGG-DCT presented in [5] with our training pipeline and show improvements of 41% on the error (23.5 points) while showing the same speed-up at $\times 2.1$ (vs $\times 2$ for [28]) when compared with RGB networks. The Deconvolution version of this network performs a bit better with an accuracy of 65.9 and improves the FPS by $\times 2.2$ when compared with RGB VGG.

If we compare the VGG based networks with the ResNet50 based ones, we can see that the VGG networks are about a 50% faster but increase the error by 9 points (34%). We attribute these differences to the hard pruning done on the first layers of the RGB architecture.

b) Training with only Y DCT input: The related classification networks are respectively denoted as VGG-DCT Y, LC-RFA Y and LC-RFA-Thinner Y. The obtained results lead to the following remarks: the accuracy slightly decreases while the networks' speed increases. The smallest decrease is for the LC-RFA Y architecture with 3.1% drop in accuracy while the biggest is for the VGG-DCT Y network with a decrease of 4.4%. This seems to be consistent with the fact that ResNet50 classifier is more accurate than the VGG. The speed improvements, ranging from 6 to 30 FPS, are due to the reduction in computation entailed by only using the Y input. While FPS gain may be negligible when comparing with the loss in accuracy, the reduced bandwidth due to the usage of the Y component only makes such architecture attractive in case of limited computation resources.

TABLE II: Classification results on ImageNet. The top panel refers to RGB-based networks, the second one corresponds to DCT-based architectures. The two last panels refer to our implementations. In bold are the best results of our trainings.

Network	top-1 accuracy	top-5 accuracy	FPS
<u>State of the Art:</u>			
VGG [33]	73.0	91.2	N/A
VGG-DCT [5]	42.0	66.9	N/A
Resnet50 [28]	75.78	92.65	N/A
LC-RFA (DCT) [28]	75.92	92.81	N/A
LC-RFA-thinner (DCT) [28]	75.39	92.57	N/A
Deconvolution-RFA (DCT) [28]	76.06	92.02	N/A
<u>our trainings (VGG based):</u>			
VGG	71.9	90.8	267
VGG-DCT	65.5	86.4	553
VGG-DCT Y	62.6	84.6	583
VGG-DCT Deconvolution	65.9	13.3	571
<u>our trainings (ResNet50 based):</u>			
Resnet50	74.73	92.33	324
LC-RFA (DCT)	74.82	92.58	318
LC-RFA Y (DCT)	73.25	91.40	329
LC-RFA-Thinner (DCT)	74.62	92.33	389
LC-RFA-Thinner Y (DCT)	72.48	91.04	395
Deconvolution-RFA (DCT)	74.55	92.39	313

C. Detection

We train the detection networks using the classification networks discussed previously. For fairness of comparison, we also retrain the RGB networks with our trained classification networks. The results for the PASCAL VOC evaluation are reported in table III and the MS-COCO results are reported in table V. We evaluate the speed of inference on the networks trained on the 07+12 PASCAL VOC dataset. Accuracy vs Speed is detailed in figure 4 and the speed results are shown in table IV.

a) Comparison of the two RGB backbones: Except when training on the Pascal VOC 2007 set, we manage to reproduce the detection results for the SSD300 architecture. Regarding the ResNet50 based architecture, it performs worse than original SSD300 on the Pascal VOC dataset but performs better on the MS-COCO dataset. These results seem to indicate that the ResNet50 based architecture has a better convergence when provided with enough data. For both of the networks, FPS are mostly comparable.




Fig. 3: Accuracy vs FPS for the classification networks. The starred networks are the ones presented in this paper.

TABLE III: Detection results on the PASCAL VOC 2007 test set, 07 is for trained on 2007 data and 07+12 means trained on 2007+2012 data. The last two panels report the performances of our trained networks. In bold are the best results of our trainings.

Network	mAP (07)	mAP (07+12)	FPS
<u>SoT:</u>			
SSD300 [6]	68.0	74.3	N/A
SSD300 DCT [5]	39.2	47.8	N/A
<u>our trainings (VGG based):</u>			
SSD300	65.0	74.0	102
SSD300 DCT	48.9	60.0	262
SSD300 DCT Y	50.7	59.8	278
SSD300 DCT Deconvolution	38.4	53.5	282
<u>our trainings (ResNet50 based):</u>			
SSD300-Resnet50 (retrained)	61.3	73.1	108
SSD300 DCT LC-RFA	61.7	70.7	110
SSD300 DCT LC-RFA Y	62.1	71.0	109
SSD300 DCT LC-RFA-Thinner	58.5	67.5	176
SSD300 DCT LC-RFA-Thinner Y	60.6	70.2	174
SSD300 DCT Deconvolution-RFA	54.7	68.8	104

TABLE IV: Speed inference of the tested detection networks. The tests were performed on a GTX 1080, "1 inst." means that only one instance of the model was loaded on the GPU for testing, "2 inst." means that two instances of the model were loaded on the GPU.

Network	FPS (1 inst.)	FPS (2 inst.)
VGG	SSD300	88
	SSD300 DCT	136
	SSD300 DCT Y	140
	SSD300 DCT Deconvolution	144
	SSD300-Resnet50	88
ResNet50	SSD300 DCT LC-RFA	87
	SSD300 DCT LC-RFA Y	91
	SSD300 DCT LC-RFA-Thinner	98
	SSD300 DCT LC-RFA-Thinner Y	101
	SSD300 DCT Deconvolution-RFA	87

TABLE V: Detection results on MS-COCO test-dev set. In bold are the best results of our trainings.

Network		Avg. Precision, IoU:			Avg. Precision, Area:			Avg. Recall, #Dets:			Avg. Recall, Area:		
		0.5:0.95	0.5	0.75	S	M	L	1	10	100	S	M	L
SSD300 [6]		23.2	41.2	23.4	5.3	23.2	39.6	22.5	33.2	35.3	9.6	37.6	56.5
VGG	SSD300 (our training)	24.5	42.4	25.2	7.8	25.3	38.0	23.0	34.0	35.7	12.3	38.1	54.4
	SSD300 DCT	14.3	27.0	13.7	2.1	12.1	26.2	15.8	22.4	23.4	3.4	21.1	42.5
	SSD300 DCT Y	14.4	27.0	14.0	2.1	12.0	26.5	15.8	22.2	23.3	3.5	20.8	42.3
	SSD300 DCT Deconvolution	13.5	26.0	12.6	2.5	11.3	23.8	15.3	21.9	23.1	4.5	21.1	39.6
ResNet50	SSD300 Resnet50	26.8	43.8	28.3	6.2	28.2	45.2	24.6	35.6	37.1	10.0	40.4	60.0
	SSD300 DCT LC-RFA	25.8	42.4	27.1	5.1	27.0	44.4	23.9	34.2	35.6	8.0	38.8	59.0
	SSD300 DCT LC-RFA-Y	25.2	41.6	26.5	5.2	25.7	43.7	23.6	33.7	35.0	8.1	37.4	58.2
	SSD300 DCT LC-RFA-Thinner	25.4	41.8	26.9	4.7	26.3	44.6	23.7	33.8	35.1	7.2	38.0	59.4
	SSD300 DCT LC-RFA-Thinner-Y	24.6	40.6	25.8	4.7	24.8	43.4	23.1	32.8	34.1	7.2	36.3	57.6
	SSD300 DCT Deconvolution-RFA	25.9	42.5	27.2	5.4	26.7	44.4	24.0	34.5	36.0	8.5	39.0	59.4




Fig. 4: mAP vs FPS for the detection networks. Networks with a star at the end of their names are the one presented in this paper.

b) Training with YC_bC_r DCT inputs: We detail all the architectures using the full compressed inputs, namely SSD300 DCT, SSD300 DCT LC-RFA and SSD300 DCT LC-RFA-Thinner. Deconvolution approaches are treated in the next section. For SSD300 DCT, we largely improve the results when compared with [5]. We relate these improvements to a better trained backbone and to the incorporation of the dense layers as pre-trained convolutional layers. On the Pascal VOC 07+12 dataset we reach 60.0 mAP and are 14.0 points behind the RGB method (18.9% decrease). On the MS-COCO dataset, we lose 10.1 points when compared with the RGB method but this time it represents a 42.2% decrease. As could be expected, the hard pruning of the network is a limitation factor on more complex dataset such as MS-COCO.

Regarding the ResNet50 based methods, SSD300 DCT LC-RFA and SSD300 DCT LC-RFA-Thinner outperform the VGG-based network by a large margin on the two datasets. They even outperform the original SSD on MS-COCO dataset. While this might be expected as the classification backbones used provide similar accuracy performances, the gap is sharper for detection. When comparing the SSD300 DCT LC-RFA with the SSD300 DCT LC-RFA-Thinner, we see that the second approach is 3 points behind the first on the Pascal

VOC datasets and about the same accuracy on MS-COCO. The main advantage of the SSD300 DCT LC-RFA-Thinner is the number of FPS it can process, $\times 1.63$ more images than the SSD300 DCT LC-RFA while maintaining an equivalent accuracy.

c) Influence of the Deconvolution on detection performance: While the deconvolution networks tend to perform better for classification we can see more mitigated results for detection. For the networks trained on the Pascal 2007 data only, we have a mAP of 38.4 for the VGG based network and 54.7 for the ResNet50 based one. They are respectively 12.3 points and 7.4 points lower than the mAP of the best performing DCT networks at equivalent speed (same backbone type, i.e VGG or ResNet). While the gap is reduced for the network trained on the 07+12 data, they still lag behind. However for the MS-COCO dataset the SSD300 Deconvolution-RFA is the best performing of all the DCT based architectures. Moreover, when looking at results by size of area (Small, Medium or Large), we can see that both of the deconvolution architectures improve the accuracy for small objects. Overall, it seems that when provided with a enough data, the network will reach accuracy level equivalent to the non-deconvolution networks.

The SSD300 DCT Deconvolution is the fastest of all the detectors with a speed of 282 FPS but with the worst overall accuracy. The SSD300 Deconvolution-RFA as a speed equivalent to the RGB networks while not performing better than the SSD300 Resnet50.

d) Evaluation of the detection networks using only the Y input: On the Pascal VOC dataset, we get similar mAP in comparison with the networks using the full YC_bC_r input. The reverse holds for the MS-COCO dataset, where performances using only the Y input tend to be lower than their full input counter-parts (1 point below). Yet, it appears that the C_b and C_r components are not critical to correctly detect objects within images. Moreover, networks for detection using only the Y component have equivalent or higher speed than the networks using the YC_bC_r inputs and require less bandwidth. The SSD300 DCT LC-RFA-Thinner Y network is $\times 1.70$ faster (with 2 instances, $\times 1.15$ when using only one instance of the network) than the original SSD while being only 3.8 point less accurate on the Pascal VOC dataset and more accurate on the MS-COCO dataset. The SSD300 DCT Y is even faster with a $\times 2.72$ speed improvement (with 2 instances, $\times 1.59$ when

using only one instance) but at the cost of 14.2 points drop in the mAP for the Pascal VOC dataset and 10.1 points for MS-COCO dataset.

e) On the speed of the networks: The FPS ratio the networks can process were computed either with one instance or with two instances of the model on one GPU. While this may seem anecdotal, we see that all the DCT network scale effortlessly when using 2 instances. The results for the network using compressed inputs indicates that they could be deployed on GPUs half as powerful as the one we used for the experiments and maintain the FPS obtained during the tests with 1 instance of the networks. This is not true for the RGB based network, which were already almost using all the computation capabilities of the GPU with only one model instantiated and thus would face important loss in regard to the FPS if deployed on GPUs half as powerful. This means that the presented architectures using the compressed inputs are good matches for usage in limited resources environment or on small remote computation devices. By combining our approach with other methods such as MobileNet [22] or TinySSD [21] we expect to even better fit such conditions.

VI. CONCLUSION

In this paper, we investigate object detection in JPEG compressed images. We devise several deep architectures based on the SSD detection network framework. The architectures we explore differ in the classification backbones they rely on. Experimental evaluations evidence that they are not all equal for detection performance. When using ResNet50 as backbone we witness a slight performance drop of 5.5% and 5.3% on the Pascal VOC and MS-COCO datasets respectively in comparison to the best RGB network. Furthermore, we demonstrate an effective speed up of $\times 1.7$ when using compressed input. Also, we demonstrate that using only the Y input leads to detection performances similar to those of networks using the YC_bC_r input. The benefit is the reduced bandwidth for image transfer. These findings are promising and may prove useful for the deployment of large real-time monitoring application.

ACKNOWLEDGMENT

This research is supported by ANRT and ACTEMIUM Paris Transport. We thank ACTEMIUM Paris Transport the funding. We thank CRIANN for the GPU computation facilities.

REFERENCES

- [1] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010.
- [2] T-Y Lin, M Maire, S J. Belongie, L D. Bourdev, R B. Girshick, J Hays, P Perona, D Ramanan, P Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014.
- [3] T Ide, T Katsuki, T Morimura, and R Morris. City-wide traffic flow estimation from a limited number of low-quality cameras. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–10, 08 2016.
- [4] Z Wang, X Liu, J Feng, J Yang, and H Xi. Compressed-domain highway vehicle counting by spatial and temporal regression. *IEEE Transactions on Circuits and Systems for Video Technology*, PP:1–1, 10 2017.
- [5] B Deguerre, C Chatelain, and G Gasso. Fast object detection in compressed JPEG images. *CoRR*, abs/1904.08408, 2019.
- [6] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. E. Reed, C.-Y. Fu, and A. C. Berg. SSD: single shot multibox detector. *CoRR*, abs/1512.02325, 2015.
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [8] R. B. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. *CoRR*, abs/1311.2524, 2013.
- [9] J.R.R. Uijlings, K.E.A. van de Sande, T. Gevers, and A.W.M. Smeulders. Selective search for object recognition. *International Journal of Computer Vision*, 2013.
- [10] R. B. Girshick. Fast R-CNN. *CoRR*, abs/1504.08083, 2015.
- [11] S. Ren, K. He, R. B. Girshick, and J. Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *CoRR*, abs/1506.01497, 2015.
- [12] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick. Mask R-CNN. *CoRR*, abs/1703.06870, 2017.
- [13] Z Cai and N Vasconcelos. Cascade R-CNN: high quality object detection and instance segmentation. *CoRR*, abs/1906.09756, 2019.
- [14] Y Liu, Y Wang, S Wang, T Liang, Q Zhao, Z Tang, and H Ling. Cbnet: A novel composite backbone network architecture for object detection. *ArXiv*, abs/1909.03625, 2019.
- [15] J. Redmon, S. Kumar Divvala, R. B. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. *CoRR*, abs/1506.02640, 2015.
- [16] J. Redmon and A. Farhadi. YOLO9000: better, faster, stronger. *CoRR*, abs/1612.08242, 2016.
- [17] J. Redmon and A. Farhadi. Yolov3: An incremental improvement. *CoRR*, abs/1804.02767, 2018.
- [18] T-Y Lin, P Goyal, R B. Girshick, K He, and P Dollár. Focal loss for dense object detection. *CoRR*, abs/1708.02002, 2017.
- [19] K Lee, J Choi, J Jeong, and N Kwak. Residual features and unified prediction network for single stage detection. 07 2017.
- [20] Z Yang, S Liu, H Hu, L Wang, and S Lin. Reppoints: Point set representation for object detection. *CoRR*, abs/1904.11490, 2019.
- [21] A. Wong, M. Javad Shafiee, F. Li, and B. Chwyl. Tiny SSD: A tiny single-shot detection deep convolutional neural network for real-time embedded object detection. *CoRR*, abs/1802.06488, 2018.
- [22] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017.
- [23] M Tan, R Pang, and Q V. Le. Efficientdet: Scalable and efficient object detection. *ArXiv*, abs/1911.09070, 2019.
- [24] S Wang, H Lu, P A. Dmitriev, and Z Deng. Fast object detection in compressed video. *CoRR*, abs/1811.11057, 2018.
- [25] C-Y Wu, M Zaheer, H Hu, R. Manmatha, A J. Smola, and P Krähenbühl. Compressed video action recognition. *CoRR*, abs/1712.00636, 2017.
- [26] Z Shou, Z Yan, Y Kalantidis, L Sevilla-Lara, M Rohrbach, X Lin, and S-F Chang. Dmc-net: Generating discriminative motion cues for fast compressed video action recognition. *CoRR*, abs/1901.03460, 2019.
- [27] X-D Yu, L-Y Duan, and Q Tian. Highway traffic information extraction from skycam mpeg video. pages 37 – 42, 02 2002.
- [28] L. Gueguen, A. Sergeev, B. Kadlec, R. Liu, and J. Yosinski. Faster neural networks straight from jpeg. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 3933–3944. Curran Associates, Inc., 2018.
- [29] S-Y Lo and H-M Hang. Exploring semantic segmentation on the DCT representation. *CoRR*, abs/1907.10015, 2019.
- [30] L D. Chamain and Z Ding. Faster and accurate classification for jpeg2000 compressed images in networked applications. *ArXiv*, abs/1909.05638, 2019.
- [31] O Russakovsky, J Deng, H Su, J Krause, S Satheesh, S Ma, Z Huang, A Karpathy, A Khosla, M Bernstein, A C. Berg, and L Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015.
- [32] P Goyal, P Dollár, R B. Girshick, P Noordhuis, L Wesolowski, A Kyrola, A Tulloch, Y Jia, and K He. Accurate, large minibatch SGD: training imagenet in 1 hour. *CoRR*, abs/1706.02677, 2017.
- [33] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [34] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE CVPR*, pages 770–778, June 2016.