

SCUOLA INTERNAZIONALE SUPERIORE DI STUDI AVANZATI

SISSA Digital Library

A Proper Orthogonal Decomposition Approach for Parameters Reduction of Single Shot Detector Networks

Original

A Proper Orthogonal Decomposition Approach for Parameters Reduction of Single Shot Detector Networks / Meneghetti, L.; Demo, N.; Rozza, G. - (2022). (Intervento presentato al convegno IEEE International Conference on Image Processing tenutosi a Bourdeaux, France nel 16-19 OCTOBER 2022) [10.1109/ICIP46576.2022.9897513].

Availability:

This version is available at: 20.500.11767/129551 since: 2022-09-21T13:32:56Z

Publisher: IEEE

Published DOI:10.1109/ICIP46576.2022.9897513

Terms of use:

Testo definito dall'ateneo relativo alle clausole di concessione d'uso

Publisher copyright Springer

This version is available for education and non-commercial purposes.

note finali coverpage

(Article begins on next page)

A Proper Orthogonal Decomposition approach for parameters reduction of Single Shot Detector networks

Laura Meneghetti^{*}, Nicola Demo[†], and Gianluigi Rozza[‡]

Mathematics Area, mathLab, SISSA, via Bonomea 265, I-34136 Trieste, Italy

July 28, 2022

Abstract

As a major breakthrough in artificial intelligence and deep learning, Convolutional Neural Networks have achieved an impressive success in solving many problems in several fields including computer vision and image processing. Real-time performance, robustness of algorithms and fast training processes remain open problems in these contexts. In addition object recognition and detection are challenging tasks for resource-constrained embedded systems, commonly used in the industrial sector. To overcome these issues, we propose a dimensionality reduction framework based on Proper Orthogonal Decomposition, a classical model order reduction technique, in order to gain a reduction in the number of hyperparameters of the net. We have applied such framework to SSD300 architecture using PASCAL VOC dataset, demonstrating a reduction of the network dimension and a remarkable speedup in the fine-tuning of the network in a transfer learning context.

1 Introduction

Computer vision has attracted much research interest in recent years due to a growing attention to the problems of video analysis and image understanding. The task of estimating the location and the class of objects contained within an image is referred as *object detection* [3, 29, 32, 14]. In such procedure, it is not only required to classify objects contained in pictures but also to detect their position. This additional requirement to find all instances of objects leads to more complex and deep networks than traditional Convolutional Neural Networks (CNNs) [26, 34, 14], such as Faster R-CNN [20], SSD [15], YOLO [18, 19]. In fact net architectures for object detection are typically composed of a base net, i.e. a CNN, and some additional structures responsible for the classification and localization of the different entities in images by drawing a labeled bounding box around each of them.

Training these high-dimensional deep neural networks, characterized by millions of parameters, has implications when running models in practice, especially for the possible long training time [7, 28] and for the required storage space. There exists object detector's implementation designed for memory constrained systems maintaining their original image classifier, but based on well-established hardware-efficient CNNs as backbone, e.g. MobileNet [10], ShuffleNet [33], SqueezeNet [11], producing thus a light-weight deep neural network with reasonable performances [24, 10, 17, 27, 13, 31, 12]. Other approaches directly adopt the original object detection architecture employing input resizing and network pruning [8, 6, 1], or highly optimized base nets and convolutional feature layers [30] to lower network complexity.

Unlike what done in the aforementioned methods, to face the compute-intensive and memoryintensive issues a dimensionality reduction method is proposed, extending and adapting the one proposed for Artificial and Convolutional Neural Networks [16, 2]. Starting from the original structure of an object detector, the method performs a compression on the network by reducing the base net. This dimensionality reduction is carried out using Proper Orthogonal Decomposition

^{*}laura.meneghetti@sissa.it

[†]nicola.demo@sissa.it

 $^{^{\}ddagger}$ gianluigi.rozza@sissa.it

(POD) [9], a well established method for model order reduction [21, 22, 23]. Such method allows for the computation of the optimal space to linearly represent the input data, here the inner features of a certain inner layer, as well as an estimation of the error introduced by projecting onto this space. We are thus retaining only a certain number of layers of the backbone and substituting the remaining part with this reduction layer, decreasing the number of hyperparameters in order to obtain benefits both in the space needed to store the architecture and in fine tuning time. It must be said that the amount of layers — and so of information — we discard controls the final accuracy and final dimension of the network.

The article is organized as follows: Section 2 provides a detailed description of the reduced framework for object detectors, while in Section 3 we present the results obtained by applying the proposed methodology to a benchmark deep neural network designed for the problem of object detection. Finally Section 4 summarizes the entire procedure and gives ideas for future developments and perspectives for our framework.

2 Reduced Object Detector

In this section we provide the detailed description of the proposed reduced technique, that is summarized in Figure 1 and in Algorithm 1. This framework is derived by the one proposed in [16] to include also more complex networks, as the ones that solve the problem of object detection [26, 34, 14]. The only assumption we are making is that we have a network composed by: i,) a base net, a convolutional neural network extracting the low-level features; ii,) some additional convolutional layers responsible for capturing the high-level features; iii,) a predictor, that will output the predicted class and localization, e.g. the coordinates of the bounding box, for each object in the image.



Figure 1: Graphical representation of the reduction method proposed for an object detector.

2.1 Network Splitting

At the beginning, as done in [16, 2], the first part of $Obj_{\mathcal{D}et}$, called *base net*, is splitted in two different blocks: the *pre-model* and the *post-model*. Therefore, describing the base net as compositions of L functions $f_j : \mathbb{R}^{n_{j-1}} \to \mathbb{R}^{n_j}$, for $j = 1, \ldots, L$, representing the different layers of the network (e.g. convolutional with ReLU, fully connected, pooling layers), it followed that for any $1 \leq l < L$:

$$basenet(\mathbf{x}^{0}) \equiv basenet_{\text{post}}^{l}(base_net_{\text{pre}}^{l}(\mathbf{x}^{0})), \tag{1}$$

where $\mathbf{x}^0 \in \mathbb{R}^{n_0}$ is an input image and the pre- and post-models are defined by:

$$basenet_{\text{pre}}^{l} = f_{l} \circ f_{l-1} \circ \dots \circ f_{1},$$

$$basenet_{\text{post}}^{l} = f_{L} \circ f_{L-1} \circ \dots \circ f_{l+1}.$$
(2)

Algorithm 1 Pseudo-code for the construction of the reduced object detector. Inputs:

- train dataset $\mathcal{D}_{\text{train}} = \{\mathbf{x}^{(0),j}, \mathbf{y}_{\text{loc}}^j, \mathbf{y}_{\text{cls}}^j\}_{j=1}^{N_{\text{train}}},$
- Obj_Det = [basenet, auxlayers, predictor],
- reduced dimension r,
- index of the cut-off layer l,
- a test dataset $\mathcal{D}_{\text{test}} = \{\mathbf{x}^i, \mathbf{y}_{\text{loc}}^i, \mathbf{y}_{\text{cls}}^i\}_{i=1}^{N_{\text{test}}}.$

Output: Reduced Object Detector Obj_Det^{red}

basenet^l_{pre}, basenet^l_{post} = splitting_net(basenet, l)
 x^(l) = basenet^l_{pre}(x⁽⁰⁾)
 z = reduce(x^(l), r)
 ŷ_{loc}, ŷ_{cls} = predictor(x^(l), z)
 Training of the constructed reduced net using D_{test}.

Then, the index l defines the *cut-off layer*, i.e. the layer at which the base net is cut. It is thus controlling how many layers of the original net, i.e. information, we are discarding. Since its important role in the final outcome, it should be chosen carefully, usually based on considerations about the network and the dataset at hand, balancing the final accuracy and the compression ratio.

2.2 Dimensionality Reduction

The base net is usually followed by an additional structure responsible for the detection of the high-level features of objects, giving thus a whole understanding of the picture. In our technique, we are substituting the auxiliary layers of Ob_j_Det with a linear layer performing dimensionality reduction on the pre-model output $\mathbf{x}^{(l)}$, as described in algorithm 1. In particular since $\mathbf{x}^{(l)}$ usually lies in a high dimensional space, we aim at projecting it onto a low-dimensional one by retaining only the most important directions and thus informations.

Based on the method presented in [16], we are exploiting the Proper Orthogonal Decomposition (POD) approach of reduce order modeling [9] in order to decrease the number of parameters. Therefore, given the pre-model outputs for each element of the training dataset $\{\mathbf{x}^{(l),i}\}_{i=1}^{N_{\text{train}}}$, we unroll them as columns of a matrix $\mathbf{S} = [\mathbf{x}^{(l),1}, \dots, \mathbf{x}^{(l),N_{\text{train}}}]$ which we decompose by applying the singular value decomposition, such that:

$$\mathbf{S} = \boldsymbol{\Psi} \boldsymbol{\Sigma} \boldsymbol{\Theta}^T, \tag{3}$$

where the left-singular vectors, i.e. the columns of the unitary matrix Ψ , are the POD modes, and the diagonal matrix Σ contains the corresponding singular values in decreasing order. The latter indicates the contribution of all the POD modes to represent the input data, allowing for a priori selection of the number of dimension of such space. The reduced solution \mathbf{z} is then obtained by applying to the full solution $\mathbf{x}^{(l)}$ the projection matrix Ψ_r , constructed by selecting the first rmodes:

$$\mathbf{z}^{i} = \mathbf{\Psi}_{r}^{T} \mathbf{x}^{(l),i}, \quad \text{for } i = 1, \dots, N_{\text{train}}.$$
 (4)

After this reduction layer, we have then in hand the reduced representation of the pre-model output for any image in our dataset.

2.3 Predictor

To obtain the reduced network we finally need to connect the reduced features $\{\mathbf{z}^i\}_{i=1}^{N_{\text{train}}}$ to the expected outputs $\{\mathbf{y}_{\text{loc}}^i, \mathbf{y}_{\text{cls}}^i\}_{i=1}^{N_{\text{train}}}$. This mapping is realized by using the original classifier: several convolutional layers for localization prediction and as many for class prediction. Regarding the structural differences, we highlight that, typically, architectures for object detection use the output of several layers — e.g. base net convolutional layers, auxiliary layers — to predict the final output. In the reduced network, instead, the only inputs are represented by the output of the pre-model $\mathbf{x}^{(l)}$ and its reduced version \mathbf{z} obtained through the reduction layer.

This last block of the object detector is then characterized by the presence of *priors* or *anchor* boxes [15, 20]: fixed size reference boxes which are placed uniformly throughout the original image

Table 1: Results obtained with the cat-dog dataset.

Network	mAP	Storage (Mb)	Training Time
SSD300	70.2%	91.09	43.5 h
$SSD300_red$	59%	77.45	26 h

in order to find the correct place and sizes (width and height) of the predicted bounding boxes. In order to discretize efficiently the space of possible predicted box length, several priors with different shapes are used for each selected input. In particular larger feature maps — e.g. base net convolutional layers output — have priors with smaller resolutions and are therefore ideal for detecting minor details of the image, whereas smaller feature maps — e.g. auxiliary layers output — will be responsible of detecting high level features, like objects' shapes. Hence, since each predictor's input is allowed to have a certain number of priors, we have performed an empirical analysis in order to find the right scale parameter, that controls the priors' resolution, to gain comparable results with the original network.

3 Results

We conduct experiments using SSD300 [15] with VGG-16 [25] as base net and the PASCAL Visual Object Classes (VOC) [5] dataset, to evaluate the ability of the full and reduced nets to detect objects in pictures. In particular, we have used both the complete PASCAL VOC and a reduced version, extracted from VOC2007, made up of 300 images subdivided in two categories: cats and dogs.

First of all, we have trained and tested the original net SSD300 with the cat-dog dataset. The results obtained are displayed in the first column of figure 2 and in table 1, where it can be seen that, after a 500 epochs of training, the net has an overall *mean Average Precision* (mAP) of 70.2%. In particular it reaches an accuracy of 86.6% for dogs and 53.8% for cats. We have then applied



(c) SSD300



(b) Reduced SSD300



(d) Reduced SSD300

Figure 2: Comparison of the results obtained using the original SSD300 and its reduced version on two test images.

the reduction method described in section 2 to the original architecture of SSD300, where we have not used the status of the network after training, but after initialization, i.e. with all the weights randomly initialized except those of the VGG-16, pre-trained on ImageNet [4]. Hence, the net has been cut at layer 16, i.e. cut-off index 7, of VGG-16 and then the output of the pre-model has been projected onto a reduced space of dimension r = 50, in analogy with [16, 2]. As highlighted in section 2.3 the inputs to the predictor part are different from the full model. In the case of SSD300 these correspond to: the output of the third layer in the fourth convolutional block conv4_3, the

Table 2: Results obtained with PASCAL VOC.

Network	mAP	Storage (Mb)	Training Time
SSD300	77.8%	100.23	48 h
SSD300_red	39%	76.23	18 h

final output of the net conv7, the outputs of each block of the auxiliary layers conv8_2, conv9_2, conv10_2, conv11_2. In our reduced network the only inputs used for predictions are the output of the pre-model, i.e. the output of layer conv3_3, and its reduced version, the output of the reduction layer. In this way, the number of priors we are taking into account is less than before: 5782 instead of 8732 for the full net. We have then adapted the scaling factor to our case, choosing two different scales: 0.1 and 0.9, i.e. one that takes into account smaller objects (10% of the picture) and one for bigger objects (90% of the picture). Despite these changes connected with inputs and default boxes, the predictor block remains the same as the original SSD300.

The reduced net has firstly been trained with the cat-dog dataset for 500 epochs. Table 1 summarizes the results obtained with our reduced method. As can be seen, in this case the mAP is decreased by 11% with respect to the original net: the accuracy for the categories dogs and cats are now 67.7% and 50.3% respectively. Figure 2 shows then how the results obtained with the full and the reduced net are comparable.

We have then performed training using all the 20 categories of PASCAL VOC. We have trained the original SSD300 and our reduced version for 232 epochs¹. As described in Table 2, in this case the lightweight version is not as accurate as the full network. In fact, there is an overall reduction in the mAP of 38%.

Despite this decreased precision, we have achieved a reduction of memory storage, 15% in the first case and 22% in the second one, and a halving of the training time. We highlight that such compression is not remarkable as the one shown in [16] since in this case the larger dimension of the pre-model output have led a huge matrix for storing the POD modes, limiting a lot the potential space gain. However, beside the space reduction, our proposed approach has allowed to accelerate the fine-tuning (performed also in the full network to optimize the auxiliary and classifier layers) of the network in the reduced version: assuming the weights of base net layers are already pre-trained², the minor number of hyperparameters results in a faster optimization task. Reduction of CNN becomes then not only a technique to compress the architecture dimension, but also accelerating the learning step of the studied network at the cost of marginally decreased accuracy.

4 Conclusion

We have applied an already explored methodology for neural network compression to one of the most common architecture for object detection, the SSD300 model. Without loosing of generality, the approach consists of a (linear) reduction layer derived from POD used to connect the selected layers of the base net (pre-model) with the original predictor block. The comparison between the original network and its reduced version shows a faster training and less required space for the latter, but at the cost of a diminished accuracy. The obtained results shows an intuitive trend between the reduction and the accuracy with respect to the complexity of the problem, in this case roughly quantified by the number of image classes. Keeping indeed the same reduced architecture for the two different datasets (2 and 20 classes), we have seen a drastic difference in the average accuracy, that demonstrates the ability of the proposed method to automatically detect redundant and superfluous information. Future works will better explore the trade-off between complexity of the problem and compression of the network, aiming hopefully to detect an automatic way to compute the cutoff layer depending on the problem in hand.

Regarding the impact of the reduced network in memory instead, this application has demonstrated that, dealing with high-dimensional features in the pre-model, can vanishing the compression of the network. The POD modes, which dimension is equal to pre-model features, require a huge amount of space to be stored, potentially reaching the full network as in this case. One

 $^{^1\}mathrm{We}$ have reduced the number of epochs for computational limits.

 $^{^2\}mathrm{Mandatory}$ for building the POD space.

possible solution to overcome this limit is the employment of hyperreduction technique or POD variants to minimize the quantity of data that has to be saved.

References

- M. Bonnaerens, M. Freiberger, and J. Dambre. Anchor pruning for object detection. arXiv preprint arXiv:2104.00432, 2021.
- [2] C. Cui, K. Zhang, T. Daulbaev, J. Gusak, I. Oseledets, and Z. Zhang. Active subspace of neural networks: Structural analysis and universal attacks. SIAM Journal on Mathematics of Data Science, 2(4):1096–1122, 2020.
- [3] B. Cyganek. Object detection and recognition in digital images: theory and practice. John Wiley & Sons, 2013.
- [4] 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.
- [5] 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.
- [6] S. Ghosh, S. K. K. Srinivasa, P. Amon, A. Hutter, and A. Kaup. Deep network pruning for object detection. In 2019 IEEE International Conference on Image Processing (ICIP), pages 3915–3919, 2019.
- [7] I. Goodfellow, Y. Bengio, and A. Courville. Deep Learning. MIT Press, 2016.
- [8] S. Han, H. Mao, and W. J. Dally. Deep Compression: Compressing Deep Neural Network with Pruning, Trained Quantization and Huffman Coding. In Y. Bengio and Y. LeCun, editors, *ICLR 2016, Conference Track Proceedings*, 2016.
- [9] J. S. Hesthaven, G. Rozza, and B. Stamm. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*. Springer Briefs in Mathematics. Springer, 1 edition, 2015.
- [10] 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. arXiv preprint arXiv:1704.04861, 2017.
- [11] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and; 0.5 MB model size. arXiv preprint arXiv:1602.07360, 2016.
- [12] L. Jiang, W. Nie, J. Zhu, X. Gao, and B. Lei. Lightweight object detection network model suitable for indoor mobile robots. *Journal of Mechanical Science and Technology*, 2022.
- [13] Y. Li and C. Lv. Ss-yolo: An object detection algorithm based on yolov3 and shufflenet. In 2020 IEEE 4th Information Technology, Networking, Electronic and Automation Control Conference (ITNEC), volume 1, pages 769–772, 2020.
- [14] L. Liu, W. Ouyang, X. Wang, P. Fieguth, J. Chen, X. Liu, and M. Pietikäinen. Deep learning for generic object detection: A survey. *IJCV*, 128(2):261–318, 2020.
- [15] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. SSD: Single Shot MultiBox Detector. In *Computer Vision – ECCV 2016*, pages 21–37. Springer International Publishing, 2016.
- [16] L. Meneghetti, N. Demo, and G. Rozza. A Dimensionality Reduction Approach for Convolutional Neural Networks. arXiv preprint arXiv:2110.09163, 2021.
- [17] T. Mudumbi, N. Bian, Y. Zhang, and F. Hazoume. An approach combined the faster RCNN and mobilenet for logo detection. *Journal of Physics: Conference Series*, 1284(1), 2019.
- [18] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, realtime object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [19] J. Redmon and A. Farhadi. Yolo9000: better, faster, stronger. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 7263–7271, 2017.
- [20] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Proceedings of NIPS'15*, page 91–99. MIT Press, 2015.
- [21] G. Rozza, M. Hess, G. Stabile, M. Tezzele, and F. Ballarin. chapter Basic Ideas and Tools for Projection-Based Model Reduction of Parametric Partial Differential Equations. 2020.
- [22] G. Rozza, M. Malik, N. Demo, M. Tezzele, M. Girfoglio, G. Stabile, and A. Mola. Advances in reduced order methods for parametric industrial problems in computational fluid dynamics. In *Proceedings of the ECCM 2018 and ECFD 2018*, pages 59–76, 2020.

- [23] F. Salmoiraghi, F. Ballarin, G. Corsi, A. Mola, M. Tezzele, and G. Rozza. Advances in geometrical parametrization and reduced order models and methods for computational fluid dynamics problems in applied sciences and engineering: overview and perspectives. In M. Papadrakakis, V. Papadopoulos, G. Stefanou, and V. Plevris, editors, *Proceedings of the ECCO-MAS Congress 2016*, 2016.
- [24] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4510–4520, 2018.
- [25] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *Proceedings 3rd ICLR*, 2015.
- [26] F. Sultana, A. Sufian, and P. Dutta. A Review of Object Detection Models Based on Convolutional Neural Network, pages 1–16. 2020.
- [27] Y. Sun, C. Wang, and L. Qu. An object detection network for embedded system. In 2019 IEEE International Conferences on Ubiquitous Computing Communications (IUCC) and Data Science and Computational Intelligence (DSCI) and Smart Computing, Networking and Services (SmartCNS), pages 506–512, 2019.
- [28] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE CVPR*, pages 1–9, 2015.
- [29] X. Wang. Deep learning in object recognition, detection, and segmentation. Foundations and Trends in Signal Processing, 8(4):217–382, 2016.
- [30] A. Womg, M. J. Shafiee, F. Li, and B. Chwyl. Tiny SSD: A Tiny Single-Shot Detection Deep Convolutional Neural Network for Real-Time Embedded Object Detection. pages 95–101, 05 2018.
- [31] B. Wu, F. Iandola, P. Jin, and K. Keutzer. SqueezeDet: Unified, Small, Low Power Fully Convolutional Neural Networks for Real-Time Object Detection for Autonomous Driving. pages 446–454, 07 2017.
- [32] X. Zhang, Y.-H. Yang, Z. Han, H. Wang, and C. Gao. Object class detection: A survey. ACM Comput. Surv., 46(1), jul 2013.
- [33] X. Zhang, X. Zhou, M. Lin, and J. Sun. ShuffleNet: An Extremely Efficient Convolutional Neural Network for Mobile Devices. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 6848–6856, 2018.
- [34] Z.-Q. Zhao, P. Zheng, S.-T. Xu, and X. Wu. Object detection with deep learning: A review. IEEE transactions on neural networks and learning systems, 30(11):3212—3232, 2019.