Jet Single Shot Detection

Adrian Alan Pol^{1*}, Thea Aarrestad¹, Katya Govorkova¹, Roi Halily⁴, Tal Kopetz⁴, Anat Klempner⁴, Vladimir Loncar^{1,3}, Jennifer Ngadiuba², Maurizio Pierini¹, Olya Sirkin⁴, and Sioni Summers¹

¹European Organization for Nuclear Research (CERN), Geneva, Switzerland

²California Institute of Technology, Pasadena, USA

³Institute of Physics Belgrade, Belgrade, Serbia

⁴CEVA, Herzliya, Israel

Abstract. We apply object detection techniques based on Convolutional Neural Networks to jet reconstruction and identification at the CERN Large Hadron Collider. In particular, we focus on CaloJet reconstruction, representing each event as an image composed of calorimeter cells and using a Single Shot Detection network, called Jet-SSD. The model performs simultaneous localization and classification and additional regression tasks to measure jet features. We investigate Ternary Weight Networks with weights constrained to $\{-1, 0, 1\}$ times a layer- and channel-dependent scaling factors. We show that the quantized version of the network closely matches the performance of its full-precision equivalent.

1 Introduction

The majority of particles produced at the CERN Large Hadron Collider (LHC) are unstable and immediately decay in different particles. When quarks and gluons are produced, QCD confinement prevents them from travelling across the detector. Instead, they shower other quarks and gluons, eventually hadronizing into particles. The result of this process is a *jet*, a collimated showers of particles with adjacent trajectories. Jets are key in many physics analyses done on the data collected by the LHC experiments, e.g. [1–4]. The procedure of classifying the origin of these jets, i.e. the nature of the particle that initiated the shower, known as *jet tagging* [5–8] is a fundamental task for collision reconstruction at the LHC. Similarly, it is important to determine the jet energy, momentum, and mass.

Traditional approaches to jet tagging rely on features designed by experts that detect characteristic energy deposition patterns [9–17]. In recent years, several studies projected the lower level detector measurements of the emanating particles into an image, known as *jet images*. This opened the path to applying computer vision and machine learning techniques [18–30], with particular attention to Convolutional Neural Networks (CNNs) [31].

The goal of this paper is to extend this approach to the problem of jet clustering, e.g., to replace FastJet [32] on computing architectures where parallel computing is more adequate. At the same time, we aim at demonstrating that jet clustering, mass measurement, and tagging could all be handled simultaneously. Besides the practical advantages of a single-shot

*e-mail: adrianalan.pol@cern.ch

Figure 1. The architecture of the SSD network, proposed by [49].

approach to jet reconstruction, one would benefit from mutual learning when accomplishing more tasks at once. For instance, a classifier and a regression running at once can learn that calibration constants depend on the nature of the jet, an issue that is not handled with ad-hoc post-processing (see [33] as an example).

With the luminosity increase expected in the future, traditional reconstruction algorithms might suffer from execution time scaling worse than linearly with the number of collisions happening in one bunch crossing. For this reason, it is worth investigating solutions that could execute many tasks at once, while retaining accuracy and benefiting from the additional speed up offered by parallel computing architectures. Deep neural networks, such as those used for computing vision tasks, are an obvious candidate.

On the other hand, memory consumption is also an important aspect to keep under control. To this purpose, we investigate the use of extreme quantization, up to ternary precision, which is applied already at training time to retain accuracy.

The remainder of this paper is structured as follows. In Sections 2 and 3 we briefly review single-shot detection and efficient model design techniques. In Section 4 we introduce the dataset and in Section 5 model architecture, implementation details and training procedure. Finally, in Section 6 we present the evaluation metric and results.

2 Single-shot object detection

Object detection is a fundamental task in computer vision. It is defined as the classification of objects from predefined categories in the image along with their precise spatial locations. The spatial location and extent of an object can be defined coarsely using a bounding box, which is an axis-aligned rectangle tightly bounding the object. Instead, a precise pixel-wise segmentation mask corresponds to the segmentation task.

Starting from Overfeat Network [34], the field of object detection focused on using primarily CNNs as a building block, achieving state-of-the-art results in tasks such as face [35] or pedestrian detection [36]. For a general survey on this subject, see [37, 38].

The deep learning-based object detection models are divided into two groups: one [39–43] or two [44–48] stage detectors. Two-stage detectors tend to achieve better accuracy, while one-stage detectors are simpler and faster, hence more suitable to online tasks.

The Single Shot Multibox Detector (SSD) [49], shown in Figure 1, is a simple one-stage, anchor-based detector. First, a set of default regions in an image with a fixed shape and size is predefined to discretize the output space of bounding boxes, called anchors. These anchors have a diverse set of shapes to detect objects with different dimensions, i.e multiple scales and aspect ratios. At each location, the same amount of anchors is defined. Based on the

ground truth, the object locations are matched with the most appropriate anchors to obtain the supervision signal for the anchor estimation.

During training, each anchor is refined by four box coordinates offsets (width, height, x and y) optimized by localization loss (a smooth L1 loss) and predict the categorical probabilities (including background), optimized by classification loss (categorical cross-entropy). To avoid a huge number of negative proposals dominating training gradients, hard negative mining is used to train the network, which fixes the foreground and background ratio.

The SSD architecture is fully convolutional, with initial layers based on a pre-trained backbone architecture, such as VGG-16 [50], followed by extra convolutional layers, progressively decreasing in size. The information in the last layer may be too coarse spatially to allow precise localization and at the same time, detecting large objects in shallow layers is non-optimal without large enough receptive fields. SSD performs detection over multiple scales by operating on multiple feature maps, i.e. at different depths of the network. Each of these feature maps is responsible for detecting objects according to their receptive field.

The final prediction is made by merging all detection results from different feature maps followed by a non-maximum suppression (NMS) step to produce the final detection. NMS removes duplicate predictions originating from multiple anchors.

3 Efficient inference

Network compression [51] is a common technique to reduce the number of operations, model size, energy consumption, and over-training of deep neural networks. As neural network synapses and neurons can be redundant, compression techniques attempt to reduce the total number of them, effectively reducing multipliers. Several approaches have been successfully deployed without much loss in accuracy, including parameter pruning [52–54] (selective removal of parameters based on a particular ranking and regularization), low-rank factorisation [55–57] (using matrix decomposition to estimate informative parameters), compact network architectures [58–61], and knowledge distillation [62] (training a compact network with distilled knowledge of a large network).

A particularly successful compression technique is weight quantization [63–71], which is reducing the precision of operations and operands. It has been observed that 32-bit floating-point calculations or full-precision (FP) are not needed at inference to achieve optimal performance. Thus, reducing the precision of the calculations, i.e. weights and biases, has little impact on performance compared to speed up and resource usage. This includes moving away from floating point to fixed point, reducing bit-width and weight sharing. An example of a very aggressive strategy is reducing weight precision to ternary values restricted to $\{-1, 0, 1\}$ only, called Ternary Weight Network (TWN) [68]. The quantization is performed during training, using a straight-through estimator [63], where ternary weights are used during the forward and backward propagation but not during the parameters update. To make the network perform well, TWNs minimize the Euclidian distance between full precision weights and the ternary ones with the use of a non-negative layer- and channel-dependent scaling factor α .

4 Dataset

The CERN LHC experiments implement a real-time selection process, called trigger [72], to store a fraction of the events for further analysis. Jets are useful for many measurements and physics searches. A truly minimal approach to perform identification and tagging is with jet images. Generally, jets need a component of tracks as well to be properly reconstructed.

However, one could reconstruct the calorimeter part alone (known as CaloJet). The energy measurements of the emanating particles can be projected onto a cylindrical detector and represented as images by unfolding the inner surface of the calorimeter on a rectangle, and using the crystals as pixels, as in [73].

The detector effects and hadronization have an important effect on the jet substructure. In this work, we use an emulation of the Compact Muon Solenoid (CMS) apparatus as a reference. There are two calorimeters within the solenoid volume of the CMS detector. A lead tungstate crystal Electromagnetic Calorimeter (ECAL) is designed to stop particles whose main interaction is electromagnetic (photons, electrons). A brass and scintillator Hadronic Calorimeter (HCAL) is designed to stop hadrons. They give a measurement of the energy of particles (charged and neutrals). Each of them is composed of a barrel and two endcap sections. Forward calorimeters extend the pseudorapidity range (η) coverage provided by the barrel ($\eta \leq 1.4$) and endcap detectors ($1.4 < |\eta| \leq 3.0$). A more detailed description of the CMS detector, together with a definition of the coordinate system used and the relevant kinematic variables, can be found in [74].

This study aims at identifying different kinds of jets. To this purpose, we consider 13 TeV proton-proton collision events, in which RS gravitons decay to $b\bar{b}$, HH, WW, ZZ, or $t\bar{t}$ final states. Events are generated with Pythia [75] and the CMS detector effects are emulated using the Delphes [76] library. In addition to the hard collision, parasitic *pileup* collisions are also simulated, overlapping minimum bias events. The number of pileup collisions is sampled from a Poisson distribution. The calorimeter cells (towers) in the barrel region are arranged in a fixed discrete space with fine segmentation in η , ϕ , where ϕ is the translated azimuthal angle. The final image is formed by translating the calorimeter energy deposits into pixels, which results in a 340×360 pixel image. The intensity of each pixel is proportional to the sum of the energy of the corresponding cell. The previous studies on jet images implemented data pre-processing steps such as translation, rotation, re-pixelation, or inversion. However, in our study we only limit the input to barrel and endcap section, $\eta \in (-3, 3)$, and normalize pixel intensities to a fixed range $[0, 1]$, using maximum scaling. The ground truth labels for jets above threshold momentum (30 GeV/c for b and 200 GeV/c for the jets from boosted heavy particles) are obtained using a simple cone algorithm, i.e. associating together particles whose trajectories lie within a circle of radius $R = 0.4$ from the jet centre.

As a proof of concept, we investigate the tagging of the bottom (b) W boson (W), Higgs boson (H), or top quark (t) jet. An example input, energy deposits translated to two-dimensional images with two channels (corresponding to ECAL and HCAL) together with marked ground truth bounding boxes is shown in Figure 2.

5 Model, implementation and training procedure

The Jet-SSD architecture is shown in Figure 3. Several modifications are applied to the original architecture [49]. Due to target hardware constraints, all filters in convolution layers are of size 3×3 with no dilatation and all pooling layers have 2×2 filters. Each convolution block is followed by batch normalization [77, 78] and parametric rectified linear unit (PReLU) layers. To compress the model we use half of the channels of the VGG-16 in each layer. We also remove bias from all convolution layers. The extra layers proposed by the original paper do not contribute to accurate detection due to the size of jets and thus they are removed at the training. Retaining the deeper layers in the base network does not show improvements in the final detection results either, but they are critical during training due to additional signal during back-propagation. Hence, we only purge them at inference.

The Jet-SSD network is implemented on an NVidia Tesla GPU using PyTorch [79]. For training, we use stochastic gradient descent with an initial learning rate of 10^{-3} with momen-

Figure 2. Energy deposits in CMS ECAL and HCAL translated to a two-dimensional image, an example input to the SSD network. The red bounding boxes correspond to ground truth with target label and momentum.

Figure 3. Jet-SSD architecture: input (in blue), convolution block, i.e. convolution layer followed by batch normalization and PReLU activation (in yellow), average pooling (in red) and output layer (in green). The numbers indicate the number of output channels in each block. The part of the network highlighted by dashed lines is used only during the training step.

tum set to 0.9 and weight regularization to 0.0005. We train the network for 100 epochs with a batch size of 25, decreasing the learning rate by a factor of 2 after 20, 30, 50, 60, 70, 80 and 90 epochs. We use 90k and 30k samples for training and validation, respectively. The training is performed in mixed-precision to speed up computation and distributed across 3 GPUs.

The full precision network (FPN) is trained from scratch using Xavier uniform initialization [80] (which helps with the sparsity of the input) as the pre-trained classification models on the real-world ImageNet [81] dataset have little relation to our calorimeter images. A common challenge when training models from scratch is the insufficient amount of training data which may lead to overfitting. However, it is not a problem in our case: the training dataset is large enough and, if overfitting occurred, we can go back and generate an even larger one. For TWN training we find out that pre-loading trained FPN weights greatly speeds up the process. And per-layer and per-channel scaling factor α improves the results.

The final detection layer returns a classification label (background, b, W/H or t jet) and three regression values. Two of them correspond to the centre of the jet, i.e. offset in η and ϕ plane from the anchor. The last one is jet mass regression which is an example of an auxiliary function that Jet-SSD can be tasked with.

Figure 4. An example of the Jet-SSD at inference for one event with the input calorimeter image and highlighted true labels (top) and prediction bounding boxes (bottom). The fractional number next to the categorical label corresponds to the network confidence score.

6 Results

An example of the Jet-SSD in action is shown in Figure 4. Jet-SSD outputs predicted categorical label of the object, confidence and bounding boxes. In object detection true positive is defined as prediction with category equal to the ground truth label and Intersection over Union (IoU) above the predefined threshold, in our case 0.5. Successful prediction meets both criteria, otherwise, it is considered as a false negative.

To evaluate the model we use precision and recall (true positive rate), and average precision (AP) metric, which is computed for each category separately. Classification tasks usually report on the receiver operator characteristic (ROC) curve, which is a function of the false positive rate (fall-out or the background efficiency) as a function of the true positive rate (sensitivity or signal efficiency). In the case of object detection, the false positive rate is not very informative as there is a big imbalance between positive and negative class (there are no objects in most locations). Thus, the false positive rate is replaced by precision or positive predictive value (PPV). Intuitively, precision measures how accurate the predictions while recall measures the quality of the positive predictions. To draw a precision-recall (PR) curve, the predictions are first sorted in order of confidence followed by calculation of posi-

Figure 5. Jet-SSD PR curve and AP values for three target categories (b, W-H and t) and two weights precision (FPN and TWN).

tive predictive value and true positive rate for each confidence threshold. For the relationship between ROC and PR curve, see [82].

The PR curve of Jet-SSD, evaluated on a held-out test dataset consisting of 90k samples, is shown in Figure 5. The TWN results are closely matching the results of the FPN, which is reflected in an AP score. To calculate the value of AP, the maximum precision is calculated for the recall values that range from 0 to 1 with a step size of 0.1 and finally averaging over the results. From the PR curve, we can conclude that t jets are the easiest to identify while b jets detection is lacking. The result is not surprising for two reasons. Firstly, b jets have a lower momentum threshold, making the energy deposits more challenging to detect. Secondly, CNN based object detection is more challenging as the scale of the target object decreases; and b jets have a smaller radius than t, W and H jets. The latter issue can be further mitigated as small scale object detection is an active research field in machine learning (for example [83]).

Finally, we report the mean and median localization error in ϕ and η and the relative error in mass regression. These results are shown in Figure 6. The ϕ localization error is smaller than η due to input information loss. Remind that we limit input in η dimension. In the case when the jet centre is close to the edge, i.e. $|\eta| \approx 3.0$, part of the information is lost beyond image boundaries. Due to the cylindrical structure of the detector, this is not happening in ϕ dimension. Furthermore, we notice that the error does not decrease with p_T for η for which we do not find a reason. Finally, the mass regression relative error can be further decreased with re-balancing of the SSD training loss, i.e. increasing regression error contribution to back-propagation by introducing a new scaling hyper-parameter β : $loss = classification + localization + \beta \times auxiliary$, where $\beta > 1$.

7 Conclusions

In this paper, we introduced Jet-SSD, a deep learning network able to simultaneously localize, tag and estimate the mass of jets, a collimated spray of particles produced in high energy

Figure 6. The mean (μ) and median (Q2) localization error in η (left column), ϕ (centre column) and relative error in mass regression between ground truth (GT) and Jet-SSD output (SSD) for FPN and TWN versions. The results are reported for each class independently: b jets (top row), W-H jets (centre row) and t jest (bottom row). All results are calculated as a function of p_T .

physics experiments. We showed that the compressed model via quantized weights to ternary values with layer- and channel-dependent scaling factor closely matches the performance of the full precision model. We seek to examine the performance of the network on dedicated hardware.

Acknowledgments

A. A. P., M. P., S. S. and V. L. are supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement n^o 772369). V. L. is supported by Zenseact under the CERN Knowledge Transfer Group. A. A. P. is supported by CEVA under the CERN Knowledge Transfer Group. We thank Simons Foundation, Flatiron Institute and Ian Fisk for granting access to computing resources used for this project.

References

- [1] J.M. Butterworth, A.R. Davison, M. Rubin, G.P. Salam, Physical review letters **100**, 242001 (2008)
- [2] W. Skiba, D. Tucker-Smith, Physical Review D **75**, 115010 (2007)
- [3] V. Khachatryan, A.M. Sirunyan, A. Tumasyan, W. Adam, T. Bergauer, M. Dragicevic, J. Erö, C. Fabjan, M. Friedl, R. Fruehwirth et al., Journal of High Energy Physics **2014**, 173 (2014)
- [4] G. Aad, B. Abbott, J. Abdallah, R. Aben, M. Abolins, O. AbouZeid, H. Abramowicz, H. Abreu, R. Abreu, Y. Abulaiti et al., Journal of High Energy Physics **2015**, 1 (2015)
- [5] D. Adams, A. Arce, L. Asquith, M. Backovic, T. Barillari, P. Berta, D. Bertolini, A. Buckley, J. Butterworth, R.C. Toro et al., The European Physical Journal C **75**, 1 (2015)
- [6] A. Abdesselam, A. Belyaev, E.B. Kuutmann, U. Bitenc, G. Brooijmans, J. Butterworth, P.B. de Renstrom, D.B. Franzosi, R. Buckingham, B. Chapleau et al., The European Physical Journal C **71**, 1 (2011)
- [7] A. Altheimer, S. Arora, L. Asquith, G. Brooijmans, J. Butterworth, M. Campanelli, B. Chapleau, A. Cholakian, J. Chou, M. Dasgupta et al., Journal of Physics G: Nuclear and Particle Physics **39**, 063001 (2012)
- [8] A. Altheimer, A. Arce, L. Asquith, J.B. Mayes, E.B. Kuutmann, J. Berger, D. Bjergaard, L. Bryngemark, A. Buckley, J. Butterworth et al., The European Physical Journal C **74**, 1 (2014)
- [9] T. Plehn, M. Spannowsky, M. Takeuchi, D. Zerwas, Journal of High Energy Physics **2010**, 1 (2010)
- [10] A.J. Larkoski, S. Marzani, G. Soyez, J. Thaler, Journal of High Energy Physics **2014**, 146 (2014)
- [11] J. Thaler, K. Van Tilburg, Journal of High Energy Physics **2011**, 15 (2011)
- [12] A.J. Larkoski, G.P. Salam, J. Thaler, Journal of High Energy Physics **2013**, 108 (2013)
- [13] D. Krohn, J. Thaler, L.T. Wang, Journal of High Energy Physics **2010**, 84 (2010)
- [14] S.D. Ellis, C.K. Vermilion, J.R. Walsh, Physical Review D **81**, 094023 (2010)
- [15] M. Dasgupta, A. Fregoso, S. Marzani, G.P. Salam, Journal of High Energy Physics **2013**, 29 (2013)
- [16] M. Dasgupta, A. Fregoso, S. Marzani, A. Powling, The European Physical Journal C **73**, 1 (2013)
- [17] M. Dasgupta, A. Powling, A. Siodmok, Journal of High Energy Physics **2015**, 1 (2015)
- [18] J. Cogan, M. Kagan, E. Strauss, A. Schwartzman, Journal of High Energy Physics **2015**, 118 (2015)
- [19] J. Pearkes, W. Fedorko, A. Lister, C. Gay, arXiv preprint arXiv:1704.02124 (2017)
- [20] P. Baldi, K. Bauer, C. Eng, P. Sadowski, D. Whiteson, Physical Review D **93**, 094034 (2016)
- [21] S. Macaluso, D. Shih, Journal of High Energy Physics **2018**, 1 (2018)
- [22] L.G. Almeida, M. Backović, M. Cliche, S.J. Lee, M. Perelstein, Journal of High Energy Physics **2015**, 1 (2015)
- [23] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, A. Schwartzman, Journal of High Energy Physics **2016**, 1 (2016)
- [24] D. Guest, J. Collado, P. Baldi, S.C. Hsu, G. Urban, D. Whiteson, Physical Review D **94**, 112002 (2016)
- [25] J. Barnard, E.N. Dawe, M.J. Dolan, N. Rajcic, Physical Review D **95**, 014018 (2017)

- [26] A. Butter, G. Kasieczka, T. Plehn, M. Russell, *SciPost Phys* **5**, 028 (2018)
- [27] P.T. Komiske, E.M. Metodiev, M.D. Schwartz, *Journal of High Energy Physics* **2017**, 110 (2017)
- [28] J. Lin, M. Freytsis, I. Moutl, B. Nachman, *Journal of High Energy Physics* **2018**, 1 (2018)
- [29] G. Kasieczka, T. Plehn, A. Butter, K. Cranmer, D. Debnath, B.M. Dillon, M. Fairbairn, D.A. Faroughy, W. Fedorko, C. Gay et al., *arXiv preprint arXiv:1902.09914* (2019)
- [30] G. Kasieczka, T. Plehn, M. Russell, T. Schell, *Journal of High Energy Physics* **2017**, 6 (2017)
- [31] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, *Proceedings of the IEEE* **86**, 2278 (1998)
- [32] M. Cacciari, G.P. Salam, G. Soyez, *The European Physical Journal C* **72**, 1 (2012)
- [33] A.M. Sirunyan et al. (CMS), *Comput. Softw. Big Sci.* **4**, 10 (2020), 1912.06046
- [34] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, Y. LeCun, *arXiv preprint arXiv:1312.6229* (2013)
- [35] K. Zhang, Z. Zhang, Z. Li, Y. Qiao, *IEEE Signal Processing Letters* **23**, 1499 (2016)
- [36] L. Zhang, L. Lin, X. Liang, K. He, *Is faster R-CNN doing well for pedestrian detection?*, in *European conference on computer vision* (Springer, 2016), pp. 443–457
- [37] Z. Zou, Z. Shi, Y. Guo, J. Ye, *arXiv preprint arXiv:1905.05055* (2019)
- [38] L. Liu, W. Ouyang, X. Wang, P. Fieguth, J. Chen, X. Liu, M. Pietikäinen, *International journal of computer vision* **128**, 261 (2020)
- [39] J. Redmon, A. Farhadi, *YOLO9000: better, faster, stronger*, in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017), pp. 7263–7271
- [40] J. Redmon, S. Divvala, R. Girshick, A. Farhadi, *You only look once: Unified, real-time object detection*, in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 779–788
- [41] C.Y. Fu, W. Liu, A. Ranga, A. Tyagi, A.C. Berg, *arXiv preprint arXiv:1701.06659* (2017)
- [42] X. Zhou, D. Wang, P. Krähenbühl, *arXiv preprint arXiv:1904.07850* (2019)
- [43] T.Y. Lin, P. Goyal, R. Girshick, K. He, P. Dollár, *Focal loss for dense object detection*, in *Proceedings of the IEEE international conference on computer vision* (2017), pp. 2980–2988
- [44] R. Girshick, J. Donahue, T. Darrell, J. Malik, *Rich feature hierarchies for accurate object detection and semantic segmentation*, in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2014), pp. 580–587
- [45] S. Ren, K. He, R. Girshick, J. Sun, *arXiv preprint arXiv:1506.01497* (2015)
- [46] R. Girshick, *Fast r-cnn*, in *Proceedings of the IEEE international conference on computer vision* (2015), pp. 1440–1448
- [47] J. Dai, Y. Li, K. He, J. Sun, *arXiv preprint arXiv:1605.06409* (2016)
- [48] H. Xu, X. Lv, X. Wang, Z. Ren, N. Bodla, R. Chellappa, *Deep regionlets for object detection*, in *Proceedings of the European Conference on Computer Vision (ECCV)* (2018), pp. 798–814
- [49] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.Y. Fu, A.C. Berg, *Ssd: Single shot multibox detector*, in *European conference on computer vision* (Springer, 2016), pp. 21–37
- [50] K. Simonyan, A. Zisserman, *arXiv preprint arXiv:1409.1556* (2014)
- [51] Y. Cheng, D. Wang, P. Zhou, T. Zhang, *arXiv preprint arXiv:1710.09282* (2017)

- [52] Y. LeCun, J.S. Denker, S.A. Solla, R.E. Howard, L.D. Jackel, *Optimal brain damage*, in *NIPs* (Citeseer, 1989), Vol. 2, pp. 598–605
- [53] S. Han, H. Mao, W.J. Dally, arXiv preprint arXiv:1510.00149 (2015)
- [54] C. Louizos, M. Welling, D.P. Kingma, arXiv preprint arXiv:1712.01312 (2017)
- [55] A. Sironi, B. Tekin, R. Rigamonti, V. Lepetit, P. Fua, *IEEE transactions on pattern analysis and machine intelligence* **37**, 94 (2014)
- [56] E. Denton, W. Zaremba, J. Bruna, Y. LeCun, R. Fergus, arXiv preprint arXiv:1404.0736 (2014)
- [57] M. Jaderberg, A. Vedaldi, A. Zisserman, arXiv preprint arXiv:1405.3866 (2014)
- [58] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, *Going deeper with convolutions*, in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2015), pp. 1–9
- [59] A.G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, H. Adam, arXiv preprint arXiv:1704.04861 (2017)
- [60] F.N. Iandola, S. Han, M.W. Moskewicz, K. Ashraf, W.J. Dally, K. Keutzer, arXiv preprint arXiv:1602.07360 (2016)
- [61] T. Cohen, M. Welling, *Group equivariant convolutional networks*, in *International conference on machine learning* (PMLR, 2016), pp. 2990–2999
- [62] C. Buciluă, R. Caruana, A. Niculescu-Mizil, *Model compression*, in *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining* (2006), pp. 535–541
- [63] M. Courbariaux, Y. Bengio, J.P. David, arXiv preprint arXiv:1511.00363 (2015)
- [64] M. Courbariaux, I. Hubara, D. Soudry, R. El-Yaniv, Y. Bengio, arXiv preprint arXiv:1602.02830 (2016)
- [65] S. Zhou, Y. Wu, Z. Ni, X. Zhou, H. Wen, Y. Zou, arXiv preprint arXiv:1606.06160 (2016)
- [66] M. Rastegari, V. Ordonez, J. Redmon, A. Farhadi, *Xnor-net: Imagenet classification using binary convolutional neural networks*, in *European conference on computer vision* (Springer, 2016), pp. 525–542
- [67] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, Y. Bengio, *The Journal of Machine Learning Research* **18**, 6869 (2017)
- [68] F. Li, B. Zhang, B. Liu, arXiv preprint arXiv:1605.04711 (2016)
- [69] C. Zhu, S. Han, H. Mao, W.J. Dally, arXiv preprint arXiv:1612.01064 (2016)
- [70] E.H. Lee, D. Miyashita, E. Chai, B. Murmann, S.S. Wong, *Lognet: Energy-efficient neural networks using logarithmic computation*, in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (IEEE, 2017), pp. 5900–5904
- [71] Z. Cai, X. He, J. Sun, N. Vasconcelos, *Deep learning with low precision by half-wave gaussian quantization*, in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017), pp. 5918–5926
- [72] CMS Collaboration, arXiv preprint arXiv:1609.02366 (2016)
- [73] W. Bhimji, S.A. Farrell, T. Kurth, M. Paganini, E. Racah et al., **1085**, 042034 (2018)
- [74] CMS Collaboration, *JInst* **3**, S08004 (2008)
- [75] T. Sjöstrand, S. Mrenna, P. Skands, *Computer Physics Communications* **178**, 852 (2008)
- [76] J. De Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaitre, A. Mertens, M. Selvaggi, D. Collaboration et al., *Journal of High Energy Physics* **2014**, 57 (2014)
- [77] S. Ioffe, C. Szegedy, *Batch normalization: Accelerating deep network training by reducing internal covariate shift* (2015), 1502.03167

- [78] E. Sari, M. Belbahri, V.P. Nia, *How does batch normalization help binary training?* (2020), 1909.09139
- [79] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshein, L. Antiga et al., *PyTorch: An Imperative Style, High-Performance Deep Learning Library* (2019)
- [80] X. Glorot, Y. Bengio, *Understanding the difficulty of training deep feedforward neural networks*, in *Proceedings of the thirteenth international conference on artificial intelligence and statistics* (JMLR Workshop and Conference Proceedings, 2010), pp. 249–256
- [81] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, Li Fei-Fei, *ImageNet: A large-scale hierarchical image database*, in *2009 IEEE Conference on Computer Vision and Pattern Recognition* (2009), pp. 248–255
- [82] J. Davis, M. Goadrich, *The relationship between Precision-Recall and ROC curves*, in *Proceedings of the 23rd international conference on Machine learning* (2006), pp. 233–240
- [83] J. Rabbi, N. Ray, M. Schubert, S. Chowdhury, D. Chao, *Remote Sensing* **12**, 1432 (2020)