

3DQ: Compact Quantized Neural Networks for Volumetric Whole Brain Segmentation

Magdalini Paschali^{1*}, Stefano Gasperini^{1*}, Abhijit Guha Roy²,
Michael Y.-S. Fang³, and Nassir Navab^{1,4}

¹ Computer Aided Medical Procedures, Technical University of Munich, Germany

² Department of Child and Adolescent Psychiatry, Psychosomatic and Psychotherapy, Ludwig Maximilian University, Munich, Germany

³ Department of Physics, University of California Berkeley, USA

⁴ Computer Aided Medical Procedures, Johns Hopkins University, USA

Abstract. Model architectures have been dramatically increasing in size, improving performance at the cost of resource requirements. In this paper we propose 3DQ, a ternary quantization method, applied for the first time to 3D Fully Convolutional Neural Networks (F-CNNs), enabling 16x model compression while maintaining performance on par with full precision models. We extensively evaluate 3DQ on two datasets for the challenging task of whole brain segmentation. Additionally, we showcase our method’s ability to generalize on two common 3D architectures, namely 3D U-Net and V-Net. Outperforming a variety of baselines, the proposed method is capable of compressing large 3D models to a few MBytes, alleviating the storage needs in space-critical applications.

Keywords: Quantization · Volumetric · Segmentation · Deep Learning

1 Introduction

Fully Convolutional Neural Networks (F-CNNs) have been incorporated successfully into numerous Computer Assisted Diagnosis (CAD) systems performing various medical image analysis tasks with increasing difficulty and requirements. Hence, their size has grown drastically reaching commonly hundreds of layers and several million parameters.

An additional factor towards the size explosion of F-CNNs in CADs is that medical data is in most cases volumetric by nature and has been continuously increasing in resolution. This growth of F-CNNs has shortcomings regarding resource requirements, such as computation, energy consumption and storage. Towards this end, we propose, for the first time in 3D F-CNNs, a quantization-based technique that offers model compression up to 16 times without any loss in performance.

To this day, various segmentation methods process volumetric data per slice using 2D F-CNNs [1]. Such methods yield satisfying results but are not able

* The authors contributed equally.

to fully exploit the contextual information from adjacent slices. Existing 3D F-CNNs, like V-Net [2], 3D U-Net [3] and VoxResNet [4] have achieved state-of-the-art performance in a plethora of segmentation tasks utilizing 3D kernels. However, their millions of parameters require a large amount of space for storage.

Deep learning-based CADs deploying 3D F-CNNs are being integrated into the medical workflow [5]. Hospitals, which were already burdened with storing a myriad of large medical records, now have to allocate additional storage for the trained models in CADs. On top of that, the emerging field of patient-specific care [6], while improving personalized diagnosis and monitoring, will increase the need for storage in medical facilities even further.

Model compression has been an active area of research in the past few years aiming at deploying state-of-the-art F-CNNs in low-power and resource-limited devices, like smartphones and embedded electronics. An additional potential use-case is represented by the decentralized training approach of Federated Learning [7]. Despite maintaining the privacy of client data, iteratively sending over the internet millions of parameters to train a global model becomes unstable with unreliable connections, and compressed models would improve the process.

Multiple techniques, like parameter pruning, low-rank factorization, knowledge distillation and weight quantization [8], have been proposed to compress the size of F-CNNs without compromising their performance. Specifically, weight quantization to binary [9] and ternary values [10], [11], [12] has been among the most popular methods for this task. This can be attributed to its additional advantage of allowing for impressive speed-up during training and inference by approximating convolutions with XNOR and bitcounting operations [9]. Even though it is generally accepted that XNOR-Net revolutionized this field [9], there is a significant trade-off between performance and speed, not ideal for medical applications. TernaryNet [10] was the first attempt in medical imaging to create compact and efficient F-CNNs utilizing ternary weights, where a 2D U-Net was employed for the task of per-slice pancreas CT segmentation. However, extending quantization to 3D F-CNNs has yet to be explored.

In this paper, our contribution is two-fold: 1) We propose, for the first time in 3D F-CNNs, a quantization mechanism with a novel bit-scaling scheme which we name 3DQ. 3DQ integrates two trainable scaling factors and a normalization parameter that increases the learning capacity of a model while maintaining compression. 2) We extensively evaluate 3DQ on the challenging task of 3D whole brain segmentation, showcasing that state-of-the-art performance can be combined with impressive compression rates.

2 Method

2.1 Weight Quantization

The main goal of our quantization method is approximating the full precision weights W by their ternary counterparts $\{-1, 0, 1\}$, \tilde{W} , as formalized in Eq. 1.

The first step is computing the threshold Δ , based on which W will be assigned into three quantization bins. Various approaches leverage a single Δ for

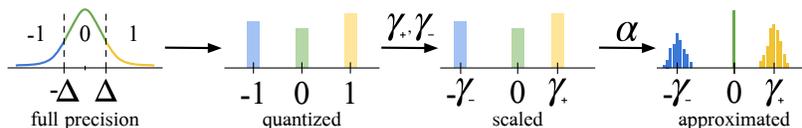


Fig. 1: Overview of 3DQ. The full precision weights are quantized into ternary values, scaled by γ_{\pm} and then further diversified by the factor α .

the entire network [10]. However, 3DQ computes one Δ per layer l , to maintain the variability in weight range values within each layer and avoid weight sparsity [11]. Specifically, each Δ_l is computed as $\Delta_l = t \cdot \max(|W_l|)$: the maximum absolute value of the weights in each layer is multiplied by a constant factor t which moderates the weight sparsity and is consistent among all layers. We set t to 0.05, as suggested by [11] as it proved to be the optimal trade-off between sparsity and accuracy.

After thresholding, the acquired ternary weights \tilde{W} are multiplied by a set of scaling factors, since training an entire F-CNN with $\{-1, 0, 1\}$ weight values would lead to substantial suboptimal performance. 3DQ utilizes two scaling factors, γ_l^+ and γ_l^- [11], which are variables learned for each layer l during training, in contrast to previous methods such as [10, 9].

$$W_l \approx \tilde{W}_l = \begin{cases} +\gamma_l^+ \cdot \alpha & \text{if } W_l > \Delta_l \\ 0 & \text{if } |W_l| < \Delta_l \\ -\gamma_l^- \cdot \alpha & \text{otherwise.} \end{cases} \quad (1)$$

Furthermore, unlike [11], we incorporated an additional scaling factor α into 3DQ [13]. α is calculated from W as the average of the weights with an absolute value larger than Δ_l , as $\alpha = \frac{1}{n_{\Delta_l}} \sum |\tilde{W}_l| |W_l|$, where $n_{\Delta_l} = \sum |\tilde{W}_l|$. α enhances the approximation of the full precision weights, since it spreads the quantized values within the same bin, increasing the expressivity and diversity of the weights in between the various channels of each layer. The quantization process is schematically summarized in Fig. 1.

2.2 Compact Weight Storing

As explained above, both W and \tilde{W} are required during training in order to achieve the optimization of the model weights and the learned scaling factors γ^+ and γ^- . However, during inference, the full precision weights are no longer necessary and therefore there is no need to store them.

Since after the scaling of \tilde{W} , the values still use 32 bits, it is important to store the model ensuring the 16x compression rate offered by the ternary weights. Towards this end, we separate each kernel into three components, as can be seen in Fig. 2: 1) The trained scaling factors γ^+ and γ^- , which are as many as the layers of the chosen F-CNN architecture. 2) The α values, which are computed from the full precision weights and sum up to as many as the

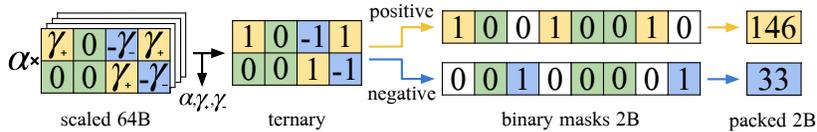


Fig. 2: Overview of 3DQ compression method to store the model parameters. The scaling factors are separated from the ternary values, which are further split into binary vectors, then packed from 8 bits to a byte.

channels at each layer. 3) The ternary weights, which constitute the majority of the model parameters, summing up to millions of values in cases of 3D models.

The scaling factors are stored as full precision variables, requiring 32 bits of disk space per value. Meanwhile, each ternary weight kernel is split into two binary masks, one for the positive weights and one for the negatives. The unmarked locations in both masks stand for the zero weights. Afterwards, the masks go through bit packing and 8 weight bits are stored in 1 byte. The same process is followed backwards to restore the saved models: first unpack the weight values, then multiply them with the saved scaling factors. This method achieves impressive compression rates, particularly crucial for large 3D models, where oftentimes 45M parameters need to be stored for each network [2].

3 Experimental Setup

Datasets We evaluated 3DQ on two challenging and publicly available medical imaging 3D segmentation datasets, namely the Multi-Atlas Labelling Challenge (MALC) [14] and the Hippocampus (HC) Segmentation dataset from the Medical Decathlon challenge [15]. MALC is part of the OASIS dataset and contains 30 whole brain MRI T1 scans with manual annotations. The input volumes are sized $256 \times 256 \times 256$, which were sampled in cubic patches of size 64. Maintaining the original challenge split, we used 15 scans for training and 15 for testing. We considered 28 classes for the segmentation, following [1], and we repeated all the experiments 5 times.

HC includes 263 training samples sized on average $36 \times 50 \times 35$, which we padded to cubes sized 64. Due to the public unavailability of the test set, we performed 5-fold cross-validation for all our experiments, dividing the given dataset to 80/20 patient-level splits. In this dataset, the volumes are segmented into 3 classes, 2 parts of the hippocampus (hippocampus proper and hippocampal formation) and the background [15].

Model Training To highlight the generalizability of our method, for the first time, we quantized common 3D F-CNN architectures, namely 3D U-Net [3] on MALC and HC, and V-Net [2] on MALC. The aforementioned models are highly suitable candidates for quantization and compression, since they have 16M and 45M trainable parameters respectively and require up to 175MB of storage space.

For both datasets, we trained the models with a composite equally-balanced loss function comprised of Dice loss and weighted cross entropy. The weights were

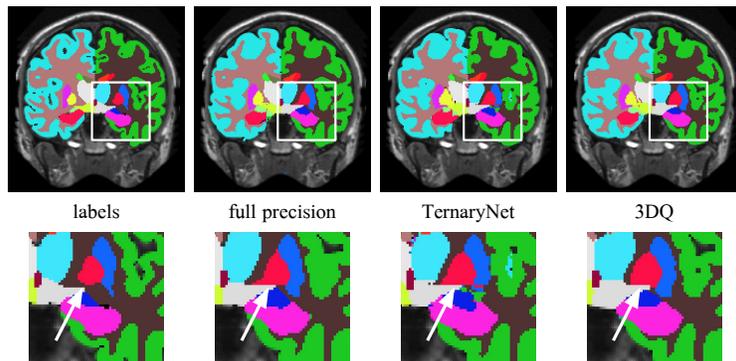


Fig. 3: Qualitative results of 3DQ compared with baseline methods. White arrows on the zoomed views show the superior segmentation performance of 3DQ.

computed with median frequency balancing [16] to circumvent class imbalance. We used an Adam optimizer, initialized with learning rate 0.0001 for 3D U-Net and 0.00005 for V-Net. We trained all models on an NVIDIA Titan Xp Pascal GPU and implemented 3DQ on PyTorch. Although all 3D U-Net models were trained from scratch, we found beneficial starting the quantized experiments on V-Net from a pretrained version.

Evaluation Metrics The main idea of our approach is compressing quantized models, without sacrificing performance with respect to their full precision counterparts. Therefore, the evaluation, which we will discuss in Section 4, is based on two different criteria: the Dice coefficient achieved by the models across volumes and the storage space required to save them.

Ablative testing In order to showcase the effectiveness of the main components of 3DQ, namely the ternary weights and the addition of the scaling factor α , we performed ablative testing. Comparing 3DQ with BTQ, an adjusted binarized version of Trained Ternary Quantization (TTQ) [11], is an essential experiment to highlight the benefits of ternary weights. Furthermore, we compared 3DQ against TTQ to demonstrate the contribution of α as a scaling factor.

Baseline comparison Furthermore, 3DQ was evaluated against its full precision counterpart in order to investigate whether quantized models are able to match the performance of full precision networks. 3DQ was also compared with TernaryNet [10], which was recently proposed for the compression of 2D U-Net [17] for pancreas CT segmentation. Additionally, as an alternative compression method, we deployed knowledge distillation with a temperature $T = 40$ to train scaled-down variants [18] of 3D U-Net and V-Net, that take up exactly the same storage space as the original-size models compressed with 3DQ.

4 Results and Discussion

Ablative Testing Table 1 highlights that models quantized with ternary weights outperform their binary counterparts by 3-11% for MALC and 7% for HC, thanks

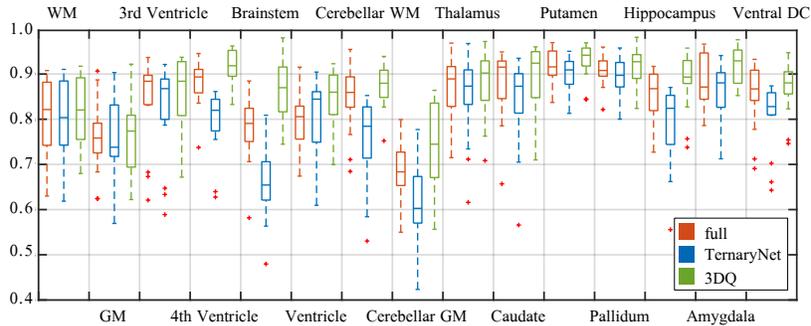


Fig. 4: Box plot of Dice scores achieved by 3D U-Net comparing 3DQ with baseline methods on the right hemisphere classes of MALC.

to the higher learning capacity, motivating the choice of ternary weights in 3DQ. Moreover, Table 1 showcases the positive impact of the scaling factor α we introduced, allowing 3DQ to perform up to 1-2% better than TTQ on both datasets, with a lower standard deviation. α mitigates the quantization constraints, enabling the ternary weights to assume a larger variety of values, better resembling their full precision counterparts.

Comparative Methods As can be seen in Table 2, 3D U-Net quantized with 3DQ performs as good as the full precision model in the case of HC, while it outperforms it by over 2% for MALC. This can be attributed to the quantization acting as a regularization technique by limiting the dynamic range of the weight values.

3DQ also outperformed TernaryNet across all experiments for MALC and HC, with a margin ranging from 7 to over 10%, thanks to the learned scaling factors γ_{\pm} and the absence of the hyperbolic ternary tangent which bounds the activation values and limits the learning capacity of the model.

Figure 3 shows sample segmentations for a slice of a MALC volume from 3D U-Net. A zoomed view of the segmentations, indicates important subcortical structures with a white arrow. Both full precision and TernaryNet predictions suffer from over-inclusion of small structures and spurious misclassified regions. The box plot in Figure 4 confirms the higher quality of the segmentations produced by 3DQ, reporting the Dice scores on the right hemisphere structures. 3DQ outperformed both full precision and TernaryNet, with fewer outliers, demonstrating more uniform results throughout the samples.

Comparison with Knowledge Distillation Another experiment reported in Table 2 is the comparison of 3DQ with knowledge distillation. In order to

	BTQ	TTQ	3DQ
HC			
$_{3DU}$	0.847 ± 0.009	0.912 ± 0.008	0.915 ± 0.006
MALC			
$_v$	0.770 ± 0.013	0.790 ± 0.010	0.802 ± 0.004
$_{3DU}$	0.735 ± 0.005	0.828 ± 0.007	0.844 ± 0.006

Table 1: Comparison of Dice scores of 3DQ with TTQ and its binarized version BTQ on HC and MALC, with 3D U-Net and V-Net.

		Full	Distilled	TernaryNet	3DQ
HC	3DU	0.914 ± 0.005	0.908 ± 0.019	0.845 ± 0.013	0.915 ± 0.006
	V	0.815 ± 0.008	0.715 ± 0.001	0.696 ± 0.016	0.802 ± 0.004
MALC	3DU	0.822 ± 0.005	0.730 ± 0.008	0.774 ± 0.012	0.844 ± 0.006

Table 2: Comparison of Dice scores of 3DQ with baseline methods. Tests performed on HC and MALC, with 3D U-Net and V-Net.

match the 3DQ model sizes while keeping the full precision weights, the distilled networks have 16x fewer parameters than the full models. Even though the smaller networks achieve almost equal performance to the full model for HC, the margin is increased for MALC, where the student networks achieved a 9-10% worse Dice score than the full models for both 3D U-Net and V-Net. This drop in performance can be attributed to the fact that distilled models are 16x smaller than the original ones and additionally have to rely on the predictions of a teacher network, limiting their learning capacity. 3DQ also outperforms the distilled models across the board by a substantial 8-11% on MALC, constituting a successful model compression choice.

Quantization on Different Architectures Table 2 highlights the effects of quantization in the two different 3D model architectures, specifically 3D U-Net and V-Net. Although 3D U-Net is 3x smaller than V-Net, it achieved higher Dice scores in our experiments on MALC, especially when quantized. While the full models performed similarly with a difference of 1%, the quantized 3D U-Net achieved a 4% higher Dice than the quantized V-Net. We attribute this difference in performance to the fact that MALC consists of 15 training volumes, which are considered limited data for V-Net, a large model with 45M trainable parameters, in contrast with 3D U-Net that has 16M.

Compression The models storage requirements are shown in Table 3. By using ternary weights, TernaryNet, TTQ and 3DQ reduce the storage size by a factor of 16, compared to a full precision model. The different scaling factors impact the storage by only a few KBytes. Binary weights additionally reduce the storage required by 2x in comparison to ternary ones, at the cost of lower Dice scores, due to limited learning capacity.

	Full	Distilled	Ternary	Binary
3D U-Net	63MB	3.9MB	3.9MB	2.0MB
V-Net	175MB	11MB	11MB	5.5MB

Table 3: Model size in MBytes for full precision and compressed models.

5 Conclusion

In this paper we proposed 3DQ, a ternary quantization method, which was applied for the first time to 3D F-CNNs on the challenging task of volumetric whole brain segmentation. The models quantized with 3DQ achieved equal or better Dice scores than the baselines, including the full precision models, across two datasets. Thanks to 16x model compression, 3DQ constitutes a valid ap-

proach for storage-critical applications, as patient-specific networks or weights transferring during Federated Learning.

References

1. A. G. Roy, S. Conjeti, N. Navab, and C. Wachinger. QuickNAT: A fully convolutional network for quick and accurate segmentation of neuroanatomy. *NeuroImage*, 186:713–727, 2019.
2. F. Milletari, N. Navab, and S.-A. Ahmadi. V-Net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*. IEEE, 2016.
3. Ö. Çiçek, A. Abdulkadir, S. S Lienkamp, T. Brox, and O. Ronneberger. 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In *MICCAI*, pages 424–432. Springer, 2016.
4. H. Chen, Q. Dou, L. Yu, J. Qin, and P.-A. Heng. VoxResNet: Deep voxelwise residual networks for braisegmentation from 3d MR images. *NeuroImage*, 170, 2018.
5. S. Thaler and V. Menkovski. The role of deep learning in improving healthcare. In *Data Science for Healthcare - Methodologies and Applications*, pages 75–116. 2019.
6. R. Vivanti, L. Joskowicz, N. Lev-Cohain, A. Ephrat, and J. Sosna. Patient-specific and global convolutional neural networks for robust automatic liver tumor delineation in follow-up ct studies. *MBEC*, 56(9):1699–1713, 2018.
7. J. Konečný, H B. McMahan, F. X Yu, P. Richtárik, A. T. Suresh, and D. Bacon. Federated Learning: Strategies for improving communication efficiency. *NIPS Workshop*, 2016.
8. Y. Cheng, D. Wang, P. Zhou, and T. Zhang. A survey of model compression and acceleration for deep neural networks. *arXiv/1710.09282*, 2017.
9. M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi. XNOR-Net: Imagenet classification using binary convolutional neural networks. In *ECCV*, 2016.
10. M. P Heinrich, M. Blendowski, and O. Oktay. TernaryNet: Faster deep model inference without GPUs for medical 3d segmentation using sparse and binary convolutions. *IJCARS*, 13(9):1311–1320, 2018.
11. C. Zhu, S. Han, H. Mao, and W. J Dally. Trained ternary quantization. *ICLR*, 2017.
12. J. Lin, C. Gan, and S. Han. Defensive quantization: When efficiency meets robustness. In *ICLR*, 2019.
13. Li F. and Liu B. Ternary weight networks. *NIPS Workshop (EMDNN)*, 2016.
14. D. S. Marcus, T. H. Wang, J. Parker, J. G. Csernansky, J. C. Morris, and R. L. Buckner. Open Access Series of Imaging Studies (OASIS): Cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. *J. Cognitive Neuroscience*, 19(9):1498–1507, 2007.
15. A. L Simpson, M. Antonelli, S. Bakas, M. Bilello, K. Farahani, B. van Ginneken, et al. A large annotated medical image dataset for the development and evaluation of segmentation algorithms. *arXiv preprint arXiv:1902.09063*, 2019.
16. A. G. Roy, S. Conjeti, D. Sheet, A. Katouzian, N. Navab, and C. Wachinger. Error corrective boosting for learning fully convolutional networks with limited data. In *MICCAI*, pages 231–239. Springer, 2017.
17. O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241. Springer, 2015.
18. G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *NIPS Workshop*, 2015.