QReg: On Regularization Effects of Quantization

MohammadHossein AskariHemmat[1], Reyhane Askari Hemmat[2], Alex Hoffman[3], Ivan Lazarevich[3], Ehsan Saboori[3], Olivier Mastropietro[3], Sudhakar Sah[3], Yvon Savaria[1], Jean-Pierre David [1]

[1]Ecole Polytechnique Montreal, [2]Mila, Universite de Montreal, [3]Deeplite Inc. Montreal

TL;DR

We hypothesize that quantization has regularization effect that is related to quantization level. We model quantization as a form of additive noise. We show that the magnitude of this noise is correlated with the level of quantization. Our empirical results show that the regularization effect of quantization exist regardless of model, dataset, vision task and quantization technique.

Related Work

Regularization effect of quantization has been studied before. We have categorized earlier works into three main lines of work (refer to the paper for full list of references):

- **Effect of Quantization on Accuracy Improvements:** Courbariaux et al.(2015), Mishchenko et al. (2019). (2018), Abdolrashidi et al. (2021)
- **Analytical Studies:** Merolla et al. (2016), Alizadeh et al. (2020), Zhang et al. (2022)
- **Using Quantization for its Regularization Effect:** Hou et al. (2017), Hou & Kwok (2018)

Unlike the previous work, we provide an extensive list of experiments performed on different models, tasks and datasets. Also, to the best of our knowledge, none of the previous works provide a study on how different quantization level is related to induced regularization.

Quantization Noise as a Regularizer

Weight quantization can be characterized as noise perturbing the weights. Such noise can be described as:

$$f(x, w_q) = f(x, w + \delta)$$

Given the mean square error loss, assuming that the quantization noise follows a normal distribution, $\delta \sim \mathcal{N}(0, \sigma I)$, applying a first-order Taylor approximation around the weights of the full precision model, results in the following approximation on the $\widetilde{\mathcal{L}}$:

$$\widetilde{\mathcal{L}} \approx \mathcal{L} + \frac{\sigma\delta^2}{m} \sum_{i=1}^{m} \|\nabla w \hat{y_i}\|_2^2$$

In the above equation, $\frac{\sigma\delta^2}{m}$ shows the amount of regularization that is induced into the loss function due to quantization.

Results

Model	2 bits					4 bits					8 bits				
Resnet20	-9.0	-9.0	-10.5	-8.4	-7.8	-0.2	0.3	-0.2	0.5	0.7	0.7	1.1	0.7	1.7	1.4
Resnet20	-10.2	-5.8	-8.6	6.4	27.1	-0.1	1.1	0.7	0.7	7.0	0.7	1.2	1.6	-0.1	2.1
Resnet20	-1.9	-1.4	-5.8	-6.3	-9.2	2.9	0.7	0.0	0.9	-0.1	0.3	1.1	1.9	1.3	0.5
Resnet20	-0.2	-4.4	-6.3	-0.4	-6.3	1.3	2.2	2.2	0.7	0.2	1.1	1.9	2.1	1.2	2.3
Resnet18	-0.6	0.0	-0.5	-0.7	-0.3	2.7	3.1	2.3	2.4	1.2	3.6	4.2	3.4	3.8	3.5
Resnet18	-0.3	-0.7	-0.6	-4.4	0.3	2.6	1.3	2.7	-2.2	-1.5	3.9	3.1	3.7	-0.4	0.2
Resnet18	1.9	-0.7	-0.3	-0.6	-2.8	3.8	1.7	1.0	1.4	1.7	2.8	2.6	2.0	3.8	2.2
Resnet18	-2.7	-0.6	-0.8	-3.2	-1.5	-0.3	2.0	2.4	-0.2	0.4	1.3	3.6	3.1	1.1	2.5
YOLOv5n	-17.8	-17.4	-13.4	-12.0	-12.3	-4.2	-4.9	-4.3	-5.0	-4.2	3.3	3.2	-1.0	1.0	0.7
YOLOv5n	-13.3	-12.7	-17.1	-15.0	-15.7	-4.4	-1.3	-5.9	-2.4	-5.4	-0.4	2.7	2.1	2.0	2.3
YOLOv5n	-14.0	-17.6	-15.3	-10.8	-17.4	-2.3	-4.6	-5.4	-4.0	-4.8	1.3	2.9	1.4	-0.2	2.1
YOLOv5n	-15.3	-8.8	-17.5	-15.6	-7.1	-2.6	-1.0	-4.7	-3.5	-3.1	2.1	3.1	3.2	1.6	-0.5

The value in each cell shows the test accuracy difference between the quantized and full precision models. The position of each cell shows the corresponding augmentation. ■ quantized model generalized better than full precision model, ■ shows the opposite.

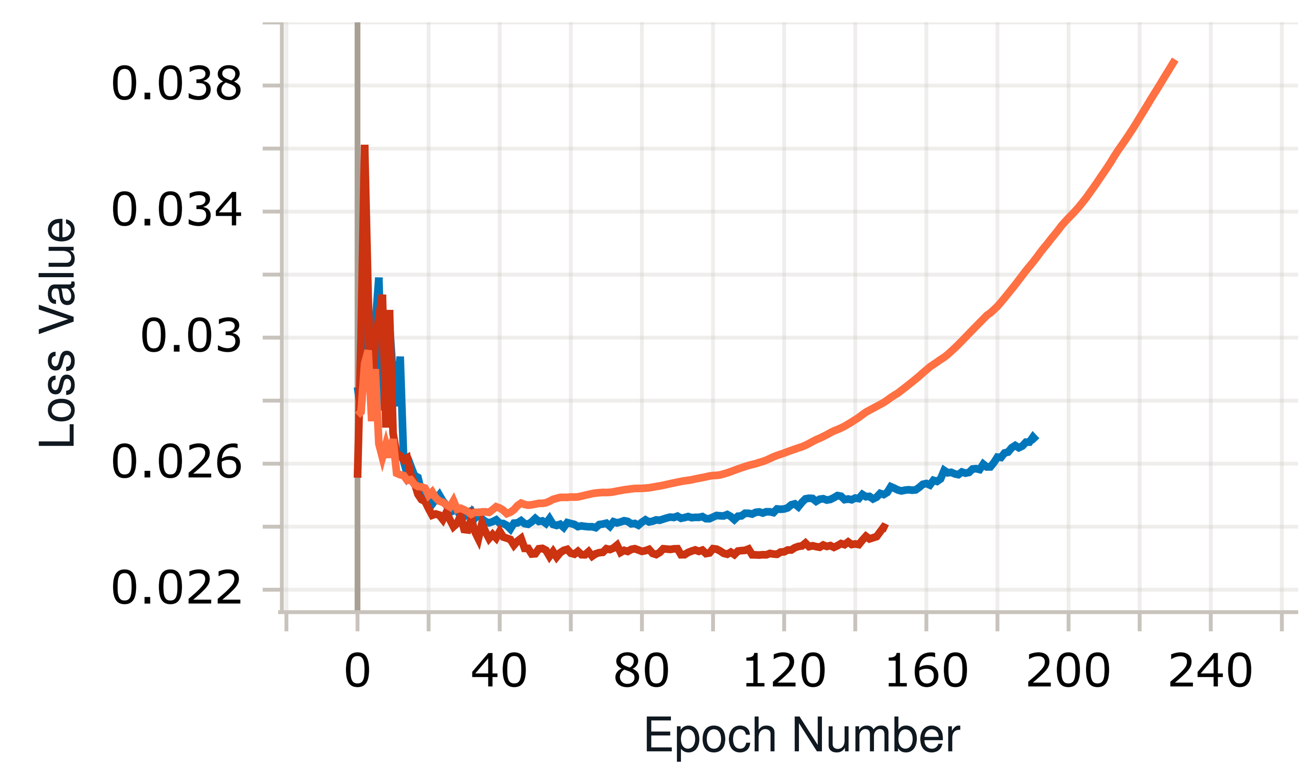

Validation objective loss of 4-bit quantized ■, 8-bit quantized ■ and full precision model ■ for the YOLOv5n model on the VOC dataset. Unlike the quantized model, the full precision model exhibits heavy overfitting. We believe that the quantized model has better generalization performance due to its regularization effect.

Remarks

Our experiments show that the regularization effect of quantization exist regardless of model, dataset, vision task and quantization technique. Finally, we analytically and empirically confirm that this regularization effect is correlated with quantization level. Our empirical study shows that moderate quantization (8-bit) helps models generalize better.

Acknowledgements

The authors would like to thank Mohammad Pezeshki and Anush Sankaran for helpful discussion about designing empirical studies presented in this paper. The authors would also like to acknowledge the support for this project from Fonds de Recherche du Québec Nature et technologies (FRQNT) and Natural Sciences and Engineering Research Council of Canada (NSERC).