Hessian Aware Quantization of Spiking Neural Networks

Hin Wai Lui and Emre Neftci

{hwlui,eneftci}@uci.edu

University of California, Irvine

Irvine, California, USA

ABSTRACT

To achieve the low latency, high throughput, and energy efficiency benefits of Spiking Neural Networks (SNNs), reducing the memory and compute requirements when running on a neuromorphic hardware is an important step. Neuromorphic architecture allows massively parallel computation with variable and local bit-precisions. However, how different bit-precisions should be allocated to different layers or connections of the network is not trivial. In this work, we demonstrate how a layer-wise Hessian trace analysis can measure the sensitivity of the loss to any perturbation of the layer's weights, and this can be used to guide the allocation of a layer-specific bit-precision when quantizing an SNN. In addition, current gradient based methods of SNN training use a complex neuron model with multiple state variables, which is not ideal for compute and memory efficiency. To address this challenge, we present a simplified neuron model that reduces the number of state variables by 4-fold while still being compatible with gradient based training. We find that the impact on model accuracy when using a layer-wise bit-precision correlated well with that layer's Hessian trace. The accuracy of the optimal quantized network only dropped by 0.3%, yet the network size was reduced by 58%. This reduces memory usage and allows fixed-point arithmetic with simpler digital circuits to be used, increasing the overall throughput and energy efficiency.

CCS CONCEPTS

• Computing methodologies → Neural networks.

KEYWORDS

Spiking Neural Networks, Neuromorphic Engineering, Mixed Precision Quantization, Hessian Trace

ACM Reference Format:

Hin Wai Lui and Emre Neftci. 2021. Hessian Aware Quantization of Spiking Neural Networks. In *International Conference on Neuromorphic Systems 2021 (ICONS 2021)*, July 27–29, 2021, Knoxville, TN, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3477145.3477158>

1 INTRODUCTION

Spiking Neural Networks (SNNs) promise low latency, high throughput, energy efficient, and event-driven information processing when paired with appropriate neuromorphic hardware [14]. Unlike Deep Neural Networks (DNNs) with dense activity patterns and dense

matrix multiplications, SNNs process and communicate information with sparse discrete spikes and are naturally equipped with internal memory states (*i.e.* they are recurrent neural networks).

Dedicated neuromorphic hardware can realize the benefits of spike based computation of SNNs. For this work, we focus on digital implementations of such hardware as they are more widespread today. Neuromorphic hardware allows massively parallel computation of neuron dynamics, with memory local to computation, which prevents the Von Neumann bottleneck [9]. To make efficient use of energy and silicon area, digital implementations of neuromorphic hardware often only provide integer or low-precision arithmetic. For example, the TrueNorth chip [6] uses 2-bit trinary weights and performs classification tasks with frame rates exceeding 1,200 frames/s using only 25mW. The Loihi chip [3] supports synaptic plasticity states (*e.g.* weights) with bit-precisions ranging between 1 and 9 bits. Furthermore, the bit-precision of the weights on the Loihi is connection-specific, and this capability is only possible with the parallel computation and local memory architecture of Neuromorphic hardware. Traditional Von Neumann or Single Instruction Multiple Data (SIMD) processors such as CPUs or GPUs cannot efficiently provide connection-specific bit-precision with dense weight matrices, due to the entailing instruction overhead. On the other hand, there is evidence to suggest that the brain adapts such variable weight bit-precision strategy, through synapses with variable sizes [2]. The locality and adaptability of computation and memory resources are what inspired the field of neuromorphic engineering at the first place.

Existing hardware implementations and investigations on neural network quantization [5, 6, 15] argued that integer arithmetic could significantly reduce the memory usage while only marginally increasing the error, provided that the quantization is accounted for during training. [5] showed that fine-tuning a quantized DNN that was fully trained with full precision weights could achieve full precision accuracy, when a learnt quantizer step size was used. [15] achieved near full precision accuracy on a quantized SNN with the same fine-tuning approach, using a uniform bit-precision with 10-12 bits assigned for neuron state variables, and 6 bits for the weights. However, these approaches cannot take full advantage of the connection-specific bit-precision capability provided by neuromorphic hardware. Recently, a Hessian aware approach showed good results on mixed precision quantization of DNNs [4]. It allowed different layers of the network to have a different bit-precision, guided by the Hessian trace of that layer. The layer-wise Hessian trace provides a measure of the sensitivity of the loss to any perturbation of the layer's parameters. Therefore, a layer with a low Hessian trace could afford to have a lower bit-precision than one with a high Hessian trace. With this approach, each layer within the network could use an optimally low bit-precision with only small impact on the loss.

This work is licensed under a Creative Commons Attribution International 4.0 License.

ICONS 2021, July 27–29, 2021, Knoxville, TN, USA

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8691-3/21/07.

<https://doi.org/10.1145/3477145.3477158>

In this work, we explore the use of Hessian aware quantization on SNNs, and adapt the Hutchinson method for fast estimations of Hessian trace on SNNs. We also present a simplified version of a Leaky Integrate-and-Fire (LIF) neuron model that could further compress the memory usage and speed up computation. We introduce this neuron model in section 2, the method of quantization in section 3, the method of computing the Hessian trace of an SNN in section 4, the experiments and results in section 5, and conclusion and future work in section 6.

2 SIMPLIFIED NEURON MODEL

The standard model of an LIF neuron [7] has the following form:

$$\begin{aligned}\rho_i[n] &= \alpha U_i[n-1] S_i[n-1], \\ U_i[n] &= \alpha U_i[n-1] + \sum_j W_{ij} S_j[n-1] - \rho_i[n], \\ S_i[n] &= \Theta(U_i[n] - U_{thres}),\end{aligned}\quad (1)$$

where $U_i[n]$ is the internal membrane potential of the post-synaptic neuron at time step n , S_j is the pre-synaptic spike, S_i is the post-synaptic spike, W_{ij} is the synaptic weight between pre and post-synaptic neuron, α is the exponential decay constant, U_{thres} is the reset threshold, and Θ is the step function. The post-synaptic neuron receives pre-synaptic spikes and integrates them to its membrane potential U_i . When U_i is above the threshold U_{thres} , it generates a spike S_i at that time step. For an LIF neuron, there is a leakage current that results in U_i decaying by the factor α between each time step. In this model, the neuron resets its potential to 0 after a spike is generated. Here we represent this reset behavior with the reset variable ρ_i , which is the decayed membrane potential just before the spike generation. It is subtracted from U_i after the spike generation to reset U_i back to 0.

In this article, we use Deep Continuous Local Learning to train the SNN. DECOLLE enables the scalable training of SNNs locally with very long time sequences. DECOLLE uses an approximation of forward-mode autodiff in which $\frac{\partial U_i}{\partial W_{ij}}$ are traces of past pre-synaptic spiking activities [16] and gradients are computed using layer-wise local classifiers. Implementing DECOLLE [10] with Equation (1) and automatic differentiation is not straightforward. Instead, DECOLLE uses a discrete-time Spike Response Model formulation of Equation (1) [8] which consists of 5 state variables, the eligibility trace of the current based synapse Q_j , the membrane potential trace of the pre-synaptic neuron P_j , the reset dynamical variable of the post-synaptic neuron R_i , the membrane potential U_i , and the spike train S_i . While the DECOLLE neuron model offers some level of biological realism by including synaptic dynamics, it requires more memory and computation than compared to Equation (1). Therefore, we modify it to the following:

$$\begin{aligned}P_j[n] &= \alpha P_j[n-1] + S_j[n-1], \\ R_i[n] &= \alpha R_i[n-1] + \alpha U_i[n-1] S_i[n-1], \\ U_i[n] &= \sum_j W_{ij} P_j[n] - R_i[n], \\ S_i[n] &= \Theta(U_i[n] - U_{thres}).\end{aligned}\quad (2)$$

Equation (2) is simplified from the DECOLLE neuron model with the state variable Q_i and the factor $(1 - \alpha)$ removed. It is also equivalent

to (1). We can prove this by subtracting the state variables of (2) between two subsequent time steps:

$$P_j[n] - \alpha P_j[n-1] = S_j[n-1], \quad (3)$$

$$R_i[n] - \alpha R_i[n-1] = \alpha U_i[n-1] S_i[n-1] = \rho_i[n], \quad (4)$$

$$U_i[n] - \alpha U_i[n-1] = \sum_j W_{ij} (P_j[n] - \alpha P_j[n-1]) + (R_i[n] - \alpha R_i[n-1]). \quad (5)$$

Substituting (3) and (4) into (5):

$$U_i[n] = \alpha U_i[n-1] + \sum_j W_{ij} S_j[n-1] - \rho_i[n]. \quad (6)$$

Equation (6) has the same form as U_i of (1). The difference is that (2) allows us to compute $\frac{\partial U_i}{\partial W_{ij}}$ as the pre-synaptic traces without gradient backpropagation, which makes it compatible with DECOLLE:

$$\frac{\partial U_i}{\partial W_{ij}} = P_j[n], \quad (7)$$

Hence the loss gradient with respect to the weight $\frac{\partial \mathcal{L}}{\partial W_{ij}}$ is:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial W_{ij}} &= \frac{\partial \mathcal{L}}{\partial S_i} \frac{\partial S_i}{\partial U_i} \frac{\partial U_i}{\partial W_{ij}} \\ &= E_i[n] \sigma'(U_i[n]) P_j[n],\end{aligned}\quad (8)$$

where $E_i[n]$ is the propagated loss gradient to the neuron, and σ' is the surrogate gradient function. This neuron model allows us to use (2) during training, and (1) during inference to gain the benefits of reduced memory usage and computation. Furthermore, since the reset mechanism of (1) can be implemented by simply setting $U_i = 0$, the state variable ρ_i is not actually needed during inference, further simplifying the computation. Overall, only one state variable, the internal membrane potential U_i , is needed during inference.

3 STATE AND WEIGHT QUANTIZATION

Low precision arithmetic can bring a number of benefits to SNN running on a neuromorphic hardware. Synaptic weight quantization can reduce the memory usage and storage space required, allowing for a higher throughput given the same memory bandwidth. If both the neuron state variables and synaptic weights are quantized to low precision fixed-point representation, simpler digital circuits with less silicon can be used. This allows more neurons to be simulated in parallel within the same silicon and power budget, increasing the overall throughput and energy efficiency.

We used the QPyTorch [17] library to simulate the effect of low precision fixed-point quantization with various bit-precision settings for both the neuron state variables and synaptic weights. Figure 1 shows how the quantization was performed on each layer of the SNN. The quantizer function Q shown in figure 1 used the low precision fixed-point quantizer provided by QPyTorch. It performed simulated quantization by rounding or clipping a value to the desired bit-precision, but still used the full precision PyTorch tensor to hold the quantized values. The quantizers used stochastic rounding [12], with the probability of rounding up or down depending on the distance between the true value to the two nearest

quantized values as follows,

$$Q(x) = \begin{cases} \lfloor x \rfloor, & \text{with probability } 1 - \frac{x - \lfloor x \rfloor}{\epsilon} \\ \lfloor x \rfloor + \epsilon, & \text{with probability } \frac{x - \lfloor x \rfloor}{\epsilon} \end{cases}, \quad (9)$$

where ϵ was the smallest represented value of a given bit-precision, and $\lfloor x \rfloor$ is the rounded-down value. This allowed the discarded precision to be represented by a stochastic process. We also enabled value clamping, whereby values outside of the represented range were clamped to the maximum or minimum represented value. For this, the backward loss gradient was zeroed if the forward value was clamped. Without gradient zeroing, there would be a mismatch between the loss gradient and the forward clamped value. Finally, due to the reduced precision of the weights, a gradient scaling factor of $1e3$ was applied to magnify the step size of the weight updates. This avoided the updates being much smaller than ϵ of a given bit-precision.

Figure 1: State diagram of a quantized SNN layer (grouped by dashed box) for one time step. Q quantizes the state variables P , U , R , and the weights W .

4 HESSIAN TRACE FOR LAYER-WISE BIT PRECISION QUANTIZATION

The Hessian matrix is the second order differential of the loss with respect to the weights, and represents the sensitivity of the loss to any perturbation of the weights. Suppose that a fully converged network is at a global or local minimum of the loss space, the gradient at this point should be close to zero. Therefore the gradient itself does not provide information on how weight quantization affects the loss. However, if the converged network has a Hessian matrix with very large eigenvalues, any small perturbation of the weights would result in large loss gradients, and hence the loss itself. Therefore, the relative magnitude of the Hessian of each layer can tell us how the numerical errors introduced by weight quantization at that layer would lower the accuracy of the network. We can use this information to assign an optimal bit-precision to each layer according to each layer's Hessian.

Except for very small networks, computing the entire Hessian matrix is prohibitive. However, for the purpose of mixed precision quantization, the Hessian *trace* has been found to be sufficient [4].

The Hessian trace is the sum of all eigenvalues of the Hessian matrix, and can be evaluated by randomized numerical linear algebra methods. In particular, the Hutchinson algorithm [1] enables fast estimations of the trace using only the Hessian-vector product. For any symmetric matrix $H \in R^{d \times d}$, and a random vector $v \in R^d$ with i.i.d components sampled from a Gaussian distribution ($\mathcal{N}(0, 1)$), the trace of the matrix can be computed by the expectation of vector-matrix-vector product as follows,

$$Tr(H) = \mathbb{E}[v^T H v]. \quad (10)$$

Algorithm 1 shows the Hessian trace computation adapted for our SNN model. Note that the Hessian matrix is never computed explicitly, and the PyTorch autograd function is used to evaluate the Hessian-vector product Hv for each batch of the data and for each time step.

Algorithm 1 Computing the Hessian Trace of an SNN layer

Require: Layer parameters W_i

```

1:  $vHv \leftarrow []$ 
2: for  $i \leftarrow 1$  to  $MaxIter$  do
3:   Initialize a random vector  $v$  with same shape as  $W_i$ 
4:   Initialize a zero vector  $z$  with same shape as  $W_i$ 
5:   Normalize  $v$ ,  $v = \frac{v}{\|v\|_2}$ 
6:   for  $b \leftarrow 1$  to  $MaxBatch$  do
7:     for  $n \leftarrow 1$  to  $MaxSeq$  do
8:       Compute  $g_i = \frac{\partial \mathcal{L}}{\partial W_i}$  by autograd
9:       Compute  $Hv = \frac{\partial (g_i^T v)}{\partial W_i}$  by autograd
10:      Accumulate  $z \leftarrow z + Hv$ 
11:    end for
12:  end for
13:  Append  $v^T z$  to  $vHv$ 
14: end for
15:  $trace \leftarrow mean(vHv)$ 
16:
17: return  $trace$ 

```

5 EXPERIMENTS AND RESULTS

We performed all experiments on the N-MNIST [13] hand written digit recognition dataset that was recorded with an event based camera [11]. We used the same network architecture as the original DECALLE network [10], which consisted of 3 convolutional layers with filter size 7×7 and 64, 128, 128 channels respectively. Instead of the original neuron model with 5 state variables, the simplified model of (2) was used. All network configurations and hyper-parameters were kept the same, except for $lr = 1e6$ and $\alpha = 0.97$ as they were more suitable for the simplified neuron model. The same train-test split of the dataset was used, with 2000 samples in the training set and 100 samples in the testing set for each of the 10 classes. We first trained a network with full precision without any quantization for 20 epochs. The Hessian trace for each layer was then computed using algorithm 1. We then fine-tuned the same network with quantization as shown in figure 1 for another 10 epochs, using a layer-wise bit-precision for each layer. The experiments were repeated 3 times with different weight initializations, and the mean and standard deviation of the accuracy were reported.

Table 1 shows the Hessian trace of each layer after 20 epochs of full precision training. The Hessian trace increased from L1 to L3, with 4 orders of magnitude difference. However, there was only 1 order of magnitude difference between L2 and L3, indicating relatively equal sensitivity to weight perturbation between these two layers. This increasing trend of Hessian trace is compatible with the trend of the accuracy of each layer’s local classifier, which also increased from L1 to L3.

Table 1: Hessian Trace and accuracy of different SNN layers

	L1	L2	L3
Trace	3.57e2	2.86e5	3.45e6
Accuracy	47.8	95.0	98.2

Table 2 shows the accuracy of the final layer after fine tuning the network with layer-wise bit-precision quantization, and the corresponding size of the network parameters after the quantization. One surprising observation is that assigning a 16-bit-precision for each layer did not lower the accuracy at all, while reducing the network size by 50%. This shows that full precision arithmetic is not needed. As expected, the accuracy dropped when lowering the bit-precision of all 3 layers from 16 to 8 and 4 bits, from 98.1% to 96.3% and 81.1% respectively. We also compared the accuracy when using a layer-wise bit-precision, with a (16, 8) and (8, 4) bit combination. As indicated by the Hessian trace shown on table 1, we expected assigning L1 with a low bit-precision to have the least impact to accuracy, followed by L2 and L3. This was confirmed by the results, with the accuracy dropping from 98.0% to 97.9% and 97.6% as we moved an 8-bit quantizer from L1 to L3 in a 16-bit network. The same observation was found when moving a 4-bit quantizer from L1 to L3 in an 8-bit network, with the accuracy dropping from 95.0% to 94.0% and 93.6%. Finally we quantized the first two layers of the network with lower bit-precision, and found that the (8, 8, 16) combination provided a good balance between accuracy and network size, with the accuracy dropping by only 0.3% while the network size was reduced by 58%.

6 CONCLUSION AND FUTURE WORK

In this work we presented a method of quantizing SNNs for high throughput, energy efficient inference. We used a simplified LIF neuron model that has two equivalent forms, with one form being compatible with gradient based DECOLLE training, and another form suitable for inference with reduced computation and memory usage. We also applied Hessian aware quantization on SNN and used the layer-wise Hessian trace to evaluate the sensitivity of the loss to the quantization of each layer’s weights. We found that the Hessian trace increased 4-fold from layer L1 to L3, and confirmed that this provided good information to determine the optimal layer-wise bit-precision for quantization. Our results show that when assigning a low bit-precision to a layer with low Hessian trace, it had a smaller impact on accuracy than reducing the bit-precision of a layer with a high Hessian trace.

There are a number of limitations of the present work. The network only had 3 layers and could not take full advantage of the Hessian aware quantization method. However, this is the limitation

Table 2: Accuracy of the final SNN layer after fine tuning

L1	L2	L3	Accuracy	Size (MB)
FP32	FP32	FP32	98.1 ± 0.2	4.84
16	16	16	98.1 ± 0.2	2.42
8	16	16	98.0 ± 0.2	2.41
16	8	16	97.9 ± 0.1	2.02
16	16	8	97.6 ± 0.2	1.62
8	8	16	97.8 ± 0.2	2.01
8	8	8	96.3 ± 0.3	1.21
4	8	8	95.0 ± 0.5	1.21
8	4	8	94.0 ± 0.9	1.01
8	8	4	93.6 ± 0.8	0.81
4	4	8	92.6 ± 1.1	1.01
4	4	4	81.1 ± 1.0	0.61

of SNNs in general, as having many layers has not been shown to result in a higher accuracy using the available benchmarks [10]. After the Hessian trace analysis, the bit-precision allocation step was performed manually, due to the small number of layers available, as opposed to the Pareto Frontier method used in [4]. Finally, the present work was performed with simulated quantization provided by QPyTorch, and the underlying computation was still performed in FP32. Future work should explore using Hessian aware quantization of SNNs with native fixed-point arithmetic on a neuromorphic hardware to more accurately assess the trade off between bit-precision, compute time, memory usage and energy efficiency.

ACKNOWLEDGMENTS

This work was supported by the Korean Institute for Science and Technology (EN), the National Science Foundation under grant 1652159 (EN) and 1823366 (HL, EN).

REFERENCES

- [1] Haim Avron and Sivan Toledo. 2011. Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix. *J. ACM* (2011). <https://doi.org/10.1145/1944345.1944349>
- [2] Thomas Bartol, Cailey Bromer, Justin Kinney, Michael Chirillo, Jennifer Bourne, Kristen Harris, and Terrence Sejnowski. 2015. Hippocampal Spine Head Sizes are Highly Precise. *bioRxiv* (2015). <https://doi.org/10.1101/016329>
- [3] M. Davies, N. Srinivasa, T. H. Lin, G. Chinya, P. Joshi, A. Lines, A. Wild, and H. Wang. 2018. Loihi: A Neuromorphic Manycore Processor with On-Chip Learning. *IEEE Micro PP*, 99 (2018), 1–1. <https://doi.org/10.1109/MM.2018.112130359>
- [4] Zhen Dong, Zhewei Yao, Yaohui Cai, Daiyaan Arfeen, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. 2019. HAWQ-V2: Hessian aware trace-weighted quantization of neural networks.
- [5] Steven K. Esser, Jeffrey L. McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S. Modha. 2019. Learned step size quantization.
- [6] Steven K Esser, Paul A Merolla, John V Arthur, Andrew S Cassidy, Rathinakumar Appuswamy, Alexander Andreopoulos, David J Berg, Jeffrey L McKinstry, Timothy Melano, Davis R Barch, et al. 2016. Convolutional Networks for Fast, Energy-Efficient Neuromorphic Computing. *PNAS* 113 (2016), 11441–11446. <https://doi.org/10.1073/pnas.1604850113>
- [7] Wulfram Gerstner, Werner M. Kistler, Richard Naud, and Liam Paninski. 2014. *Neuronal dynamics: From single neurons to networks and models of cognition*. <https://doi.org/10.1017/CBO9781107447615>
- [8] Wulfram Gerstner, Werner M Kistler, Richard Naud, and Liam Paninski. 2014. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press.

- [9] Giacomo Indiveri, Bernabé Linares-Barranco, Tara Julia Hamilton, André van Schaik, Ralph Etienne-Cummings, Tobi Delbruck, Shih Chii Liu, Piotr Dudek, Philipp Häfliger, Sylvie Renaud, Johannes Schemmel, Gert Cauwenberghs, John Arthur, Kai Hynna, Fopefolu Folowosele, Sylvain Saighi, Teresa Serrano-Gotarredona, Jayawan Wijekoon, Yingxue Wang, and Kwabena Boahen. 2011. Neuromorphic silicon neuron circuits. <https://doi.org/10.3389/fnins.2011.00073>
- [10] Jacques Kaiser, Hesham Mostafa, and Emre Neftci. 2020. Synaptic Plasticity Dynamics for Deep Continuous Local Learning (DECOLLE). *Frontiers in Neuroscience* (2020). <https://doi.org/10.3389/fnins.2020.00424>
- [11] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. 2008. A 128×128 120 dB 15 μ s latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits* (2008). <https://doi.org/10.1109/JSSC.2007.914337>
- [12] Lorenz K. Muller and Giacomo Indiveri. 2015. Rounding Methods for Neural Networks with Low Resolution Synaptic Weights. (4 2015). <http://arxiv.org/abs/1504.05767>
- [13] Garrick Orchard, Ajinkya Jayawant, Gregory K. Cohen, and Nitish Thakor. 2015. Converting Static Image Datasets to Spiking Neuromorphic Datasets Using Saccades. *Frontiers in Neuroscience* 9 (nov 2015). <https://doi.org/10.3389/fnins.2015.00437>
- [14] Michael Pfeiffer and Thomas Pfeil. 2018. Deep Learning With Spiking Neurons: Opportunities and Challenges. *Frontiers in Neuroscience* (2018). <https://doi.org/10.3389/fnins.2018.00774>
- [15] Clemens J.S. Schaefer and Siddharth Joshi. 2020. Quantizing Spiking Neural Networks with Integers. In *ACM International Conference Proceeding Series*. <https://doi.org/10.1145/3407197.3407203>
- [16] Friedemann Zenke and Emre O. Neftci. 2021. Brain-Inspired Learning on Neuromorphic Substrates. *Proc. IEEE* (2021). <https://doi.org/10.1109/JPROC.2020.3045625>
- [17] Tianyi Zhang, Zhiqiu Lin, Guandao Yang, and Christopher de Sa. 2019. QPyTorch: A low-precision arithmetic simulation framework.