Supplement: Band-limited Training and Inference for Convolutional Neural Networks

1. Implementation Details

We present details on the map reuse, CUDA implementation and shifting of the DC coefficient.

1.1. Map Reuse

We divide the input map M (with half of the map already removed due to the conjugate symmetry) into two parts: upper D1 and lower D2. We crop out the top-left (S1) corner from D1 and bottom-left (S2) corner from D2. The two compressed representations S1 and S2 can be maintained separately (small saving in computation time) or concatenated (more convenient) for the backward pass. In the backward pass, we pad the two corners S1 and S2 to their initial sizes D1 and D2, respectively. Finally, we concatenate D1 and D2 to get the FFT map M', where the high frequency coefficients are replaced with zeros.

If the memory usage should be decreased as much as possible and the filter is small, we can trade the lower memory usage for the longer computation time and save the filter in the spatial domain at the end of the forward pass, followed by the FFT re-computation of the filter in the backward pass. The full frequency representation of the input map (after padding) is bigger than its spatial representation, thus the profitability of re-computing the input to save the GPU memory depends on the applied compression rate.

We also contribute a fast shift of the DC coefficients either to the center or to the top-left corner. The code for the element-wise solution uses two for loops and copy each element separately. For the full FFT map, we divide it into quadrants (I - top-right, II - top-left, III - bottom-left, IV - bottom-right). Then, we permute the quadrants in the following way: I $\rightarrow$ III, II $\rightarrow$ IV, III $\rightarrow$ I, IV $\rightarrow$ II.

1.2. CUDA

We use *min(max threads in block*, n^2*)* threads per block and the total number of GPU blocks is Sf', where S is the mini-batch size, f' is the number of output channels, and n is the height and width of the inputs. Each block of threads is used to compute a single output plane. Intuitively, each thread in a block of threads incrementally executes a complex multiplication and sums the result to an aggregate for all f input channels to obtain a single output cell (x, y).

Additional optimizations, such as maintaining the filters only in the frequency domain or tiling, will be implemented in our future work.

2. Experiments

2.1. Experimental Setup

For the experiments with ResNet-18 on CIFAR-10 and DenseNet-121 on CIFAR-100, we use a single instance of P-100 GPU with 16GBs of memory.

We also use data from the UCR archive, with the main representative: 50 words time-series dataset with 270 values per data point, 50 classes, 450 train data points, 455 test data points, 2 MB in size. One of the best peforming CNN models for the data is a 3 layer Fully Convolutional Neural Network (FCN) with filter sizes: 8, 5, 3. The number of filter banks is: 128, 256, 128. [1].

Our methodology is to measure the memory usage on GPU by counting the size of the allocated tensors. The direct measurement of hardware counters is imprecise because PyTorch uses a caching memory allocator to speed up memory allocations and incurs much higher memory usage than is actually needed at a given point in time.

2.2. DenseNet-121 on CIFAR-100

We train DenseNet-121 (with growth rate 12) on the CIFAR-100 dataset.

In Figure 1 we show small differences in test accuracy during training between models with different levels of energy preserved for the FFT-based convolution.

In Figure 2 we show small differences in accuracy and loss between models with different convolution implementations. The results were normalized with respect to the values obtained for the standard convolution used in PyTorch.

2.3. Reduced Precision and Bandlimited Training

In Figure 4 we plot the maximum allocation of the GPU memory during 3 first iterations. Each iteration consists of training (forward and backward passes) followed by test-

[1] http://bit.ly/2FbdQNV

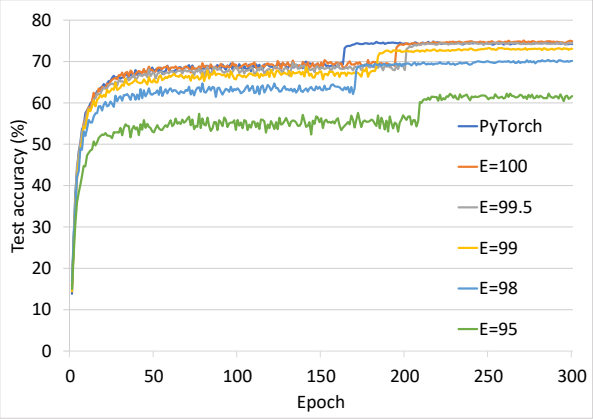

Figure 1. Comparing test accuracy during training for CIFAR-100 dataset trained on DenseNet-121 (growth rate 12) architecture using convolution from PyTorch and FFT-based convolutions with different energy rates preserved.

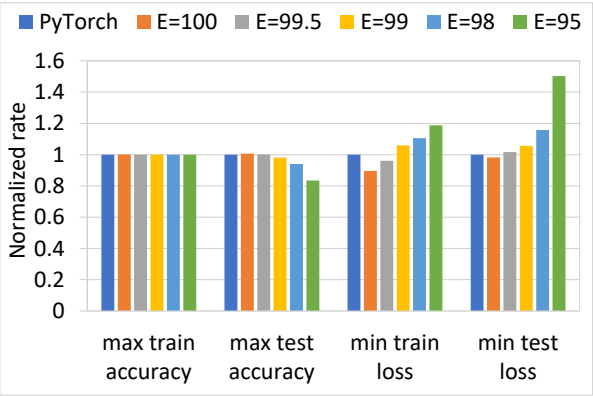

Figure 2. Comparing accuracy and loss for test and train sets from CIFAR-100 dataset trained on DenseNet-121 (growth rate 12) architecture using convolution from PyTorch and FFT-based convolutions with different energy rates preserved.

ing (a single forward pass). We use CIFAR-10 data on ResNet-18 architecture. We show the memory profiles of RPA (Reduced Precision Arithmetic), bandlimited training, and applying both. A detailed convergence graph is shown in Figure 3.

2.4. Resource Usage vs Accuracy

The full changes in normalized resource usage (GPU memory or time for a single epoch) vs accuracy are plotted in Figure 5.

2.5. Dynamic Changes of Compression

Deep neural networks can better learn the model if the compression is fixed and does not change with each iteration depending on the distribution of the energy within the frequency coefficients of a signal.

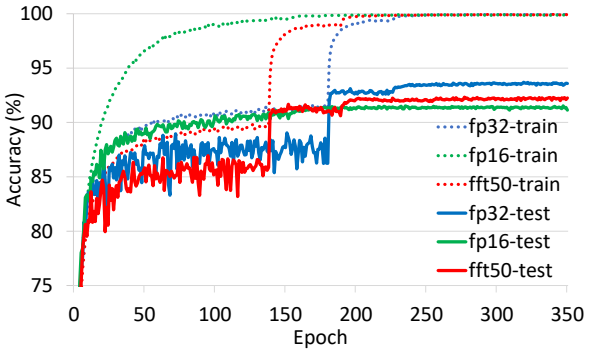

Figure 3. Train and test accuracy during training for CIFAR-10 dataset trained on ResNet-18 architecture using convolution from PyTorch (fp32), mixed-precision (fp16) and FFT-based convolutions with 50% of compression for intermediate results and filters (fft50). The highest test accuracy observed are: 93.69 (fp32), 91.53 (fp16), 92.32 (fft50).

We observe that the compression can be applied more effectively to the first layers and the deeper the layers the less compression can be applied (for a given energy level preserved).

The dynamic and static compression methods can be combined. We determine how much compression should be applied to each layer via the energy level required to be saved in each layer and use the result to set the static compression for the full training. The sparsification in the Winograd domain requires us to train a full (uncompressed) model, then inspect the Winograd coefficients of the filters and input maps and zero-out these of them which are the smallest with respect to their absolute values, and finally retrain the compressed model. In our approach, we can find the required number of coefficients to be discarded with a few forward passes (instead of training the full network), which can save time and also enables us to utilize less GPU memory from the very beginning with the dynamic compression.

2.6. Compression Based on Preserved Energy

There are a few ways to compress signals in the frequency domain for 2D data. The version of the output in the frequency domain can be compressed by setting the DC component in the top left corner in the frequency representation of an image or a filter (with the absolute values of coefficients decreasing towards the center from all its corners) and then slicing off rows and columns. The heat maps of such a representation containing the absolute value of the coefficients is shown in Figure 6.

The number of preserved elements even for 99% of the preserved energy is usually small (from 2X to 4X smaller than

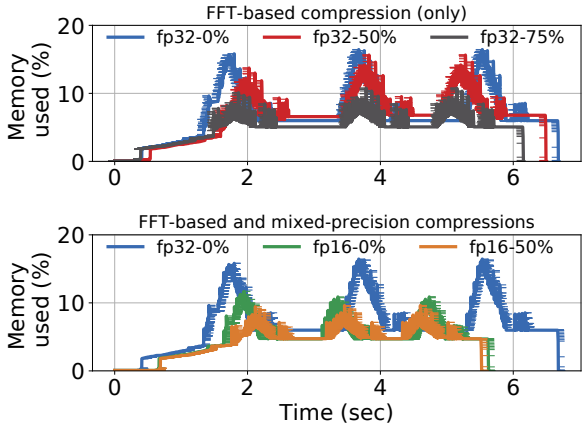

Figure 4. Memory used (%) for the first 3 iterations (train and test) with mixed-precision and FFT-based compression techniques. Mixed precision allows only a certain level of compression whereas with the FFT based compression we can adjust the required compression and accuracy. The two methods can be combined (fp16-50%).

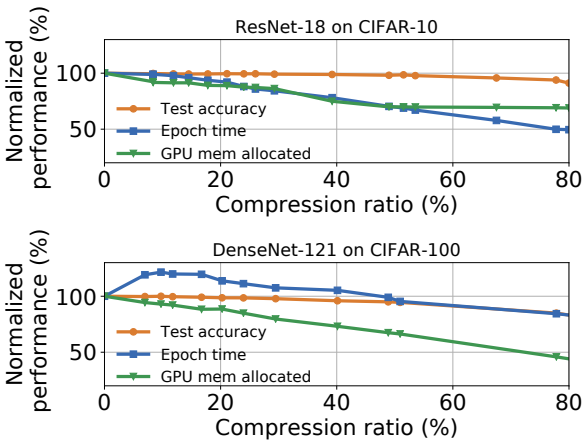

Figure 5. Normalized performance (%) between models trained with different FFT-compression ratios.

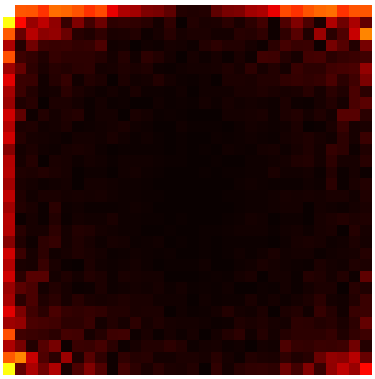

Figure 6. A heat map of absolute values (magnitudes) of FFT coefficients with linear interpolation and the max value colored with white and the min value colored with black. The FFT-ed input is a single (0-th) channel of a randomly selected image from the CIFAR-10 dataset.

the initial input). Thus, for the energy based compression, we usually proceed starting from the DC component and then adding rows and columns in the vertically mirrored L fashion. It can be done coarse-grained, where we just take into account the energy of the new part of row or column to be added, or fine-grained, where we add elements one by one and if not the whole row or column is needed, we zero-out the remaining elements of both an activation map and a filter.

2.7. Visualization of the Compression in 1D

We present the visualization of our FFT-based compression method in 7. The magnitude is conveniently plotted in a logarithmic scale (dB).

2.8. Energy Based Compression for ResNet-18

Figure 8 shows the linear correlation between the accuracy of a model and the energy that was preserved in the model during training and testing. Each point in the graph requires a fool training of a model for the indicated energy level preserved.

Figure 9 shows the test accuracy during the training process of the ResNet-18 model on the CIFAR-10 dataset.

Figure 10 shows the train accuracy during the training process of the ResNet-18 model on the CIFAR-10 dataset.

2.9. Training vs. Inference Bandlimiting

To further corroborate our points, consider a scheme where we train the network with one compression ratio and test with another (Figure 19).

We observe that the network is most accurate when the compression used for training is the same that is used during testing. We used the Friedman statistical test followed by the post-hoc Nemenyi test to assess the performance of multiple compression ratios during inference over multiple datasets. Figure 14 shows the average rank of the test accuracies of different compression ratios during inference across 25 randomly chosen time-series data from the UCR Archive. The training was done while preserving 90% of the energy. Inference with the same compression ratio (90%) is ranked first, meaning that it performed the best in the majority of the datasets. The Friedman test rejects the null hypothesis that all measures behave similarly, and, hence, we proceed with a post-hoc Nemenyi test, to evaluate the significance of the differences in the ranks. The wiggly line in the figure connects all approaches that do not perform statistically

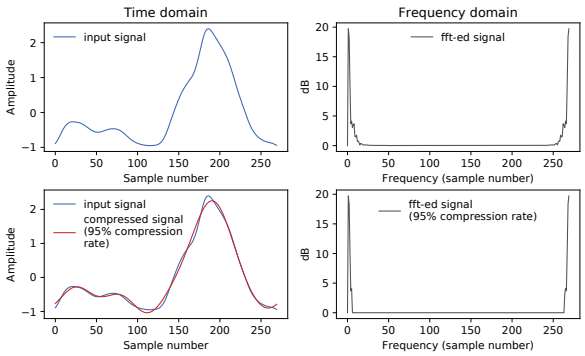

Figure 7. We present a time series (signal) from the UCR archive and fifty words dataset in the top-left quadrant. Its frequency representation (as power spectrum) after normalized FFT transformation is shown in the top-right quadrant. The signal is compressed by 95% (we zero out the middle Fourier coefficients) and presented in the bottom-right quadrant. We compare the initial signal and its compressed version in the bottom-left quadrant. The magnitudes of Fourier coefficients are presented in the logarithmic (dB) scale.

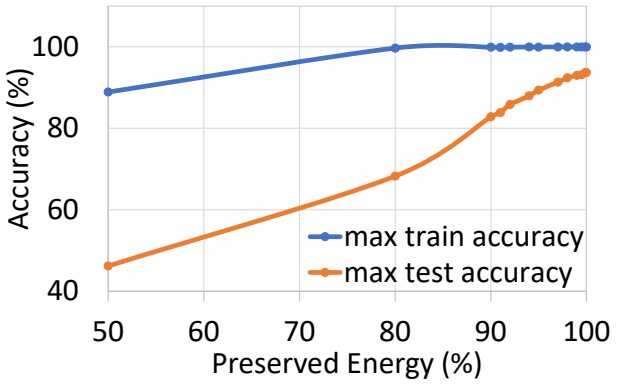

Figure 8. The linear correlation between the accuracy of a model and the energy that was preserved in the model during training and testing.

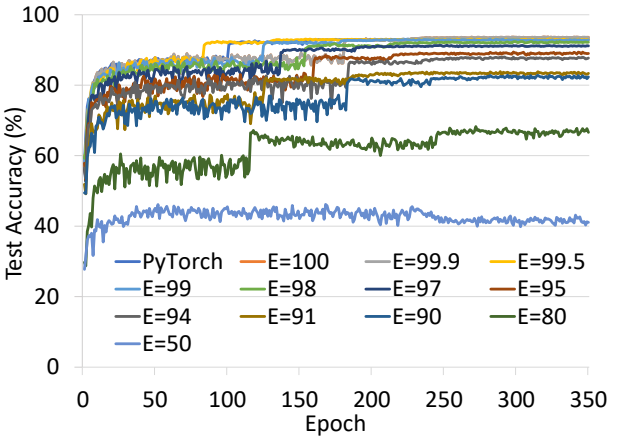

Figure 9. The test accuracy during the training process of the ResNet-18 model on the CIFAR-10 dataset.

differently according to the Nemenyi test. We had similar findings when training was done using no compression but compression was later applied during inference (see Figure 15). In other words, the network *learns how to best leverage a band-limited operation to make its predictions*.

Even so its performance degrades gracefully for tests with the compression level further from the one used during training. In our opinions, the smooth degradation in performance is a valuable property of band-limiting. An outer optimization loop can tune this parameter without worrying about training or testing instability.

2.10. Error Incurred by 2D Convolution with Compression

We tried to measure how accurate the computation of the convolution result is when the compression is applied. An image from CIFAR-10 dataset (3x32x32) was selected and an initial version of a single filter (3x5x5, Glorot initialization). We did convolution using PyTorch, executed our convolution with compression for different compression ratios, and compared the results. The compression was measured relatively to the execution of our FFT-based convolution without any compression (100% of the energy of the input image is preserved). The results show that for 2D convolution the relative error is already high (about 22.07%) for a single index discarded (the smallest possible compression of about 6%). However, after the initial abrupt change we observe a linear dependence between compression ratio and relative error until more than about 95% of compression ratio, after which we observe a fast degradation of the result.

We plot in Figure16 fine-grained compression using the top method (the coefficients with lowest values are zeroed-out first). For a given image, we compute its FFT and its spectrum. For a specified number k of elements to be zeroed-out, we find the k smallest elements in the spectrum and zero-out the corresponding elements in the in the image. The same procedure is applied to the filter. Then we compute the 2D convolution between the compressed filter and the image. We do not remove the elements from the tensors (the sizes of the tensors remain the same, only the smallest coefficients are zeroed-out). The plots of the errors for a given compression (rate of zeroed-out coefficients) are relatively smooth. This shows that our method to discard coefficients and decrease the tensor size is rather coarse-grained and especially for the first step, we remove many elements.

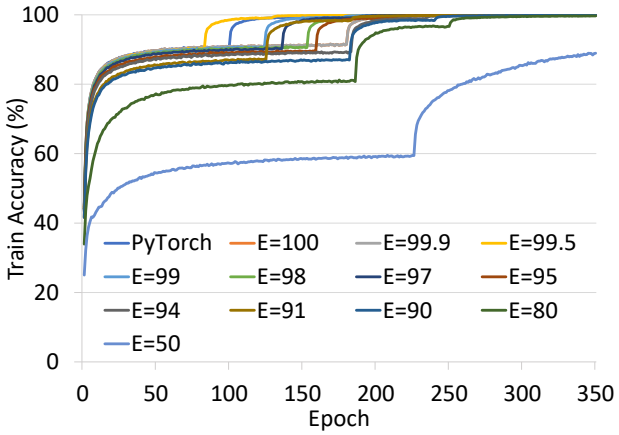

Figure 10. The train accuracy during the training process of the ResNet-18 model on the CIFAR-10 dataset.

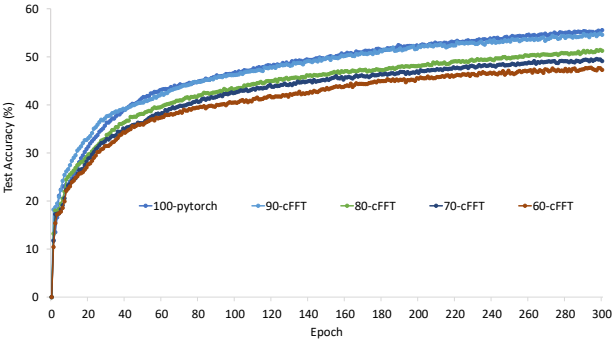

Figure 11. A comparison of 2D convolution operation implemented in PyTorch and FFT version for different percentage of preserved energy (on the level of a batch).

We have an input image from CIFAR-10 with dimensions (3x32x32) which is FFT-ed (with required padding) to tensor of size (3x59x30). We plot the graph in 10 element zero-out step, i.e. first we zero-out 10 elements, then 20, and so on until we reach 5310 total elements in the FFT-ed tensors). The compression ratio is computed as the number of zeroed-out elements to the total number of elements in FFT-ed tensor. There are some dips in the graph, this might be because the zeroed-out value is closer to the expected value than the one computed with imprecise inputs. With this fine-grained approach, after we zero-out a single smallest coefficients (in both filter and image), the relative error from the convolution operation is only 0.001%. For the compression ratio of about 6.61%, we observe the relative error of about 8.41%. In the previous result, we used the lead method and after discarding about 6.6% of coefficients, the relative error was 22.07%. For the lead method, we were discarding the whole rows and columns across all channels.

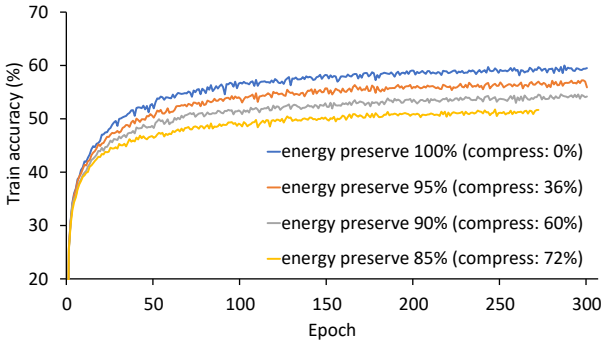

Figure 12. Train accuracy for CIFAR-10 dataset on LeNet (2 conv layers) architecture Momentum 0.9, batch size 64, learning rate 0.001 .

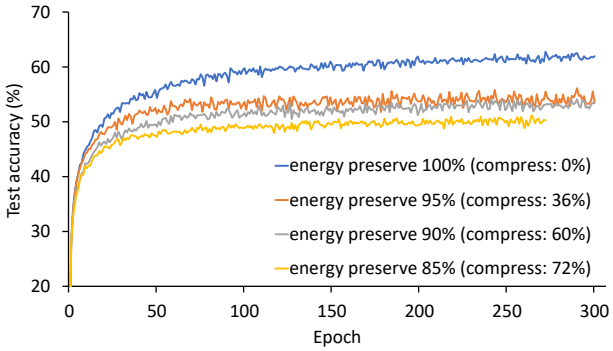

Figure 13. Test accuracy for CIFAR-10 dataset on LeNet architecture (2 conv layers, momentum 0.9, batch size 64, learning rate 0.001.

For the fine-grained method, we select the smallest elements within the whole tensor.

2.11. Time-series data

We show the accuracy loss of less than 1% for 4X less average GPU memory utilization (Figures: 17 and 18) when training FCN model on 50 words time-series dataset from the UCR archive.

In Figure 19 we show the training compression vs. inference compression for time-series data. This time we change the compression method from static to the energy based, however, the trend remains the same. The highest test accuracy is achieved by the model with the same energy preserved during training and testing.

2.12. Robustness to Adversarial Examples

We present the most relevant adversarial methods that were executed using the foolbox library. Our method is robust to

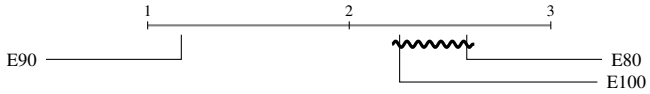

Figure 14. Ranking of different compression ratios (80%, 90%, and 100% energy preserved) during inference with model trained using no compression (90% of energy preserved)

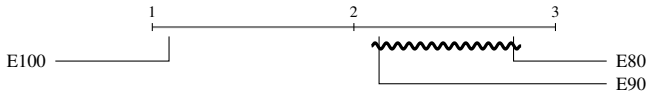

Figure 15. Ranking of different compression ratios (80%, 90%, and 100% energy preserved) during inference with model trained using no compression (100% of energy preserved)

decision-based attacks (GaussianBlur, Additive Uniform or Gaussian Noise) but not to the gradient-based (white-box and adaptive) attacks (e.g., Carlini & Wagner or FGSM) since we return proper gradients in the band-limited convolutions. If an adversary is not aware of our band-limiting defense, then we can recover the correct label for many of the adversarial examples by applying the FFT compression to the input images.

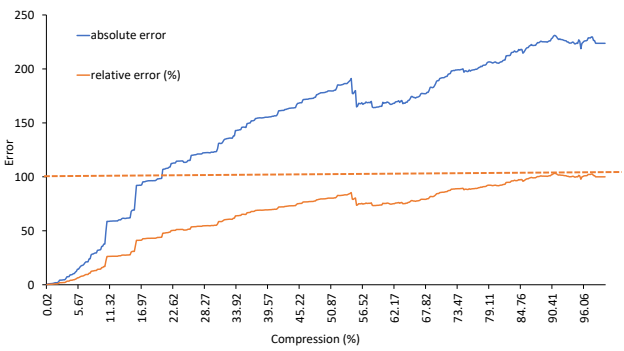

Figure 16. A comparison of the relative (in %) and absolute errors between 2D convolution from PyTorch (which is our gold standard with high numeric accuracy) and a fine-grained top compression method for a CIFAR-10 image and a 5x5 filter (with 3 channels).

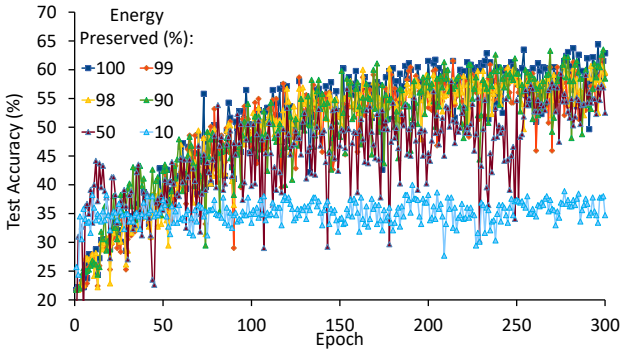

Figure 17. Test accuracy on a 3 layer FCN architecture for 50 words time-series dataset from the UCR archive.

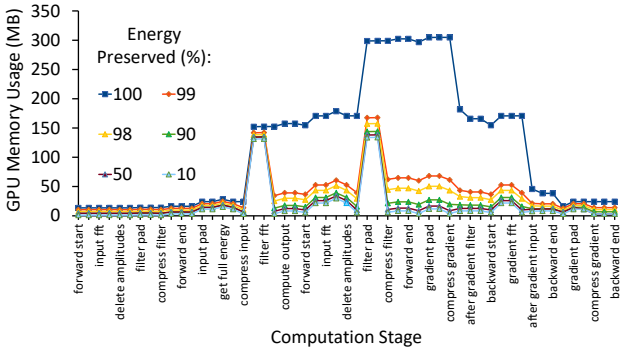

Figure 18. GPU memory usage (in MB) during training for a single forward and backward pass through the FCN network using 50 words dataset.

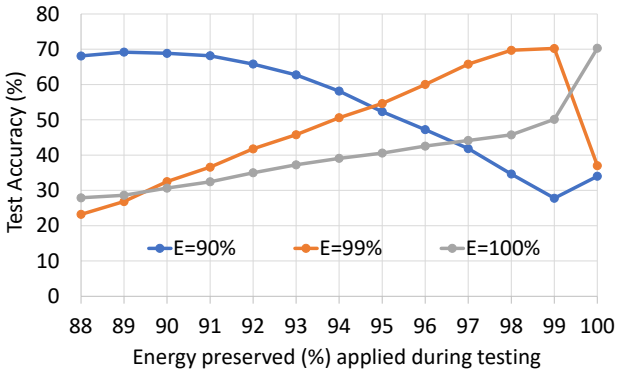

Figure 19. We train three models on the time-series dataset uWaveGestureLibrary_Z. The preserved energy during training for each of the models is 90%, 99% and 100% (denoted as E=X% in the legend of the graph). Next, we test each model with energy preserved levels ranging from 88% to 100%. We observe that the highest accuracy during testing is for the same energy preserved level as the one used for training and the accuracy degrades smoothly for higher or lower levels of energy preserved.

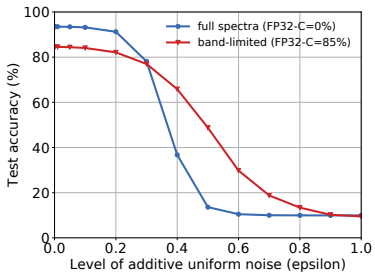

Figure 20. Input test images are perturbed with additive uniform noise, where the epsilon parameter is changed from 0 to 1. The more band-limited model, the more robust it is to the introduced noise. We use ResNet-18 models trained on CIFAR-10.