

# Rank and run-time aware compression of NLP Applications

**Urmish Thakker**  
SambaNova Systems  
uthakker@cs.wisc.edu

**Jesse Beu**  
Arm ML Research

**Dibakar Gope**  
Arm ML Research

**Ganesh Dasika**  
AMD Research

**Matthew Mattina**  
Arm ML Research

## Abstract

Sequence model based NLP applications can be large. Yet, many applications that benefit from them run on small devices with very limited compute and storage capabilities, while still having run-time constraints. As a result, there is a need for a compression technique that can achieve significant compression without negatively impacting inference run-time and task accuracy. This paper proposes a new compression technique called Hybrid Matrix Factorization that achieves this dual objective. HLF improves low-rank matrix factorization (LMF) techniques by doubling the rank of the matrix using an intelligent hybrid-structure leading to better accuracy than LMF. Further, by preserving dense matrices, it leads to faster inference run-time than pruning or structure matrix based compression technique. We evaluate the impact of this technique on 5 NLP benchmarks across multiple tasks (Translation, Intent Detection, Language Modeling) and show that for similar accuracy values and compression factors, HLF can achieve more than  $2.32\times$  faster inference run-time than pruning and 16.77% better accuracy than LMF.

## 1 Introduction

Sequence based (LSTMs/GRUs) NLP Applications are being increasingly run on mobile phones and smart watches. They are typically enabled by querying a cloud-based system to do most of the computation. The energy, latency, and privacy implications associated with running a query on the cloud is changing where users run a neural network application. We should, therefore, expect an increase in the number of NLP applications running on embedded devices. Due to the energy and power constraints of edge devices, embedded SoCs frequently use lower-bandwidth memory technologies and smaller caches compared to desktop and server

processors. Thus, there is a need for good compression techniques to enable large NLP models to fit into a smaller edge device or ensure that they run efficiently on devices with smaller caches (Thakker et al., 2019b). Additionally, compressing models should not negatively impact the inference run-time as these tasks may have real-time deadlines to provide a good user experience.

In order to choose a compression scheme for a particular network, one needs to consider 3 different axes – the compression factor, the inference run-time speedup over the baseline, and the accuracy. Ideally, a good compression algorithm should not sacrifice improvement along one axis for improvement along another. For example, network pruning (Han et al., 2016) has shown to be an effective compression technique, but pruning creates a sparse matrix representation that is inefficient to execute on most modern CPUs. Our analysis shows that pruned networks can achieve a faster run-time than the baseline only for significantly high compression factors. Low-rank matrix factorization (LMF) is another popular compression technique that can achieve speedup proportional to the compression factor. However, LMF has had mixed results in maintaining model accuracy (Grachev et al., 2017; Chen et al., 2018; Lu et al., 2016). This is because LMF reduces the rank of a matrix significantly, reducing its expressibility (Yang et al., 2018). Lastly, structured matrices (Ding et al., 2018) can also be used to compress neural networks. While these techniques show a significant reduction in computation, this reduction only translates to a realized run-time improvement for large matrices (Thomas et al., 2018) or while using specialized hardware (Li et al., 2018; Sindhvani et al., 2015). For benchmarks evaluated in this paper, HLF gets  $30\times$  speed-up improvement over structured matrix based technique (Sindhvani et al., 2015).

Given LMF’s good run-time characteristics, it can potentially act as an alternative to pruning. However, LMF leads to an accuracy loss. **To overcome the problem of finding an alternative to pruning, which preserves the run-time benefit of dense structures of LMF and the accuracy benefits of pruned networks, we introduce a new compression technique called Hybrid Matrix Factorization (HLF). HLF can act as an effective compression technique for NLP edge use cases on embedded CPUs.** The results are very promising – HLF achieves iso-accuracy for a large compression factor ( $2\times$  to  $4\times$ ), improves the CPU run-time over pruning by a factor of  $2.32\times$  and can achieve 16.77% better model accuracy than LMF and 9% better accuracy than smaller baselines.

## 2 Related Work

**Pruning** (Han et al., 2016; Zhu and Gupta, 2017; Sanh et al., 2020) has been the most successful compression technique for all types of neural networks. Poor hardware characteristics of pruning has led to research in block based pruning technique (Narang et al., 2017). However, block based pruning technique also requires certain amount of block sparsity to achieve faster run-time than baseline. Having a strict compression factor requirement to get better run-time is a stringent constraint that HLF manages to avoid.

**Structured matrices** have shown significant potential for compression of NN (Sindhwani et al., 2015; Ding et al., 2018; Wang et al., 2018; Ding et al., 2017; Cheng et al., 2015; Thakker et al., 2020). Block circular compression is an extension of structured matrix based compression technique, converting every block in a matrix into structured matrix. We will show in this paper that HLF is a superior technique than block circular decomposition.

**Tensor decomposition** (CP decomposition, Kronecker, Tucker decomposition etc) based methods have also shown significant reduction in parameters (Tjandra et al., 2017; Thakker et al., 2019c,a). Matrix Factorization (Kuchaiev and Ginsburg, 2017; Chen et al., 2018; Grachev et al., 2017; Thakker et al., 2019) can be categorized under this topic. We will show in this paper that HLF can lead to better accuracy than LMF compressed RNNs.

**Quantization** is another popular technique for compression (Hubara et al., 2017, 2016; Gope et al., 2020a; Sanh et al., 2019; Liu et al., 2018; Gope

et al., 2020b, 2019). Networks compressed using HLF can be further compressed using quantization.

**Dynamic techniques** are used to improve inference run-time of RNNs by skipping certain RNN state updates (Campos et al., 2018; Seo et al., 2018; Yu et al., 2017; Tao et al., 2019). These techniques are based on the assumption that not all inputs to a RNN are needed for final classification task. Thus we can learn a small and fast predictor that can learn to skip certain inputs and its associated computation. HLF technique is orthogonal to this technique and networks compressed using HLF can be further optimized using this technique.

**Design of efficient structures** for LSTM/GRU cells like SRU (Lei et al., 2018), QRNN (Bradbury et al., 2016) and PRU (Mehta et al., 2018) have also led to networks with faster inference run-time benefits or lesser number of parameters. These structures are different from structured matrices and are hand-crafted after better understanding the application domain. HLF can be further used to optimize the matrices in these architectures to make the resultant network more parameter and run-time efficient.

Finally, any technique used to reduce the parameter footprint of embedding matrices in NLP can further optimize RNN networks optimized using HLF (Acharya et al., 2019; Mehta et al., 2019). In this paper, we show that HLF can compress networks with compressed word embedding layers.

## 3 Hybrid Matrix Factorization

### 3.1 Why LMF can potentially lead to loss in accuracy

LMF (Kuchaiev and Ginsburg, 2017) expresses a larger matrix  $A \in R^{m \times n}$  as a product of two smaller matrices  $U \in R^{m \times r}$  and  $V \in R^{r \times n}$ , respectively. Parameter  $r$  controls the compression factor. Unlike pruning, matrix factorization is able to improve the run-time over the baseline for most compression factors. Unfortunately, compression via LMF can lead to loss in accuracy. We believe, this is because of two closely related reasons:

- **Rank-Loss:** The rank of a matrix is a measure of the expressibility of a matrix. A lower rank matrix means less expressibility, limiting its learning capacity. This can potentially lead to some accuracy loss. LMF compression leads to a lower ranked matrix. While before compression, the rank of matrix  $A$  is  $\min(m, n)$ ,

after compression, it becomes  $\min(m, n, d)$ . Eg - If  $A \in R^{256 \times 256}$ , compression using LMF by a factor of 2 leads to  $U \in R^{256 \times 64}$  and  $V \in R^{64 \times 256}$ . The resultant compressed matrix  $A (= U * V)$  is a 64 rank matrix. Thus, in order to compress the matrix by a factor of 2, LMF reduces the rank of a matrix by a factor of 4.

- **Less expressive output features:** A closely related argument can be viewed when we extend the idea of low-rank matrix and its impact on the output features. Without loss of generality, an LSTM/GRU layer calculates a matrix-vector product during inference. If we assume the parameters of a LSTM/GRU layer are represented by a matrix  $A \in R^{m \times n}$  and the input to the matrix is  $x \in R^{n \times 1}$ , then the output feature calculated is -

$$y = f(A * x)$$

where,  $y \in R^{m \times 1}$

.  $f$  is a non-linear function. Thus, each element of  $y$  is a dot product of a row of  $A$  and the vector  $x$  followed by non-linearity. LMF expresses  $A$  in a lower dimensional space using the  $U$  and  $V$  matrix. If we rewrite the equation to calculate  $y$ , when  $A$  is expressed using LMF, we get -

$$y = f(U * V * x) \quad (1)$$

$$y = f(U * k) \quad (\text{assuming } k = V * x) \quad (2)$$

$$\text{where, } k \in R^{r \times 1} \quad (3)$$

Generally, for compression  $r < m, n$ . Thus,  $x \in R^{n \times 1}$  is projected to a lower dimensional embedding of size  $R^{d \times 1}$  and expanded again to  $R^{m \times 1}$  to create  $y$ . Thus, compressing  $A$  to a lower rank leads to output features calculated from a lower dimension embedding vector.

### 3.2 Hybrid Matrix Factorization

This paper introduces a new compression technique that uses dense matrix representation to ensure fast run-time properties and avoids making the strong assumptions made by LMF. This technique is based on three assumptions -

- **A1:** Rank of a matrix is important to create a high-task accuracy LSTM/GRU network (Yang et al., 2018)




Figure 1: Representation of a matrix using hybrid factorization

- **A2:** LMF makes a strong assumption that all the elements of the output feature vector of a LSTM/GRU Cell can be expressed from a low-dimensional embedding vector (Equation 2). HLF is based on the assumption that a more relaxed constraint of having a hybrid output feature vector, where some elements are calculated from a lower dimensional embedding space and other's from a higher dimensional embedding space can lead to better accuracy.
- **A3:** Most LSTM/GRU networks are followed by a fully-connected softmax layer or another LSTM/GRU layer. Even if the order of the elements in the output of a particular RNN layer changes, the weights in the subsequent fully connected or LSTM/GRU layers can adjust to accommodate that. Thus, the order of the elements of the output vector of LSTM/GRU layer is not strictly important.

These three intuitions of a LSTM/GRU layer can be used to create a more hardware-friendly compression scheme. This paper introduces one such scheme – Hybrid Matrix Factorization.

Hybrid Matrix Factorization (HLF) splits the input and recurrent matrices in an LSTM/GRU layer into two parts – a fully parameterized upper part and low-rank lower part.

Figure 1 shows the strategy we use to decompose the matrix - an unconstrained upper half  $A'$  and a lower half that is composed of  $k$  rank-1 blocks. If we decompose the weight matrix using this tech-

**Algorithm 1** Matrix vector product when a matrix uses the HLF technique

**Input 1:** Matrices  $A'$  of dimension  $j \times n$ ,  $B$  of dimension  $(m - j) \times k$ ,  $C$  of dimension  $k \times n$ )

**Input 2:** Vector  $I$  of dimension  $n \times 1$

**Output:** Matrix  $O$  of dimension  $m \times 1$

- 1:  $O_{1:j} \leftarrow A' \times I$
- 2:  $Temp1 \leftarrow C \times I$
- 3:  $O_{j+1:m} \leftarrow B \times Temp1$
- 4:  $O = concatenate\{O_{1:j}, O_{j+1:m}\}$

Matrix of Size (256,256)		
Compression Factor	LMF	HLF
1.25	102	103 - 204
1.67	76	78 - 153
2.50	51	52 - 101
5.00	25	26 - 50

Table 1: The maximum possible rank of a  $256 \times 256$  sized matrix after it is compressed by 4 different factors using 3 different compression techniques. To compress a matrix by a given compression factor, HLF has 2 different parameters  $j$  and  $k$  to regulate the rank of the matrix. Hence, we see a range of rank values. Maximum rank is achieved when  $k=1$ . The value for  $j$  when  $k=1$  can be calculated for different compression factors using equation 4.

nique, the parameter reduction is given by:

$$\frac{m \times n}{(j \times n) + k \times (m - j + n)} \quad (4)$$

Thus, the maximum rank of the matrix becomes  $j + k$ . Different values of  $j$  and  $k$  can be used to control the amount of compression and the rank of the matrix.

Structuring a matrix as shown in Figure 1 can lead to significant increase in maximum rank of the compressed matrix. Table-1 shows the maximum possible rank of a  $256 \times 256$  matrix compressed to the same number of parameters using the two compression techniques - LMF and HLF. As shown, HLF can effectively double the rank of the matrix for the same number of parameters. To compress a matrix by a given compression factor, HLF has 2 different parameters,  $j$  and  $k$ , to regulate the rank of the matrix. Hence, we see a range of rank values. Maximum rank is achieved when  $k=1$ . The value for  $j$  when  $k=1$  can be calculated for different compression factors using equation 4.

Apart from the storage reduction, HLF also leads to a reduction in the number of computations. Assuming a batch size of 1 during inference, HLF leads to inference speed-up by using the associative property of matrix products to calculate the matrix-vector product - Algorithm 1 shows how to calculate the matrix vector product when the matrix is represented using HLF. This algorithm avoids expanding the matrix  $A'$ ,  $B$  and  $C$  into  $A$ .

Algorithm 1 uses the associative property of matrix products to gain the computation speedup. For a matrix vector product between a matrix of size  $m \times n$  and a vector of size  $n \times 1$ , the number of operations required to compute the product is  $m \times n$  (Trefethen and Bau, 1997). Referring to Algorithm 1, number of operations required to calculate  $O_{1:j}$  is  $j \times n$ . The  $Temp1$  variables need  $k * n$  operations and calculating  $O_{j+1:m}$  needs  $k*(m-j)$  operations. Thus, the reduction in number of operations when we use Algorithm 1 is:

$$\frac{m \times n}{j \times n + k \times n + k \times (m - j)} \quad (5)$$

### 3.2.1 Impact on output feature vector

Algorithm 1 shows that, HLF divides the output into two stacked sub-vectors. One is a result of a fully-parameterized multiplication,  $A' \times I$  (Line 1, Algorithm 1). The other is the result of the low rank multiplication :  $B \times C \times I$  (Line 2-3, Algorithm 1). Thus, the upper sub-vector has “richer” features created from a higher dimensional embedding, while the lower sub vector has “constrained” features created from a lower dimensional embedding. By incorporating the HLF structure during training, we force an RNN to learn “richer” features in the upper sub-vector and the “constrained” features in the lower sub-vector. Because a RNN is followed by another RNN or a softmax layer, this restructuring should not impact the subsequent layers. Thus, HLF structure combines the assumptions A2 and A3 that were discussed previously.

### 3.3 Why HLF leads to larger rank than LMF for same number of parameters?

HLF is an extension of LMF. To understand this, let us revisit Figure 1, where the matrix

$$A = [A' ; BC]$$

where  $A \in R^{m \times n}$ ,  $A' \in R^{j \times n}$ ,  $B \in R^{(m-j) \times k}$  and  $C \in R^{k \times n}$ . Then we can rewrite the matrix



as,

$$A = [I, \mathbf{0}_1; \mathbf{0}_2, B][A'; C]$$

where  $I \in R^{j \times j}$ ,  $\mathbf{0}_1 \in R^{j \times k}$  and  $\mathbf{0}_2 \in R^{(m-j) \times j}$ . The above equation could be re-written as -

$$A = U'V'$$

where  $U' \in R^{m \times (j+k)}$  and  $V' \in R^{(j+k) \times n}$ . Both  $U'$  and  $V'$  can have a maximum rank of  $j+k$ . The maximum value of this rank is achieved when  $k=1$  and  $j$  is calculated as discussed in Table 1. Let this value be  $d$ . A standard LMF decomposition of  $A$  will also lead to a representation of the form  $UV$ , but this representation will have same parameters as HLF only if the rank of both  $U$  and  $V$  is at most  $(d+1)/2$ . Thus, HLF can be regarded as a  $(d+1)$  ranked LMF of  $A$ , with a sparsity forcing mask that reduces the number of parameters to express the  $(d+1)$  ranked matrix significantly. **This is why HLF can double the rank of the matrix when compared to an iso-parameter LMF matrix.**

Neural networks seldom learn structured sparsity unless they are forced to (Narang et al., 2017), thus, an RNN trained with the LMF structure will rarely end up learning the same structure as HLF. **The pre-determined HLF structure effectively creates a sparsity forcing mask.** Such a sparsity forcing mask also leads to creation of the decoupled output feature vectors as described in section 3.2.1.

## 4 Results

We compare HLF with LMF and 3 other compression techniques – model pruning, small baseline and a structured matrix based technique called block circular decomposition. These techniques and why they need to be considered are discussed below:

- Pruning: Model pruning (Zhu and Gupta, 2017) induces sparsity in the matrices of a neural network, thereby reducing the number of non-zero valued parameters that need to be stored. Pruning creates sparse matrices which are stored in a specialized sparse data structure such as CSR. The overhead of traversing these data structures while performing the matrix-vector multiplication can lead to poorer inference run-time than when executing the baseline, non-sparse network. *Thus, while pruning*


*is an effective compression technique, its runtime performance on CPUs can make it a less appealing choice for compression.* We use the magnitude pruning framework provided by (Zhu and Gupta, 2017). While there are other possible ways to prune, recent work (Gale et al., 2019) has suggested that magnitude pruning provides state-of-the-art or comparable performance when compared to other pruning techniques (Neklyudov et al., 2017; Louizos et al., 2018).

- Small Baseline: Additionally, we train a smaller baseline with the number of parameters equal to that of the compressed baseline. This serves as a useful point of comparison because of two reasons.
  - First, to check if compression of a larger network leads to better accuracy than compressing a network by reducing its dimensions (size of hidden layer or number of layers). **This can help us verify if the network was originally over-parameterized.**
  - Second, **to establish the hypothesis whether HLF’s creation of a stacked output feature vector as described in section 3.2.1 adds any useful information in the network.** Smaller baseline creates output feature vector that is created from a high-dimensional embedding only. HLF, additionally concatenates the output features created from lower dimensional embedding. Thus, comparing the accuracy of HLF with Smaller baseline helps evaluate the usefulness of the output features created using lower dimensional embedding.


**Given the significant slow-down of inference of BCD compressed networks, we do not discuss the results sing BCD compression in the rest of the paper.**

### 4.1 Experiment Setup


**Measuring inference run-time:** In order to compare the inference run-time of RNN cells compressed using pruning, LMF and HLF, we implemented these cells in C++ using the Eigen library. This paper focuses on inference on an edge device. As a result, we make the assumption that the batch size of the application will be 1 while



(a) For both figures, HLF exists in the top-right corner, providing the best trade-off in terms of accuracy, compression and speed-up. For all 3 compression points, HLF provides better accuracy than LMF and SB and better run-time than pruning. *(Left) Language Translation: HLF improves the BLEU score achieved by LMF by 2.3% to 4.5% and by Small Baseline by 2.8% to 4.1%. At the same time, HLF improves the inference run-time over pruning by 1.5 – 1.74 $\times$ .* *(Right) Text Classification: HLF can improve the accuracy achieved by LMF by up-to 1.2% and SB by up-to 1.3% and improve the run-time achieved by pruning by up-to 1.2 $\times$*



(b) For both figures, HLF exists in the top-right corner, providing the best trade-off in terms of accuracy, compression and speed-up. *(Left) Slot Filling: For all 3 compression points, HLF provides better accuracy than LMF and SB and better run-time than pruning. HLF can improve the accuracy achieved by LMF and SB by up-to 1.2% and improve the run-time achieved by pruning by up-to 1.26 $\times$ .* *(Right) Intent Detection: For all 3 compression points, HLF provides better accuracy than LMF and SB and better run-time than pruning. HLF can improve the accuracy achieved by LMF and SB by up-to 1% and improve the run-time achieved by pruning by up-to 1.26 $\times$*



(c) Language Modeling: For LM, lower values of perplexity are better. HLF provides a viable alternative to both pruning and LMF at compression factors of 2.5x, 3.0x and 5x. **HLF can achieve 17% better perplexity than LMF, 9% better perplexity than small baseline and 2.32 $\times$  better inference run-time than pruning.**

Figure 2: Speed-up over baseline vs Accuracy comparing the baseline with a smaller baseline and the baseline compressed using different compression schemes at varying compression factors. Speed-up values  $> 1$  indicate a decrease in inference run-time and values  $< 1$  indicate an increase in inference run-time. **For each compression factor, the compression scheme that is most to the top-right is the ideal choice.** *In case of perplexity, lower values are better.* Thus, the graphs are plotted in a slightly different way to still adhere to the fact that the most ideal choice of compression is in the top-right corner. P = Pruning, LMF = Low rank matrix factorization, HLF = Hybrid matrix decomposition, SB = Smaller baseline. **The best way to view this figure is to either focus on a compression point and see how the Pareto curve of speed-up vs accuracy changes as we add HLF or focus on an accuracy region and see what compression schemes provides the best run-time at highest compression factor.**

measuring the run-time of an application. However, the observations regarding run-time should remain consistent for larger batch sizes as well. We ran our experiments on a single cortex-A73 core of the Hikey 960 board. The size of L3 cache is 2MB.

**Training infrastructure:** We use Tensorflow 1.14 to train our networks on a cluster of 2 RTX 2080 Ti NVidia GPUs with 11 GB Memory. The training settings for the benchmarks evaluated can be found in the reference paper for each benchmark.

**What do we compress?** We compress the LSTM/GRU layers in each application. We do not compress the embedding layers using HLF.

The amount of compression determines the rank of the compressed matrices when we use LMF to compress an NLP application and the sparsity of the pruned matrix when we use pruning as our choice of compression technique. Similar to LMF, the amount of compression determines the rank of the compressed matrices when we use HLF as our choice of compression. However, for HLF two parameters  $j$  and  $k$  control the rank of the matrix. We use a sweep starting with  $k=1$  to determine the exact values of  $j$  and  $k$  that help us achieve a good accuracy.

**What do we compare?:** We compare the accuracy and inference run-time of all compression technologies at iso-compression factors for various compression techniques.

## 4.2 Comparison of compression techniques across different ML tasks

The impact of compression on accuracy is compared for 5 benchmarks – machine translation, natural language understanding (intent detection and slot filling), text classification and Language Modeling. These tasks are some of the most important NLP applications that run on edge and embedded devices like smart phones, smart watches and smart homes.

### 4.2.1 Language Translation

We use the English to Vietnamese translation model in (Luong et al., 2017). The model uses 2-layer LSTMs of size 512 units with bidirectional encoder (i.e., 1 bidirectional layers for the encoder), embedding of dim 512 and an attention layer. We used the hyper-parameters in (Luong et al., 2017) to train the network while modifying the learning rate values used. We sweep the learning rates values from

$\times 0.1$  to  $3.0$  in multiples of  $3$ . For HLF, we used the value of  $k=2$ .

Figure 2a shows the results of compressing the LSTM layers in the NMT\_VIEN baseline by  $2.5\times$ ,  $3.33\times$  and  $5\times$ . **HLF improves the BLEU score achieved by LMF by 2.3% to 4.5% and by Small Baseline by 2.8% to 4.1%. At the same time, HLF improves the inference run-time over pruning by  $1.5\times - 1.74\times$ .**

### 4.2.2 Language Modeling

We use the medium LM model from (Zaremba et al., 2014) as our baseline. The PTB (Medium) baseline has 2 LSTM layers each with a hidden vector of size 650 with a vocabulary size of 10,000 words from the English vocabulary. We used the hyper-parameters in (Zaremba et al., 2014) and train the compressed networks for 50 more epochs than in baseline. The baseline network is trained for 39 epochs. For the first 6 epochs the learning rate used is of value 1, and after that we decrease it by a factor of 1.2 after each epoch. We clip the norm of the gradients at 5 and use dropout of value 0.35. For HLF, we used the value of  $k=4$ .

Figure 2c shows the results of compressing the LSTM layers in the PTB (Medium) baseline by  $2.5\times$ ,  $3.33\times$  and  $5\times$ . Lower the perplexity, better the model. Pruning achieves the same (sometimes better) perplexity than baseline and other compression techniques. LMF leads to significant loss in perplexity for all compression factors while HLF achieves better perplexity than LMF and faster inference run-time than baseline and pruned networks for all compression factors. **In fact, HLF can achieve 17% better perplexity than LMF, 9% better perplexity than small baseline and  $2.32\times$  better inference run-time than pruning.** The preferred choice of compression scheme for different compression factors will depend on whether slight loss in perplexity can be accommodated for faster inference run-time or not. However, HLF still manages to serve as a more viable alternative to pruning than LMF for inference on edge CPUs.

### 4.2.3 Text Classification

We use the text classification network in (Zhou et al., 2016) evaluated on the SemEval-2010 dataset. The baseline network has 1 bidirectional LSTM layers with hidden vector of size 256. We used the hyper-parameters in (Zhou et al., 2016) to train the baseline and as the initial hyperparameters ex-

plored for the compressed networks. We trained the compressed networks for additional 20 epochs while exploring learning rates of  $10\times$  and  $1/10$  than the baseline learning rate. The baseline model was trained using AdaDelta with a learning rate of 1.0. The model parameters were regularized with L2 regularization strength of  $10^{-5}$ .

Figure 2b shows the results for compressing the text classification network by  $2.5\times-5\times$ . **HLF can improve the accuracy achieved by LMF by up-to 1.2%, by SB by up-to 1.3% and improve the run-time achieved by pruning by almost  $1.20\times$ .**

#### 4.2.4 Intent Detection and Slot Filling

We used the benchmark published in (Liu and Lane, 2016). This benchmark is trained on the ATIS dataset and jointly trains for intent detection and slot filling. The benchmark uses 1 LSTM layer of size 128 along with attention layers. For HLF, we used the value of  $k=1$ .

Figure 2b shows the results for slot filling task. **HLF can improve the F1-accuracy achieved by LMF and SB by up-to 1.2% and improve the run-time achieved by pruning by up-to  $1.26\times$ .** Figure 2b shows the result for the intent classification task. Due to joint training, the network used for slot filling and intent classification is the same. As a result, the runtime improvement of HLF over pruning is exactly the same as for the slot filling task. **Additionally, HLF improves the intent classification accuracy by up to 1% over LMF and small baseline.**

### 4.3 Ablation Studies

#### 4.3.1 Compressing Word Embedding layers using HLF

We compressed the input word-embedding layers in the PTB-LM model discussed in section 4.2.2, without compressing other layers in the network. However, even  $3\times$  compression using HLF led to 8% loss in perplexity score.

#### 4.3.2 Orthogonality of word embedding compression methods and HLF

We ran experiments where we prune the word embedding layers in the PTB-LM model in section 4.2.2 by  $2\times$  while keeping the LSTM layers uncompressed, leading to 83.1 perplexity score. We were able to further compress this network by  $2\times$  using HLF with only 1 point loss in perplexity score, indicating that HLF is compatible with techniques used for compressing word-embedding layers.

## 5 Discussion

Effectively, HLF acts as an alternative to LMF whenever compression using pruning does not lead to the required run-time benefit and LMF leads to loss in accuracy. HLF has a better accuracy than LMF for most evaluation points, validating the assumption in the paper that rank of a matrix in a RNN is important for better task accuracy in NLP applications. Additionally, HLF has a better accuracy than smaller baseline. This validates the assumption of the importance of constrained features in addition to the richer features in a Small Baseline network.

## 6 Limitations

While HLF provides significant benefits over LMF, there are two limitations associated with the technique:

- The unique nature of RNNs (Assumption A1-A3) makes HLF a natural fit for LSTM/GRU layers. However, these assumptions are not valid for the final classification layer. In classification layer, HLF will lead to more expressive output for certain classes (the top part of HLF matrix) in the dataset and less expressive output for the rest of the classes (bottom part of HLF matrix).
- HLF is a training aware compression and cannot be applied to a pre-trained network.

## 7 Conclusion

Choosing the right compression technique requires looking at three criteria – compression factor, accuracy, and run-time. Pruning is an effective compression technique, but can sacrifice speedup over baseline for certain compression factors. LMF achieves better speedup than baseline for all compression factors, but leads to accuracy degradation. This paper introduces a new compression scheme called HLF, which preserves the dense structures of LMF while effectively doubling the rank of the matrix using an intelligent structure by design. This leads to  $2\times$  faster inference run-time than pruning and up-to 16% better accuracy than LMF.

## References

Anish Acharya, Rahul Goel, Angeliki Metallinou, and Inderjit Dhillon. 2019. Online embedding compression for text classification using low rank matrix fac-



- torization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6196–6203.
- James Bradbury, Stephen Merity, Caiming Xiong, and Richard Socher. 2016. [Quasi-recurrent neural networks](#). *CoRR*, abs/1611.01576.
- Víctor Campos, Brendan Jou, Xavier Giró-i Nieto, Jordi Torres, and Shih-Fu Chang. 2018. Skip rnn: Learning to skip state updates in recurrent neural networks. In *International Conference on Learning Representations*.
- Ting Chen, Ji Lin, Tian Lin, Song Han, Chong Wang, and Denny Zhou. 2018. [Adaptive mixture of low-rank factorizations for compact neural modeling](#). *Advances in neural information processing systems (CDNNRIA workshop)*.
- Y. Cheng, F. X. Yu, R. S. Feris, S. Kumar, A. Choudhary, and S. Chang. 2015. [An exploration of parameter redundancy in deep networks with circulant projections](#). In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 2857–2865.
- Caiwen Ding, Siyu Liao, Yanzhi Wang, Zhe Li, Ning Liu, Youwei Zhuo, Chao Wang, Xuehai Qian, Yu Bai, Geng Yuan, Xiaolong Ma, Yipeng Zhang, Jian Tang, Qinru Qiu, Xue Lin, and Bo Yuan. 2017. [Circnn: Accelerating and compressing deep neural networks using block-circulant weight matrices](#). In *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO-50 '17*, pages 395–408, New York, NY, USA. ACM.
- Caiwen Ding, Ao Ren, Geng Yuan, Xiaolong Ma, Jiayu Li, Ning Liu, Bo Yuan, and Yanzhi Wang. 2018. [Structured weight matrices-based hardware accelerators in deep neural networks: Fpgas and asics](#). In *Proceedings of the 2018 on Great Lakes Symposium on VLSI, GLSVLSI '18*, pages 353–358, New York, NY, USA. ACM.
- Trevor Gale, Erich Elsen, and Sara Hooker. 2019. [The state of sparsity in deep neural networks](#). *CoRR*, abs/1902.09574.
- Dibakar Gope, Jesse Beu, Urmish Thakker, and Matthew Mattina. 2020a. Ternary mobilenets via per-layer hybrid filter banks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Dibakar Gope, Jesse G. Beu, Urmish Thakker, and Matthew Mattina. 2020b. [Aggressive compression of mobilenets using hybrid ternary layers](#). *tinyML Summit*.
- Dibakar Gope, Ganesh Dasika, and Matthew Mattina. 2019. Ternary hybrid neural-tree networks for highly constrained iot applications. In *Proceedings of Machine Learning and Systems 2019*, pages 190–200.
- Artem M. Grachev, Dmitry I. Ignatov, and Andrey V. Savchenko. 2017. Neural networks compression for language modeling. In *Pattern Recognition and Machine Intelligence*, pages 351–357, Cham. Springer International Publishing.
- Song Han, Huizi Mao, and William J Dally. 2016. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *International Conference on Learning Representations (ICLR)*.
- Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2016. Binarized neural networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, page 4114–4122, Red Hook, NY, USA. Curran Associates Inc.
- Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2017. Quantized neural networks: Training neural networks with low precision weights and activations. *J. Mach. Learn. Res.*, 18(1):6869–6898.
- Oleksii Kuchaiev and Boris Ginsburg. 2017. [Factorization tricks for LSTM networks](#). *CoRR*, abs/1703.10722.
- Tao Lei, Yu Zhang, Sida I. Wang, Hui Dai, and Yoav Artzi. 2018. [Simple recurrent units for highly parallelizable recurrence](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4470–4481, Brussels, Belgium. Association for Computational Linguistics.
- Zhe Li, Shuo Wang, Caiwen Ding, Qinru Qiu, Yanzhi Wang, and Yun Liang. 2018. Efficient recurrent neural networks using structured matrices in fpgas. 6th International Conference on Learning Representations, ICLR 2018 ; Conference date: 30-04-2018 Through 03-05-2018.
- Bing Liu and Ian Lane. 2016. [Attention-based recurrent neural network models for joint intent detection and slot filling](#). In *Interspeech 2016*, pages 685–689.
- Xuan Liu, Di Cao, and Kai Yu. 2018. [Binarized LSTM language model](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2113–2121, New Orleans, Louisiana. Association for Computational Linguistics.
- Christos Louizos, Max Welling, and Diederik P. Kingma. 2018. [Learning sparse neural networks through l<sub>0</sub> regularization](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Zhiyun Lu, Vikas Sindhwani, and Tara N. Sainath. 2016. [Learning compact recurrent neural networks](#). *CoRR*, abs/1604.02594.

- Minh-Thang Luong, Eugene Brevdo, and Rui Zhao. 2017. Neural machine translation (seq2seq) tutorial. <https://github.com/tensorflow/nmt>.
- Sachin Mehta, Rik Koncel-Kedziorski, Mohammad Rastegari, and Hannaneh Hajishirzi. 2018. Pyramidal recurrent unit for language modeling. *CoRR*, abs/1808.09029.
- Sachin Mehta, Rik Koncel-Kedziorski, Mohammad Rastegari, and Hannaneh Hajishirzi. 2019. Define: Deep factorized input token embeddings for neural sequence modeling.
- Sharan Narang, Eric Undersander, and Gregory F. Damos. 2017. Block-sparse recurrent neural networks. *CoRR*, abs/1711.02782.
- Kirill Neklyudov, Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. 2017. Structured bayesian pruning via log-normal multiplicative noise. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 6778–6787, Red Hook, NY, USA. Curran Associates Inc.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter.
- Victor Sanh, Thomas Wolf, and Alexander M. Rush. 2020. Movement pruning: Adaptive sparsity by fine-tuning.
- Min Joon Seo, Sewon Min, Ali Farhadi, and Hannaneh Hajishirzi. 2018. Neural speed reading via skim-rnn. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Vikas Sindhvani, Tara Sainath, and Sanjiv Kumar. 2015. Structured transforms for small-footprint deep learning. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 3088–3096. Curran Associates, Inc.
- Jin Tao, Urmish Thakker, Ganesh Dasika, and Jesse Beu. 2019. Skipping rnn state updates without re-training the original model. In *Proceedings of the 1st Workshop on Machine Learning on Edge in Sensor Systems, SenSys-ML 2019*, page 31–36, New York, NY, USA. Association for Computing Machinery.
- U. Thakker, J. Beu, D. Gope, G. Dasika, and M. Mattina. 2019. Run-time efficient rnn compression for inference on edge devices. In *2019 2nd Workshop on Energy Efficient Machine Learning and Cognitive Computing for Embedded Applications (EMC2)*, pages 26–30.
- Urmish Thakker, Jesse G. Beu, Dibakar Gope, Chu Zhou, Igor Fedorov, Ganesh Dasika, and Matthew Mattina. 2019a. Compressing rnns for iot devices by 15-38x using kronecker products. *CoRR*, abs/1906.02876.
- Urmish Thakker, Ganesh Dasika, Jesse G. Beu, and Matthew Mattina. 2019b. Measuring scheduling efficiency of rnns for NLP applications. *CoRR*, abs/1904.03302.
- Urmish Thakker, Igor Fedorov, Jesse G. Beu, Dibakar Gope, Chu Zhou, Ganesh Dasika, and Matthew Mattina. 2019c. Pushing the limits of RNN compression. *CoRR*, abs/1910.02558.
- Urmish Thakker, Paul Whatamough, Matthew Mattina, and Jesse G. Beu. 2020. Compressing language models using doped kronecker products. *CoRR*, abs/2001.08896.
- Anna Thomas, Albert Gu, Tri Dao, Atri Rudra, and Christopher Ré. 2018. Learning compressed transforms with low displacement rank. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9066–9078. Curran Associates, Inc.
- Andros Tjandra, Sakriani Sakti, and Satoshi Nakamura. 2017. Compressing recurrent neural network with tensor train. In *Neural Networks (IJCNN), 2017 International Joint Conference on*, pages 4451–4458. IEEE.
- Lloyd Trefethen and David Bau. 1997. *Numerical Linear Algebra*. SIAM: Society for Industrial and Applied Mathematics.
- Shuo Wang, Zhe Li, Caiwen Ding, Bo Yuan, Qinru Qiu, Yanzhi Wang, and Yun Liang. 2018. C-lstm: Enabling efficient lstm using structured compression techniques on fpgas. In *Proceedings of the 2018 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays, FPGA '18*, pages 11–20, New York, NY, USA. ACM.
- Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, and William W. Cohen. 2018. Breaking the softmax bottleneck: A high-rank RNN language model. In *International Conference on Learning Representations*.
- Adams Wei Yu, Hongrae Lee, and Quoc Le. 2017. Learning to skim text. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1880–1890, Vancouver, Canada. Association for Computational Linguistics.
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. *CoRR*, abs/1409.2329.
- Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. Attention-based bidirectional long short-term memory networks for relation classification. In *Proceedings of the 54th*

*Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 207–212, Berlin, Germany. Association for Computational Linguistics.

Michael Zhu and Suyog Gupta. 2017. [To prune, or not to prune: exploring the efficacy of pruning for model compression.](#) *arXiv e-prints*, page arXiv:1710.01878.