Hypoformer: Hybrid Decomposition Transformer for Edge-friendly Neural Machine Translation

Sunzhu Li[1], Peng Zhang[1*], Guobing Gan[1], Xiuqing Lu[1],
Benyou Wang[2], Junqiu Wei[3], Xin Jiang[4]

[1]College of Intelligence and Computing, Tianjin University, Tianjin, China
[2]School of Data Science, The Chinese University of Hong Kong, Shenzhen, China
[3]The Hong Kong Polytechnic University, China   [4]Huawei Noah's Ark Lab, China
{lisunzhu,pzhang,ganguobing,lvxiuqing}@tju.edu.cn
junwei@polyu.edu.hk, wangbenyou@cuhk.edu.cn, Jiang.Xin@huawei.com

Abstract

Transformer has been demonstrated effective in Neural Machine Translation (NMT). However, it is memory-consuming and time-consuming in edge devices, resulting in some difficulties for real-time feedback. To compress and accelerate Transformer, we propose a *Hybrid Tensor-Train (HTT)* decomposition, which retains full rank and meanwhile reduces operations and parameters. A Transformer using HTT, named *Hypoformer*, consistently and notably outperforms the recent lightweight SOTA methods on three standard translation tasks under different parameter and speed scales. In extreme low resource scenarios, Hypoformer has a **7.1** point absolute improvement in BLEU and $1.27\times$ speedup than the vanilla Transformer on the IWSLT'14 De-En task.

1 Introduction

Transformer (Vaswani et al., 2017) is one of the most effective models in NMT. Transformer could achieve better performance by stacking more layers and increasing the embedding dimension (Liu et al., 2020; Li et al., 2021a). However, this leads to a huge scale of parameters and a slow inference speed. Various works have been recently proposed to compress Transformer, such as light-weight variants (Wu et al., 2020; Mehta et al., 2021), neural architecture search (Wang et al., 2020), parameter sharing (Reid et al., 2021), low-rank approximation (Ma et al., 2019; Liu et al., 2021; Hrinchuk et al., 2020), knowledge distillation (Kasai et al., 2020), layer reassignment (Hsu et al., 2020; Bérard et al., 2021; Li et al., 2021c).

Following the line of low-rank approximation, existing works to compress NMT models (or related deep Transformer models like BERT) 1) are specific to a part of sub-components, *e.g., self-attention networks* (Ma et al., 2019) or *embedding layer* (Hrinchuk et al., 2020); and 2) show limited

---
[*]Corresponding Author

or no potential for speedup during inference (Ma et al., 2019; Hrinchuk et al., 2020; Noach and Goldberg, 2020; Liu et al., 2021). However, unlike the cloud-based system, it needs real-time feedback and faster inference in edge devices (*e.g.,* smartphones and IoTs) even with constrained resources. Although it is challenging, we believe that low-rank approximation has the potential for compressing and accelerating NMT models in edge devices.

From a general point of view, this paper rethinks low-rank approximation methods, especially two representative methods namely Matrix Factorization (MF) and Tensor-Train decomposition (TT). The former is arguably time-efficient while the latter is parameter-efficient. However, **(1)** MF will encounter a **low-rank bottleneck** with a high compression ratio, potentially leading to a moderate performance drop (Thakker et al., 2020). **(2)** Transformer compression based on *high TT-rank* TT decelerates inference speed due to its **quadratic computational complexity**. Interestingly, TT was claimed to retain a full matrix rank by Hrinchuk et al. (2020) while the expressive power of MF is bounded to rank R (see Table 1). Although *low-rank* TT can significantly reduce parameters and operations, preliminary experimental results show that leveraging *low-rank* TT alone results in significant performance degradation (see Table 2).

To avoid performance degradation while preserving its parameter efficiency of low-rank TT, we propose to compensate *low-rank* TT with an auxiliary dense projection matrix, resulting in a hybrid decomposition (Figure 2E), called *Hybrid Tensor-Train (HTT) Decomposition*, which is inspired by the hybrid form of MF (Thakker et al., 2020). HTT can retain the full matrix rank, since its dense part is full matrix rank in definition, and its TT part is theoretically full matrix rank referring to TT embedding (Hrinchuk et al., 2020). The full matrix rank indicates that the expressivity of the original matrix can be retained in its decomposed

form. Thus, it is expected that HTT decomposition can maintain the performance of original matrix, meanwhile with fewer operations and parameters.

Based on HTT decomposition, we present a Transformer-based NMT models, called *Hypoformer* in a knowledge distillation fashion. Specifically, we replace the dense layers of sub-components (like embedding and self-attention) with HTT, which can compress and accelerate the vanilla Transformer.

We evaluate Hypoformer with three standard machine translation tasks (IWSLT'14 De-En, WMT'16 En-Ro, and WMT'14 En-De) on Raspberry Pi ARM CPU and Intel CPU. Hypoformer consistently outperforms recent lightweight state-of-the-art methods (Wu et al., 2020; Mehta et al., 2021; Wang et al., 2020) under various parameter scales and speed scales. Achieving similar performance on three tasks, Hypoformer can compress Transformer base model (Vaswani et al., 2017) $2.9 \sim 4.5\times$, and speedup it $1.9 \sim 3.3\times$ on Intel CPU and $1.8 \sim 3.9\times$ on Raspberry Pi ARM CPU. Notably, when the parameter scale is extremely low, Hypoformer achieves a **7.1** higher BLEU score and $1.27\times$ speedup than Transformer on IWSLT'14 De-En task.

The contributions of this work are as follows: **(1)** Based on the analysis of previous low-rank approximation methods (Oseledets, 2011; Thakker et al., 2020), we propose a novel *Hybrid Tensor-Train (HTT) Decomposition*, which can retain full matrix rank while with fewer operations and parameters. **(2)** Applying HTT decomposition to compress and accelerate Transformer, *Hypoformer* consistently outperforms the recent light-weight SOTA methods (Wu et al., 2020; Mehta et al., 2021) with a higher compression ratio, speedup, and better BLEU score. **(3)** In edge devices with extremely limited resources, Hypoformer can retain **95%** performance with **1/12** parameters, setting a new SOTA in extremely-small NMT models.

2 Preliminaries

Tensor Representation First, we introduce the Tensor diagrams (Penrose, 1971). An n-order tensor can be denoted as a node with n legs. Each leg is labeled by the length (an associated positive integer) of the corresponding dimension, and the orientation of the leg does not matter. As shown in Figure 1, a vector $A \in \mathbb{R}^I$, a matrix $B \in \mathbb{R}^{I \times J}$ and a 3-order tensor $C \in \mathbb{R}^{J \times K \times I}$, shown in Figure

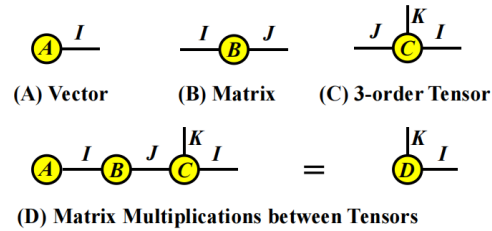1(A)(B)(C), have one, two, three legs, respectively. The tensor operation is shown in Appendix A.1.

(D) Matrix Multiplications between Tensors

Figure 1: Tensor diagrams (Penrose, 1971) for a vector (A), a matrix (B), a 3-order tensor (C), and the tensor contraction among them (D).

Matrix Factorization is a common and effective technique to compress deep neural networks. To compress the fully connected layer, the weight matrix $W \in \mathbb{R}^{I \times J}$ is decomposed into the product of two smaller matrices $U \in \mathbb{R}^{I \times R}$ and $V \in \mathbb{R}^{R \times J}$, as shown in Figure 2(B). The compression ratio of Matrix Factorization (MF) is related to the MF rank R; the larger R is, the smaller the compression ratio and the bigger the capacity.

Tensor-Train Decomposition Oseledets (2011) gives a method to represent the high-order tensor with a low-rank approximation. In Tensor-Train (TT) decomposition, each element in the original tensor can be represented as the product of two vectors and a series of matrices. For the weight matrix $W \in \mathbb{R}^{I \times J}$, where $I = \prod_{k=1}^{n} I_k$ and $J = \prod_{k=1}^{n} J_k$, it can be reshaped as a $2n$-order tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times \cdots \times I_n \times J_1 \times \cdots \times J_n}$. Then, the element in $\mathcal{T}$ is defined as follow:

$$
\begin{aligned}
&\mathcal{T}(i_1, ..., i_n, j_1, ..., j_n) \\
&= T_1(:i_1, j_1, :) T_2(:i_2, j_2, :) \cdots T_n(:, i_n, j_n, :)
\end{aligned}
\tag{1}
$$

where $T_k \in \mathbb{R}^{R_{k-1} \times I_k \times J_k \times R_k}$ are referred to as TT-cores, which are smaller tensors as shown in Figure 2(C). The set $\{R_k\}_{k=0}^{n}$ is called TT-ranks, and $R_0 = R_n = 1$. Usually, $R_1, \ldots, R_{n-1}$ are equal, denoted as R (Novikov et al., 2015). TT decomposition controls the compression ratio by adjusting TT-ranks.

3 Hybrid Tensor-Train Decomposition

In this section, we first discuss the limitation of MF and TT, as well as their expressivity comparison in the case with limited parameters. Then, we introduce *Hybrid Tensor-Train decomposition*, which can retain full matrix rank but with fewer operations and parameters than the original matrix.
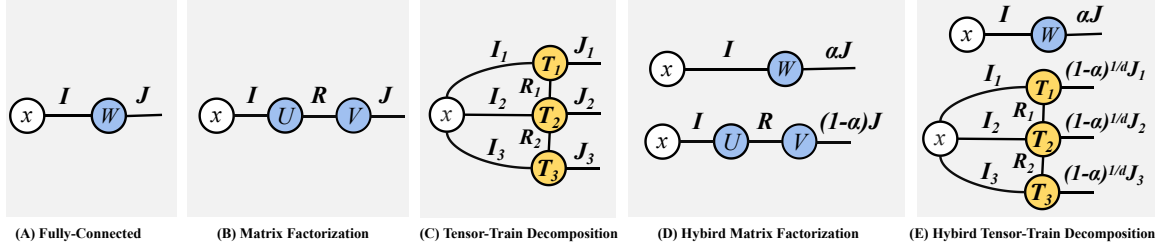
Figure 2: The operation of different decomposition using tensor graphical representations. For a given feature representation denoted as $x \in \mathbb{R}^I$, there show five ways to implement an $I \times J$ linear projection for the given x. The *hybrid* decomposition 1) conducts both a dense linear projection and MF/TT layer on an identical x; and 2) concatenates such output to obtain J-dimensional output. α denotes the ratio of dense layer, where $\alpha \in [0, 1]$. $I = I_1 \times I_2 \times I_3$. $J = J_1 \times J_2 \times J_3$. d is the number of TT-cores of Tensor-Train.

Method	Time Complexity	Space Complexity	Max Rank	Potential Risk
Fully-Connected	$\mathcal{O}(N^2)$	$\mathcal{O}(N^2)$	N	-
Matrix Factorization	$\mathcal{O}(NR)$	$\mathcal{O}(NR)$	R	low rank bottleneck
Tensor-Train	$\mathcal{O}(DN^{1+\frac{1}{D}}R^2)$	$\mathcal{O}(DN^{\frac{2}{D}}R^2)$	N	quadratic time complexity *w.r.t.* R
Hybrid Matrix Factorization	$\mathcal{O}(\alpha N^2 + 2NR - \alpha NR)$	$\mathcal{O}(\alpha N^2 + 2NR - \alpha NR)$	$\alpha N + R$	can not achieve full rank
Hybrid Tensor-Train	$\mathcal{O}(\alpha N^2 + D(max\{\alpha N, (1-\alpha)N\})^{1+\frac{1}{D}}R^2)$	$\mathcal{O}(\alpha N^2 + D(1-\alpha)^{\frac{1}{D}}N^{\frac{2}{D}}R^2)$	N	

Table 1: The comparison of Time and Space Complexity between Different decomposition methods for a $N \times N$ Fully Connected layer. **Max Rank** denotes that the approximated matrix with these decompositions can theoretically achieve the maximum rank. α indicates the ratio of dense part, where $\alpha \in [0, 1]$. R denotes the rank of MF and the TT-ranks of TT. D indicates the TT-cores number of TT.

3.1 Rethinking MF and TT

Low-rank bottleneck in MF The rank of MF is essential to its expressivity (Thakker et al., 2020). To obtain a high compression ratio, MF needs an extremely small rank, which will lead to a Low-rank bottleneck. Technically, for a matrix $W \in \mathbb{R}^{I \times J}$, the compression ratio is $\frac{IJ}{(I+J)R}$. Furthermore, a lower rank of MF has the potential to lead to a performance drop. For example, Novikov et al. (2015) found that compressing a dense layer with a rank 50 to 1 MF layer, shows a significant drop in performance. Noach and Goldberg (2020) also found a similar conclusion in compressing BERT (Devlin et al., 2019) with MF.

Quadratic computational complexity in TT Since the computational complexity of TT is $\mathcal{O}(DN^{1+\frac{1}{D}}R^2)$, TT decomposition has a *quadratic computational complexity* with respect to the TT-ranks R (Table 1). Fortunately, the quadratic computational complexity issue could be tolerated in the case TT has relatively small TT-ranks. In this paper, we avoid using TT with relatively high TT-ranks. For example, to compress a 512×512 dense layer, we used a TT-cores=3 and TT-ranks=2 TT layer, with only $1/512$ parameters and $1/8$ computational operations. As shown in Table 2, we do not directly apply original TT layers

to the Transformer in this work, since the low-rank TT Transformer suffers a performance degradation and the high-rank TT Transformer decelerates the inference speed due to quadratic computational complexity.

Expressivity between MF and TT In the case TT has relatively small TT-ranks, *TT is more expressive than MF with a comparable amount of parameters.* For example, Novikov et al. (2015) found that compressing a dense layer with a TT-ranks = 1 TT layer outperforms a rank = 50 LMF layer and achieves more significant compression. This means that the rank of LMF can not be set too small for preserving performance. Hrinchuk et al. (2020) found that TT has *full matrix rank* while compressing a matrix, and does not reduce the expressivity.

3.2 Hybrid Tensor-Train Decomposition

TT decomposition has potential since any TT-parameterized matrix has full matrix rank (Hrinchuk et al., 2020). Considering that the *high-rank* TT decomposition is time-consuming, we aim to explore *low-rank* TT in this paper. However, the preliminary experimental results in Table 2 show that purely making use of *low-rank* TT decomposition results in a large performance drop.

To compensate *low-rank* TT in terms of expressive power, we propose Hybrid Tensor Train decomposition (HTT) to mix TT decomposition with an auxiliary dense projection matrix [1]. The reasons are twofold. First, HTT can relieve the quadratic computational issue since TT-rank is small. Second, HTT is expressive thanks to the compensation from a dense matrix. HTT consists of two parts: a dense part and a Tensor-Train part (Figure 2(E)). The formula is as follows:

$$W = [W_{dense}, W_{tt}]$$
$$\mathcal{W}_{tt}(i_1, ..., i_n, j_1, ..., j_n) \qquad (2)$$
$$= \mathcal{G}_1(1, i_1, j_1, :) \cdots \mathcal{G}_n(:, i_n, j_n, 1)$$

where $W \in \mathbb{R}^{I \times J}$ can be regarded as the embedding layer or the weight matrix of Fully Connected layer. $W_{dense} \in \mathbb{R}^{I \times \alpha J}$, $W_{tt} \in \mathbb{R}^{I \times (1-\alpha)J}$, and $\alpha \in [0, 1]$ is used to control the proportion of the dense and TT part. W_{tt} can be derived from tensor $\mathcal{W}_{tt}$ by reshaping. TT-cores $\mathcal{G}_k \in \mathbb{R}^{R_{k-1} \times I_k \times J_k \times R_k}$, and n is the numbers of TT-cores. We use 2 or 3 TT-cores in this paper. $\prod_{k=1}^{n} I_k = I$, $\prod_{k=1}^{n} J_k = (1-\alpha)J$, and $k = 1, ..., n$. The space/time complexity and expressive power (in terms of rank) of the above decompositions are shown in Table 1. Note that *HTT could conduct flexible decomposition*: There exists a hyperparameter α in HTT to adjust the compression ratio of original matrices.

Theorem 3.1. *Suppose there is an original matrix $M \in \mathbb{R}^{N \times N}$, we use an HTT to approximate the original matrix. Theoretically, the max matrix rank of the original matrix and HTT are both N. Moreover, The max matrix rank of HTT is always full, which is not restricted to α and TT-ranks R.*

$$\text{MaxRank}([\mathbf{M}_{tt}; \mathbf{M}_{dense}]) = \text{MaxRank}(M) = N \qquad (3)$$

where HTT includes two parts: the dense part ($\mathbf{M}_{dense}$) is a $N \times \alpha N$ matrix and TT part ($\mathbf{M}_{tt}$) is to approximate a $N \times (1-\alpha)N$ matrix. The proof can be found in Appendix A.2.

Expressivity between HTT and HMF Theoretically, *HTT is superior to HMF in the theoretical maximum rank ($N \geq \alpha N + R$)* (Table 1). To achieve a high compression ratio, the α and R of HMF are relatively small, which may still lead to a low-rank bottleneck. The empirical results (Table 2) show that HMF leads to a performance drop at high compression.
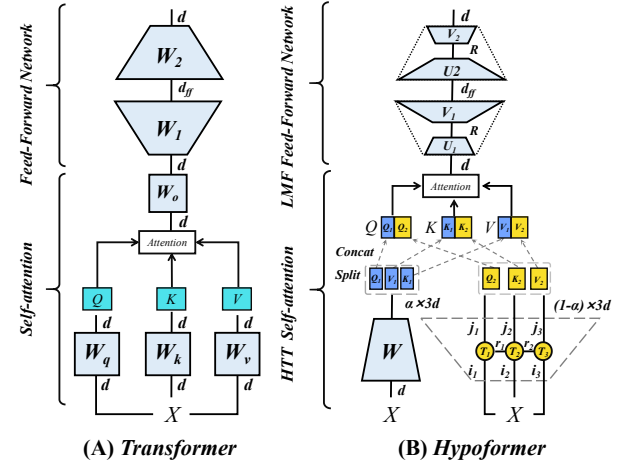

Figure 3: The comparison between vanilla Transformer and Hypoformer. The difference is that Hypoformer replaces dense layers with compressed ones.

4 Hypoformer

In this section, we first show the overall structure between the Transformer and our Hypoformer in Figure 3. Then, we introduce the three *Hypoformer* sub-layers including HTT Embedding, HTT Self-Attention, and LMF Feed-Forward Network.

4.1 HTT Embedding

The embedding layer is a very large matrix, which can not be overlooked. For example, it accounts for about $1/4$ of the model parameters in WMT'14 En-De. Previous works (Wu et al., 2019) used a joint source-target vocabulary with a large embedding $W \in \mathbb{R}^{v \times d}$, where v is the size of vocabulary and d is the embedding dimension.

Hrinchuk et al. (2020) showed that directly using Tensor-Train embedding to replace the original embedding layer would damage the performance of the machine translation model. To address this problem, we propose a new word embedding method named Hybrid Tensor-Train Embedding (HTT EMB) that mixes dense embedding and sparse embedding. Suppose the number of words in the vocabulary is v, and the dimension of word embedding is d. HTT EMB is obtained by concatenating a low-dimensional dense embedding matrix

---

[1] A related work is Thakker et al. (2020) which proposes a Hybrid Matrix Factorization (HMF). HMF decomposes a matrix into two parts — a 'narrow' matrix part and a MF part (Figure 2(D)). It is claimed that HMF can *obtain a higher rank than MF with a comparable amount of parameters.*

and a TT-matrix with 3 TT-cores, which can be formulated as follows:

$$W_E = [W_e, W_{tt}] \quad (4)$$

where $W_e \in \mathbb{R}^{v \times \alpha d}$, $W_{tt} \in \mathbb{R}^{v \times (1-\alpha)d}$, and $\alpha \in [0, 1]$ is used to control the proportion of the two kind of embeddings.

4.2 HTT Self-attention

In the original Transformer, the self-attention allows the input X to apply three projections to obtain Query, Key and Value representations. The self-attention can be computed as:

$$A = \frac{XW_q W_k^T X^T}{\sqrt{d}} \quad (5)$$
$$\text{Attention}(Q, K, V) = \text{softmax}(A)XW_v$$

where $X \in \mathbb{R}^{N \times d}$ is the input from the previous layer, N is the sequence length, and d is the model dimension. A is the attention matrix. The three weight matrices $W_q, W_k, W_v \in \mathbb{R}^{d \times d}$ can be column-wise concatenated as a single matrices $W \in \mathbb{R}^{d \times 3d}$, which can be formulated as:

$$W = [W_q, W_k, W_v] \quad (6)$$

Therefore, what we need to compress is the W matrix and it consists of only an HTT layer. HTT splits the W matrix into a dense part and a Tensor-Train part with 3 TT-cores (Figure 2(E)). This procedure can be represented as follows:

$$W = [W_{dense}, W_{tt}] \quad (7)$$

where $W_{dense} \in \mathbb{R}^{d \times 3\beta d}$, $W_{tt} \in \mathbb{R}^{d \times 3(1-\beta)d}$, and similar to HTT EMB, $\beta \in [0, 1]$ is used to control the proportion of the dense part and TT part.

In HTT SAN, the input X first goes through a dense layer, and the output is split into three equal parts Q_1, K_1, V_1. Second, the input X goes through a low-rank TT layer, and the output is also split into three equal parts Q_2, K_2, V_2. Finally, we obtain the Q, K and V by connecting related slices. The procedure can be formulated as:

$$
\begin{aligned}
Q_1, K_1, V_1 &= \text{Split}(XW_{dense}, 3) \\
Q_2, K_2, V_2 &= \text{Split}(XW_{tt}, 3) \\
Q &= \text{Concat}(Q_1, Q_2) \\
K &= \text{Concat}(K_1, K_2) \\
V &= \text{Concat}(V_1, V_2)
\end{aligned} \quad (8)
$$

4.3 LMF Feed-Forward Network

In this work, we do not apply HTT decomposition to FFN, because the performance of Transformer is less sensitive to parameter numbers in FFN than the counterparts of SAN (Hsu et al., 2020; Yang et al., 2020). Based on HTT Transformer, we apply the LMF layer to replace the HTT layer in FFN, which brings a further speed improvement (Table 2). Therefore, we decide to use the LMF layer in FFN. The function of the Feed-Forward Network (FFN) is to make the input X perform a non-linear transformation. It can be defined as:

$$\text{FFN}(X) = \text{ReLU}(XW_1 + b_1)W_2 + b_2, \quad (9)$$

where $W_1 \in \mathbb{R}^{d \times d_{ff}}$, $W_2 \in \mathbb{R}^{d_{ff} \times d}$, $b_1 \in \mathbb{R}^{d_{ff}}$ and $b_2 \in \mathbb{R}^d$. d is the model dimension, d_{ff} is the dimension of the FFN. To compress the FFN, we propose Low-rank Matrix Factorized FFN (LMF FFN). Unlike the two dense layers in the FFN, LMF FFN consists of four dense layers, which can be defined as follows:

$$\text{LMF-FFN}(X) = \text{ReLU}(XU_1 V_1 + b_1)U_2 V_2 + b_2, \quad (10)$$

where $U_1 \in \mathbb{R}^{d \times R}$, $V_1 \in \mathbb{R}^{R \times d_{ff}}$, $U_2 \in \mathbb{R}^{d_{ff} \times R}$, $V_2 \in \mathbb{R}^{R \times d}$, $b_1 \in \mathbb{R}^{d_{ff}}$ and $b_2 \in \mathbb{R}^d$. R is the rank of MF.

5 Experiment

5.1 Datasets and Evaluation

We evaluate our methods on three standard translation benchmark datasets: IWSLT'14 De-En (De-En), WMT'14 En-De (En-De), and WMT'16 En-Ro (En-Ro). For De-En, we use the same setup as in Liu et al. (2020), which consists of 160K sentence pairs and 10K joint byte pair encoding (BPE) (Sennrich et al., 2016) vocabulary. For En-De, we follow the setup as in Liu et al. (2020), which includes 3.9M training sentence pairs for and 37K joint BPE vocabulary. For En-Ro, we follow the setup of Lee et al. (2018), which includes 610K training sentence pairs.

5.2 Model Settings

Deep Encoder, Shallow Decoder The vanilla Transformer (Vaswani et al., 2017) adopts 6 encoder layers and 6 decoder layers. Besides the **6-6** setting, we choose a deep encoder shallow decoder setting that assigns 12 encoder layers and 1/2 decoder layers. Assigning more layers on encoders than decoders is beneficial for inference

speed while maintaining its performance (Li et al., 2021c) – this sometimes needs knowledge distillation (Kasai et al., 2020).

Knowledge Distillation Deep to shallow Transformer can benefit more from Sequence-level Knowledge Distillation (SKD) (Kim and Rush, 2016), which has been widely verified in previous work (Li et al., 2021b; Kasai et al., 2020). In our experiments, we utilize the method of SKD for training. Firstly, we train a teacher model with ground truth data. Then we utilize the teacher model to predict the training data and obtain SKD data. Especially, we use ground truth and SKD data together in student training. In this work, we use the 12-2 Transformer base model as our *teacher model* of knowledge distillation.

Hypoformer Hypoformer is based on a **12-2** or **12-1** Transformer base model (Vaswani et al., 2017) with the initialization of Admin (Liu et al., 2020). We use three of our proposed decomposition sub-layers (HTT EMB, HTT SAN and LMF FFN) to replace the original sub-layers (Embedding, The self-attention and FFN). We control the parameters of the model by adjusting the hyperparameters of the decomposition sub-layers and model dimensions. HTT EMB and HTT SAN have three hyperparameters: the dense part ratio α, the TT-cores D, and TT-ranks R in TT parts. LMF FFN has one hyperparameter: rank R. The detailed model settings are given in Table 5 in Appendix.

Baselines and Implementation We compare Hypoformer with decomposition methods and recent lightweight SOTA methods, including **Transformer** (Vaswani et al., 2017), Transformer with the Admin initialization (Liu et al., 2020) (denoted as **Our impl.**), **MF Transformer**, **TT Transformer**, **Lite Transformer** (Wu et al., 2020), **HAT Transformer** (Wang et al., 2020), **DeLighT** (Mehta et al., 2021), and **Subformer** (Reid et al., 2021). The implementation of all models uses Fairseq (Ott et al., 2019), including baselines and our methods. We reproduce the results of baselines following the setting from their papers or download trained models from official GitHub. The details of **measuring speed** are given in Appendix A.4.

5.3 Experimental Results

In this section, we compare our method with two different methods (decomposition methods and

lightweight SOTA methods) to verify the effectiveness of our Hypoformer.

The BLEU score of Hypoformer is superior to other decomposition-based Transformers, especially in small configurations. Table 2 shows the results of different decomposition methods applied in Transformer at two model scales. (1) On the base model scale, Hypoformer outperforms the MF Transformer and HTT Transformer in En-De. Compared with HMF Transformer, Hypoformer gets 0.7 BLEU improvement in En-De and achieves similar performance in En-Ro but using fewer parameters. (2) On the small model scale, Hypoformer achieves better BLEU scores than TT, MF, HMF, and HTT Transformers on two datasets. It indicates that Hypoformer is more effective than other decomposition methods in low-resource scenarios.

Hypoformer is smaller, faster and better-performed than recent light-weight SOTA methods. In Table 3, we first compare the results between Hypoformer with previous works in the case of compressing Transformer base model (Vaswani et al., 2017) on three tasks. Under the similar or even better performance on De-En and En-De, Hypoformer (12-2) compresses Transformer $2.9 \sim 4.4 \times$ parameters and Hypoformer (12-1) accelerates Transformer $2.5 \sim 2.8 \times$ on Intel CPU and $1.8 \sim 3.9\times$ on Raspberry Pi CPU, which also outperforms recent light-weight the SOTA methods in compression, acceleration and performance. In addition, we also conduct an experiment on a large dataset (WMT'14 En-Fr with 36M training pairs) with SKD data. The results can be found in Table 9 in Appendix.

Hypoformer consistently outperforms recent light-weight methods under different parameters and inference speeds. Previous works evaluate their methods on different model scales on De-En and En-De tasks. To further verify the effectiveness of our method, we also compare Hypoformer to these model scales. As shown in Figure 4(A) and (C), Hypoformer outperforms Transformer (Admin), Lite Transformer and DeLighT with fewer parameters and better performance matching the similar parameter scale. For the inference speedup (Figure 4(B) and (D)), Hypoformer achieves a significantly higher speedup than Lite Transformer and DeLighT when scaled down the model size. Overall, Hypoformer is more friendly than these existing methods in edge devices.

Hypoformer gets more competitive performance

	Model	E-D	WMT'14 En-De				WMT'16 En-Ro			
			Params	MACs	Speed	BLEU	Params	MACs	Speed	BLEU
Base	Transformer (Vaswani et al., 2017)	6-6	61.0 M	7.9 B	26.3 tokens/s	27.3	62.1M	7.9 B	23.2 tokens/s	**34.3**◊
	Transformer 12-2 (Our impl.)	12-2	64.9 M	3.3 B	41.5 tokens/s	**28.7**	64.1M	3.3 B	41.5 tokens/s	**34.3**
	MF Transformer	12-2	20.7 M	**1.8 B**	42.2 tokens/s	27.3	14.8 M	**1.7 B**	42.6 tokens/s	33.8
	HMF Transformer	12-2	21.0 M	2.0 B	40.2 tokens/s	26.8	20.8 M	2.0 B	41.1 tokens/s	**34.4**
	HTT Transformer	12-2	19.4 M	2.3 B	29.3 tokens/s	26.5	19.2 M	2.3 B	31.0 tokens/s	**34.3**
	Hypoformer	12-2	21.2 M	2.2 B	42.0 tokens/s	27.5	13.9 M	1.9 B	**44.0 tokens/s**	**34.3**
Small	TT Transformer (TT-ranks=2)	12-2	1.0 M	0.9 B	55.3 tokens/s	8.6	1.0 M	0.9 B	55.3 tokens/s	13.5
	TT Transformer (TT-ranks=32)	12-2	5.8 M	8.3 B	37.6 tokens/s	23.2	5.8 M	8.3 B	37.6 tokens/s	31.7
	MF Transformer	12-2	5.8 M	1.0 B	**62.7 tokens/s**	22.5	5.7 M	1.0 B	62.7 tokens/s	31.6
	HMF Transformer	12-2	5.7 M	1.0 B	59.1 tokens/s	20.5	5.6 M	1.0 B	59.1 tokens/s	29.9
	HTT Tranformer	12-2	5.7 M	1.1 B	50.5 tokens/s	22.5	5.6 M	1.1 B	50.5 tokens/s	31.5
	Hypoformer	12-2	5.7 M	**0.6 B**	60.1 tokens/s	23.6	5.6 M	**0.6 B**	**76.0 tokens/s**	**32.7**

Table 2: The performance comparison among decomposed Transformers in TT, MF, HTT, HMF, and Hypoformer on WMT'14 En-De and WMT'16 En-Ro. Hypoformer obtains the best BLEU score compared with other decomposition methods. **Speed** denotes translation speed [tokens/s] on the Intel CPU. All models apply knowledge distillation and initialization of Admin except for the vanilla Transformer. **E-D**: the numbers of the encoder (**E**) and the decoder (**D**). ◊ indicates the results is our implementation from Mehta et al. (2021). We compute multiplication-addition operations (MACs) using 20 source and 20 target tokens. **Params**: the whole model parameters including embedding layer. We control the same α of HTT and HMF for fair comparison.

Model	IWSLT'14 De-En					WMT'14 En-De					WMT'16 En-Ro				
	Params	Ratio	$S_{(Pi)}$	$S_{(Intel)}$	BLEU	Params	Ratio	$S_{(Pi)}$	$S_{(Intel)}$	BLEU	Params	Ratio	$S_{(Pi)}$	$S_{(Intel)}$	BLEU
Transformer	36.8 M	1.0×	1.0×	1.0×	34.5	61.0 M	1.0×	1.0×	1.0×	27.3	62.1 M	1.0×	1.0×	1.0×	**34.3**
Lite Transformer	13.9 M	2.7×	1.5×	1.5×	33.6	33.6 M	1.8×	0.8×	1.1×	26.5	-	-	-	-	-
HAT Transformer	35.2 M	1.1×	1.8×	1.7×	34.5	46.2 M	1.3×	1.5×	1.7×	26.9	-	-	-	-	-
DeLighT	19.9 M	1.9×	1.0×	0.8×	34.4	23.3 M	2.6×	1.3×	1.2×	26.7	22.0 M	2.8×	0.6×	1.2×	**34.3**
Subformer	-	-	-	-	-	38.0 M	1.6×	-	-	**27.7**	20.0 M	3.1×	-	-	34.1
Hypoformer 12-2	**8.4 M**	**4.4×**	3.5×	2.7×	**34.8**	21.2 M	**2.9×**	1.5×	1.6×	27.5	13.9 M	**4.5×**	3.0×	1.9×	**34.3**
Hypoformer 12-1	8.6 M	4.3×	**3.9×**	**3.3×**	34.4	23.6 M	2.6×	**1.8×**	**2.8×**	27.4	15.0 M	4.1×	**3.6×**	**3.3×**	33.7

Table 3: Comparison of compression ratio and speedup with Light-Weight Transformers on three tasks. Best performance is **bolded**. **Ratio**: dividing the parameters of Light-Weight Transformers by parameters of Transformer. $S_{(Pi)}$ and $S_{(Intel)}$: the inference speedup of Raspberry Pi-4 CPU and Intel CPU, and the speedup indicates the division between the inference speed of Light-Weight Transformers and Transformer. **Hypoformer X-X**: the X numbers of encoder and decoder Hypoformer with knowledge distillation. The results of Transformer (Vaswani et al., 2017) are from the original paper.

under extremely low parameters. Notably, Hypoformer achieves 7.1 and 4.0 higher BLEU score than Transformer (Admin) when two models scale down to about 2.8M and 5.1M parameters on De-En and En-De (Figure 4(A) and (C)). These results indicate that Hypoformer's effectiveness is more significant in extremely low-resource scenarios.

5.4 Ablation Study and Analysis

Ablation Study. In Table 4, we show how the vanilla Transformer could be ablated into Hypoformer. From scheme **1**∼**3**, we observe that Transformer with initialization of Admin and deep shallow structure brings an improvement of performance and 1.6× speedup. We further apply our Hypofomer on this setting (scheme **4**), which leads to performance degradation but can compress 3.8× parameters and 21% speedup. However, the gap

in performance between Hypoformer and Transformer can be made up with knowledge distillation (scheme **5**).

From scheme **6**∼**8**, we observe that HTT EMB (Only compressing embedding) can reduce embedding parameters without impact on speed because look-up operations from a restored embedding table are fast. HTT SAN (Only compressing self-attention) and HMF FFN (Only compressing FFN) can both reduce about 31% model parameters and promote about 11% speedup. In the aspect of performance, three components combined do not bring a significant drop of performance and get higher compression and acceleration compared to adding three components respectively.

Knowledge Distillation Analysis. As shown in Figure 5, firstly, we compared the results of Hypoformer and the model without distillation. We
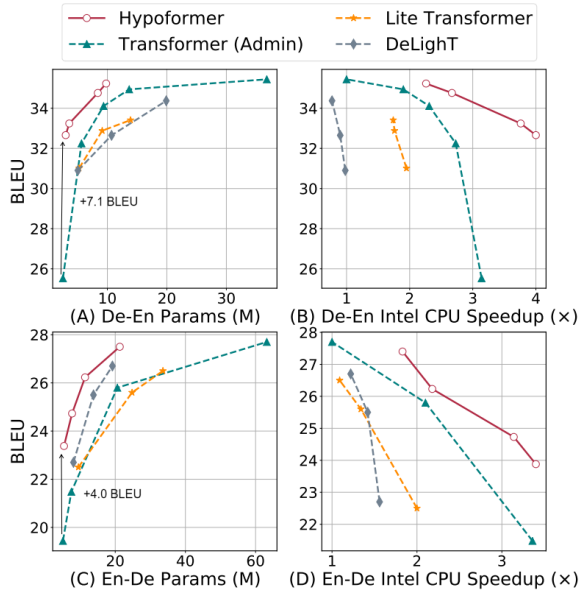
Figure 4: The comparison of performance between Hypoformer and SOTA light-weight methods on IWSLT'14 De-En (**Top**) and WMT'14 En-De (**Down**). **Left**: under different parameters. **Right**: under different inference speedup of Intel CPU. Transformer (Admin): Transformer with the initialization of Admin (Liu et al., 2020). We scaled down Transformer (Admin) by adjusting the model dimension and reproduced results of different settings in Lite Transformer and DeLighT.

	Model	#Params	#Total	$S_{(Intel)}$	BLEU
1	Transformer	31.6 M	36.8 M	1.0 ×	34.5
2	+ Admin	31.6 M	36.8 M	1.0 ×	34.9
3	+ 12-2	31.6 M	36.8 M	1.9 ×	35.4
4	+ Hypoformer	7.2 M	9.8 M	2.3 ×	33.3
5	+ Hypoformer†	7.2 M	9.8 M	2.3 ×	35.2
6	+ HTT EMB †	31.6 M	34.2 M	2.0 ×	36.0
7	+ HTT SAN †	21.7 M	24.3 M	2.1 ×	35.5
8	+ LMF FFN †	22.2 M	24.8 M	2.1 ×	35.4

Table 4: Ablation study on IWSLT'14 De-En. #Params and #Total indicate the model parameters without and with embedding respectively, † denotes that the models apply knowledge distillation. For scheme **1**∼**5**, (+) indicates that a result includes all preceding methods. For scheme **6**∼**8**, (+) indicates that a proposed decomposition applies on scheme **3** with knowledge distillation.
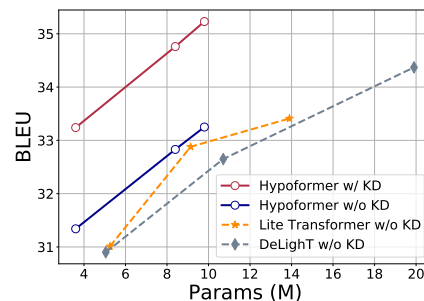

Figure 5: Distillation analysis in different model size on IWSLT'14 De-En.

found that knowledge distillation can improve our models by about 2 BLEU scores which is a significant improvement. Secondly, without using knowledge distillation, the BLEU of **Hypoformer w/o KD** is still better than Lite Transformer and DeLighT in different model scales. Moreover, we provide other analysis results in the Appendix A.6.

6 Related Work

Low-rank Approximation is a common and effective technique to compress deep neural networks. Matrix Factorization, which expresses a matrix $W \in \mathbb{R}^{I \times J}$ as a product of two smaller matrices $U \in \mathbb{R}^{I \times R}$ and $V \in \mathbb{R}^{R \times J}$. Hybrid Matrix Factorization (Thakker et al., 2020) is an extension of MF, which inspired our Hybrid Tensor Train decomposition. Tensor-Train decomposition is one part of our methods, which has proved its effectiveness in compression in previous works(Novikov et al., 2015; Hrinchuk et al., 2020; Liu et al., 2021).

Light-Weight Transformer is an increasing demand for the deployment of Transformer in resource-constrained scenarios. Recent works explore different methods to compress and accelerate

Transformer. One line of research is light-weight variants(Wu et al., 2020; Mehta et al., 2021). Our Hypoformer can also combine with another line of research, which we leave the experimental investigation in future work including efficient self-attention mechanism (Zhou et al., 2021; Tay et al., 2020), quantization (Prato et al., 2019), Neural Architecture Search (Wang et al., 2020), and parameter sharing (Reid et al., 2021).

7 Conclusion

In this paper, we present an HTT decomposition-based Transformer, named *Hypoformer*, which can compress and accelerate Transformer. We explore the application of Hypoformer on edge devices. In three standard translation tasks, Hypoformer is consistently better than recent light-weight SOTA methods under various parameter and speed scales. In edge devices with extremely limited resources, Hypoformer can retain **95%** performance with **1/12** parameters, setting a new SOTA in extremely-small NMT models.

Limitations

Since the proposed HTT is orthogonal to existing compression methods such as quantization, we leave the experimental investigation of *the hybrid between HTT and other compression methods* as future work. In addition, HTT module could be further accelerated by using a dedicated CUDA kernel, which we leave as future work since it is also the main focus in this work.

Acknowledgements

References

Alexandre Bérard, Dain Lee, Stéphane Clinchant, Kweonwoo Jung, and Vassilina Nikoulina. 2021. Efficient inference for multilingual neural machine translation. *ArXiv*, abs/2109.06679.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

Oleksii Hrinchuk, Valentin Khrulkov, Leyla Mirvakhabova, Elena Orlova, and Ivan V. Oseledets. 2020. Tensorized embedding layers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 4847–4860. Association for Computational Linguistics.

Yi-Te Hsu, Sarthak Garg, Yi-Hsiu Liao, and Ilya Chatsviorkin. 2020. Efficient inference for neural machine translation. *CoRR*, abs/2010.02416.

Jungo Kasai, Nikolaos Pappas, Hao Peng, James Cross, and Noah A. Smith. 2020. Deep encoder, shallow decoder: Reevaluating the speed-quality tradeoff in machine translation. *CoRR*, abs/2006.10369.

Yoon Kim and Alexander M. Rush. 2016. Sequence-level knowledge distillation. In *EMNLP*.

Jason Lee, Elman Mansimov, and Kyunghyun Cho. 2018. Deterministic non-autoregressive neural sequence modeling by iterative refinement. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 1173–1182. Association for Computational Linguistics.

Bei Li, Ziyang Wang, Hui Liu, Quan Du, Tong Xiao, Chunliang Zhang, and Jingbo Zhu. 2021a. Learning light-weight translation models from deep transformer. In *AAAI*.

Bei Li, Ziyang Wang, Hui Liu, Quan Du, Tong Xiao, Chunliang Zhang, and Jingbo Zhu. 2021b. Learning light-weight translation models from deep transformer. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 13217–13225. AAAI Press.

Yanyang Li, Ye Lin, Tong Xiao, and Jingbo Zhu. 2021c. An efficient transformer decoder with compressed sub-layers. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 13315–13323. AAAI Press.

Liyuan Liu, Xiaodong Liu, Jianfeng Gao, Weizhu Chen, and Jiawei Han. 2020. Understanding the difficulty of training transformers. In *EMNLP*.

Peiyu Liu, Ze-Feng Gao, Wayne Xin Zhao, Zhi-Yuan Xie, Zhong-Yi Lu, and Ji-Rong Wen. 2021. Enabling lightweight fine-tuning for pre-trained language model compression based on matrix product operators. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 5388–5398. Association for Computational Linguistics.

Xindian Ma, Peng Zhang, Shuai Zhang, Nan Duan, Yuexian Hou, Ming Zhou, and Dawei Song. 2019. A tensorized transformer for language modeling. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 2229–2239.

Sachin Mehta, Marjan Ghazvininejad, Srinivasan Iyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2021. Delight: Deep and light-weight transformer. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

---

[2] https://www.mindspore.cn/

Matan Ben Noach and Yoav Goldberg. 2020. Compressing pre-trained language models by matrix decomposition. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing, AACL/IJCNLP 2020, Suzhou, China, December 4-7, 2020*, pages 884–889. Association for Computational Linguistics.

Alexander Novikov, D. Podoprikhin, A. Osokin, and D. Vetrov. 2015. Tensorizing neural networks. In *NIPS*.

I. Oseledets. 2011. Tensor-train decomposition. *SIAM J. Sci. Comput.*, 33:2295–2317.

Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Demonstrations*, pages 48–53. Association for Computational Linguistics.

Yu Pan, Maolin Wang, and Zenglin Xu. 2021. Tednet: A pytorch toolkit for tensor decomposition networks. *CoRR*, abs/2104.05018.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.

Roger Penrose. 1971. Applications of negative dimensional tensors. *Combinatorial mathematics and its applications*, 1:221–244.

Gabriele Prato, Ella Charlaix, and Mehdi Rezagholizadeh. 2019. Fully quantized transformer for machine translation. *arXiv preprint arXiv:1910.10485*.

Machel Reid, Edison Marrese-Taylor, and Yutaka Matsuo. 2021. Subformer: Exploring weight sharing for parameter efficiency in generative transformers. *CoRR*, abs/2101.00234.

Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.

Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. 2020. Efficient transformers: A survey. *arXiv preprint arXiv:2009.06732*.

Urmish Thakker, Jesse G. Beu, Dibakar Gope, Ganesh S. Dasika, and Matthew Mattina. 2020. Rank and runtime aware compression of nlp applications. *ArXiv*, abs/2010.03193.

A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. Attention is all you need. *arXiv*.

Hanrui Wang, Zhanghao Wu, Zhijian Liu, Han Cai, Ligeng Zhu, Chuang Gan, and Song Han. 2020. Hat: Hardware-aware transformers for efficient natural language processing. In *ACL*.

Felix Wu, Angela Fan, Alexei Baevski, Yann N. Dauphin, and Michael Auli. 2019. Pay less attention with lightweight and dynamic convolutions. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.

Zhanghao Wu, Zhijian Liu, Ji Lin, Yujun Lin, and Song Han. 2020. Lite transformer with long-short range attention. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Yilin Yang, Longyue Wang, Shuming Shi, Prasad Tadepalli, Stefan Lee, and Zhaopeng Tu. 2020. On the sub-layer functionalities of transformer decoder. In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 4799–4811. Association for Computational Linguistics.

Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of AAAI*.

A Appendix for "Hypoformer: Hybrid Decomposition Transformer for Edge-friendly Neural Machine Translation"

A.1 Preliminaries: Tensor Operation

Matrix/tensor multiplication can be performed between two tensors. Two nodes share one leg means that these two tensors perform matrix/tensor multiplication in the corresponding dimension, sometimes called 'tensor contraction'. A tensor can perform contraction with one or more tensors by sharing legs together, and different orders of contraction execution result in an identical result. As shown in Figure 1(D), three tensors perform two matrix multiplications, and the result is the matrix $D \in \mathbb{R}^{K \times I}$. Figure 2(A) illustrates a fully connected layer. That is, the input vector $X \in \mathbb{R}^I$ and the weight matrix $W \in \mathbb{R}^{I \times J}$ performs matrix multiplication, and the result is a vector $\in \mathbb{R}^J$.

Tasks	E-D	d_{model}	d_{ff}	HTT EMB†			HTT SAN			LMF FFN	Speed(CPU)	Params	BLEU
				α	D_1	R_1	α	D_2	R_2	R			
De-En	12-2	512	1024	0.5	3	4	0.25	3	2	64	95 tokens/s	9.8 M	35.2
	12-1	512	1024	0.5	3	4	0.25	3	2	64	140 tokens/s	8.6 M	34.4
	12-2	512	1024	0.5	3	4	0.25	3	2	32	112 tokens/s	8.4 M	34.8
	12-2	256	1024	0.5	3	4	0.25	3	4	32	158 tokens/s	3.6 M	33.4
	12-2	256	1024	0.125	3	4	0.25	3	4	16	168 tokens/s	3.0 M	32.7
En-De	12-2	512	2048	0.25	3	16	0.5	2	4	128	42 tokens/s	21.3 M	27.5
	12-1	512	2048	0.25	3	16	0.5	2	4	256	49 tokens/s	23.6 M	27.4
	12-2	384	2048	0.25	3	16	0.5	2	3	96	59 tokens/s	11.4 M	26.2
	12-2	256	1024	0.25	3	16	0.5	2	3	96	90 tokens/s	7.7 M	24.7
	12-1	256	1024	0.25	3	16	0.5	2	3	96	96 tokens/s	7.1 M	23.9
	12-2	256	512	0.125	3	16	0.5	2	3	64	87 tokens/s	5.4 M	23.4
En-Ro	12-2	512	2048	0.125	3	32	0.25	3	4	96	44 tokens/s	13.9 M	34.3
	12-2	256	1024	0.25	3	16	0.25	3	4	64	76 tokens/s	5.6 M	32.7

Table 5: The overall results of Hypoformer in three tasks. **E-D** indicates the numbers of encoder (**E**) and decoder (**D**). d_{model} denotes the model dimension, and d_{ff} indicates the hidden state dimensions of FFN. † indicates that the TT-ranks of HTT EMB don't impact the inference speed because we restore the look-up table embedding in the processing of model inference.

Hyperparameter	Value
label smoothing	0.1
max tokens	4096
distributed world size	1
update frequency	1
dropout rate	[0.05, 0.1, 0.2]
embedding dim	[512, 256]
ffn dim	1024
Attn heads	4
optimizer	radam
Adam-betas	(0.9, 0.98)
lr	7e-4
lr scheduler	invert sqrt
warmup lr	1e-7
warmup updates	6000
max updates	20K
fp16	True
fp16 scale window	256
threshold loss scale	0.03125
initialization	Admin
beam	5
length penalty	1.0

Table 6: Training settings of De-En.

Hyperparameter	Value
label smoothing	[0.05, 0.1]
max tokens	4096
distributed world size	8
update frequency	1
dropout rate	[0, 0.05, 0.1]
embedding dim	[512, 384, 256]
ffn dim	[2048, 1024, 512]
Attn heads	8
optimizer	radam
Adam-betas	(0.9, 0.98)
lr	1e-3
lr scheduler	invert sqrt
warmup lr	1e-7
warmup updates	8000
max updates	60K
fp16	True
fp16 scale window	256
fp16 scale tolerance	0.25
threshold loss scale	0.03125
initialization	Admin
beam	4
length penalty	0.6

Table 7: Training settings of En-De and En-Ro.

A.2 Hybrid Tensor-Train Decomposition: why HTT can always keep full the theoretical maximum rank.

Referring to Hrinchuk et al. ((Hrinchuk et al., 2020)), TT can theoretically retain the full matrix rank . Therefore, *the TT part* in HTT is in principle full-rank w.r.t. the approximated matrix. The *dense part* is by definition full-rank. Thus, HTT can retain full matrix rank. Technically, the dense part

($\mathbf{M}_{dense}$) is a $N \times \alpha N$ matrix and TT part ($\mathbf{M}_{tt}$) is to approximate a $N \times (1 - \alpha)N$ matrix. We have

$$
\begin{aligned}
\text{MaxRank}(\mathbf{M}_{dense}) &= \alpha N, \\
\text{MaxRank}(\mathbf{M}_{tt}) &= (1 - \alpha)N, \\
\text{MaxRank}(\mathbf{M}_{tt}; \mathbf{M}_{dense}]) &\leq \alpha N + (1 - \alpha)N = N
\end{aligned}
\tag{11}
$$

Note that, the theoretical maximum rank is guaranteed in almost all cases excluding a zero-measure negligible set, which is irrelevant with dense layer

ratio α and TT-ranks R.

A.3 Experiment: Training Setting

Our models follow the training setups of (Liu et al., 2020). We trained our models with 1 RTX 2080Ti GPU on IWSLT'14 De-En, 8 Telsa V100 GPUs on WMT'14 En-De and 4 For WMT'16 En-Ro. The last 10 model checkpoints are averaged for testing on En-De/En-Ro and the lowest perplexity checkpoint for testing on De-En. The detailed hyperparameters are listed in Table 6 and 7. Our TT decomposition code is based on TedNet (Pan et al., 2021).

For evaluation, we use beam search decoding in three tasks, where there is beam 5 for De-En and beam 4 and length penalty 0.6 for En-De and En-Ro. The performance is measured by case-sensitive tokenized BLEU (Papineni et al., 2002) for all translation tasks.

A.4 Experiment: Measuring Speed

We do not use FLOPs as a speed metric because Wang et al. (2020) found that FLOPs do not reflect the real-time latency in autoregressive Transformer. We instead use the number of translated tokens per second, denoted as *tokens/s*, as the inference speed metric. We sample 50 sentences of an average sequence length (23 for IWSLT and 30 for WMT) to test speed. We run these samples 10 times and remove 10% for the fastest and slowest results respectively and average the rest 80% results. We test the speed on two representative devices: Raspberry Pi-4 with an ARM Cortex-A72 CPU, 1 core Intel Xeon E5-2678 v3 @ 2.50GHz CPU. We evaluate the inference speed with a batch size of 1 to simulate the inference of edge devices.

Dataset	Model	E-D	Params	Ratio	$S_{(Intel)}$	BLEU
De-En	Transformer	6-6	36.8 M	1.0×	1.0×	34.9
	Transformer	12-2	36.8 M	1.0×	1.9×	**35.4**
	Hypoformer[†]	12-2	**9.8 M**	**3.8×**	**2.3×**	35.2
En-De	Transformer	6-6	63.2 M	1.0×	1.0×	27.7
	Transformer	12-2	65.3 M	1.0×	**1.5×**	**28.3**
	Hypoformer[†]	12-2	**21.2 M**	**3.0×**	**1.5×**	27.5
En-Ro	Transformer	6-6	61.4 M	1.0×	1.0×	**34.5**
	Transformer	12-2	64.2 M	1.0×	1.6×	34.4
	Hypoformer[†]	12-2	**13.9 M**	**4.4×**	**1.9×**	34.3

Table 8: The compression ratio and the speedup comparison with deep to shallow Transformer and Hypoformer on three tasks. All models initialize with Admin. **E-D**: the numbers of encoder (**E**) and decoder (**D**). † denotes that the model utilizes knowledge distillation.
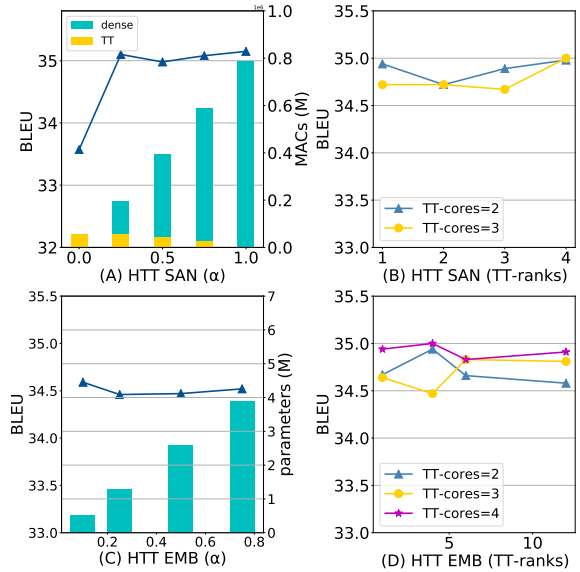

Figure 6: Sensitivity analysis on IWSLT'14 De-En. we calculate the MACs of HTT SAN and the embedding parameters in HTT EMB.

Model	Parameters	BLEU
Transformer (Vaswani et al., 2017)	62 M	38.1
Lite Transformer	40 M	39.6
DeLighT	37 M	39.6
Hypoformer 12-2	22 M	39.7

Table 9: The Comparison of BLEU on WMT'14 En-Fr

A.5 Experiment Results: WMT'14 En-Fr

In Table 9, in order to evaluate the performance our Hypoformer on a larger dataset, we conduct an experiment on WMT'14 En-Fr (36M training pairs) with SKD data. The results show that Hypoformer achieves similar performance with Lite Transformer and DeLighT but saves 15M and 18M parameters respectively. The structure of the Hypoformer we used in En-Fr is the same as in En-De. The results of Lite Transformer and DeLighT come from their original paper.

A.6 Analysis

Sensitivity Analysis. In Figure 6, we present a sensitivity analysis about the hyperparameters of Hybrid Tensor-Train: the ratio of the dense part α and the TT-ranks of the TT part. With the increase of the α, the BLEU score reaches a high point at $\alpha = 0.25$ and then tends to be flat in Figure 6(A). As shown in Figure 6(B), the BLEU score has a slight drop or improvement when TT-cores and TT-ranks have changed. For the hyperparameters of HTT EMB in Figure 6(C) and Figure 6(D), the

BLEU score has a little change by adjusting the ratio of the dense part α or the TT-cores and TT-ranks of TT part.

Deep to Shallow Setting. Table 8 shows the results of Hypoformer and Transformer (Admin) on three tasks. First, we note that the deep to shallow Transformer (12-2) can get similar performance with symmetric Transformer (6-6) but becomes roughly $1.6 \times$ faster, seeing a similar conclusion in (Li et al., 2021c). Moreover, compared with Transformer (12-2) baseline, our method can compress $3.0 \sim 4.4 \times$, accelerate $0 \sim 21\%$ in Intel CPU with a slight performance drop. Note that HTT module could be further accelerated by using a dedicated CUDA kernel, which we leave as future work since it is also the main focus in this work.

Regularization. *Low parameter models need less Regularization.* With the same model of 11.4 M parameters, Hypoformer achieves better BLEU scores (25.7→26.2) with using less dropout (0.1→0.05) and label-smoothing (0.1→0.05) on WMT'14 En-De.

A.7 Related Work: Sequence-Level Knowledge Distillation

Sequence-Level Knowledge Distillation (Kim and Rush, 2016) is a branch of knowledge distillation for seq2seq model. The generated sequences from the teacher model can be treated as the sequence-level knowledge to guide the student network training, which can make up the performance of deep shallow transformers (Kasai et al., 2020; Li et al., 2021b).

The goal of Knowledge Distillation is to transfer the knowledge learned from a large teacher network to the small student network. The procedure of KD can formulated like:

$$\mathcal{L}_{\text{KD}} = \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \mathcal{L}\left(f^T(x, y), f^S(x, y)\right) \quad (12)$$

where $\mathcal{L}(\cdot)$ is a loss function to evaluate the gap between the teacher predict $f^T(x, y)$ and the student predict $f^S(x, y)$. In machine translation, (x, y) is a pair of inputs of the source and target language. $(\mathcal{X}, \mathcal{Y})$ is the represents the training dataset.