

# Revisiting Offline Compression: Going Beyond Factorization-based Methods for Transformer Language Models

Mohammadreza Banaei<sup>\*1</sup>, Klaudia Bałazy<sup>\*2</sup>, Artur Kasymov<sup>2</sup>  
Rémi Lebret<sup>1</sup>, Jacek Tabor<sup>2</sup>, Karl Aberer<sup>1</sup>

<sup>1</sup>EPFL, <sup>2</sup>Jagiellonian University

## Abstract

Recent transformer language models achieve outstanding results in many natural language processing (NLP) tasks. However, their enormous size often makes them impractical on memory-constrained devices, requiring practitioners to compress them to smaller networks. In this paper, we explore offline compression methods, meaning computationally-cheap approaches that do not require further fine-tuning of the compressed model. We challenge the classical matrix factorization methods by proposing a novel, better-performing autoencoder-based framework. We perform a comprehensive ablation study of our approach, examining its different aspects over a diverse set of evaluation settings. Moreover, we show that enabling collaboration between modules across layers by compressing certain modules together positively impacts the final model performance. Experiments on various NLP tasks demonstrate that our approach significantly outperforms commonly used factorization-based offline compression methods.<sup>1</sup>

## 1 Introduction

The recent trend of pre-training Transformer (Vaswani et al., 2017) language models on enormous unsupervised corpus has led to outstanding performances on many downstream tasks. For downstream tasks, these pre-trained models are then either fine-tuned (Devlin et al., 2019; Liu et al., 2019; Yang et al., 2019) or the *prompting* paradigm (Brown et al., 2020) is used (especially in the so-called Large Language Models), which avoids having a different model per task (Lieber et al., 2021; Rae et al., 2021; Smith et al., 2022; Thoppilan et al., 2022). In each of the two paradigms, it has been shown

that increasing the scale of language models generally leads to better performance on a range of downstream tasks (Devlin et al., 2019; Brown et al., 2020). Indeed, for autoregressive language models, Kaplan et al. (2020) demonstrated a power-law relationship between the number of parameters and the respective performance. Wei et al. (2022) further showed that certain abilities of language models emerge only when the number of its parameters passes certain thresholds, providing an incentive to scale these models further.

Although scaling up these language models make them empirically powerful across many diverse tasks, it makes them infeasible to train for many NLP practitioners due to huge pre-training costs. More importantly, even using the available pre-trained models for inference is becoming more challenging (especially for memory-constrained applications like edge devices), with recent models having hundreds of billions of parameters (Zhang et al., 2022).

With the rise of NLP model sizes, there have been many efforts to compress transformer language models without compromising their performance. Although being inherently different, many of these efforts rely on knowledge distillation (Hinton et al., 2015) to help the compressed model (i.e., the student model) better imitate the parent model (i.e., the teacher model). However, these approaches often need costly distillation on upstream or downstream tasks (Sanh et al., 2019) as well as expensive data augmentation techniques (Jiao et al., 2019) to help improve the compressed model performance. These approaches become even less feasible when enormous language models are being distilled.

Another line of research focuses on computationally-cheap methods (i.e., offline compression) where a smaller model can be achieved from a pre-trained model without it being necessarily fine-tuned over a downstream

<sup>\*</sup>Equal contribution

Correspondence to: mohammadreza.banaei@epfl.ch, klaudia.balazy@doctoral.uj.edu.pl

<sup>1</sup>Our code is public: [github.com/MohammadrezaBanaei/auto-encoder-based-transformer-compression](https://github.com/MohammadrezaBanaei/auto-encoder-based-transformer-compression)

or upstream task. These offline methods include weight pruning (Li et al., 2016; Han et al., 2015), weight quantization (Zhou et al., 2016; Hubara et al., 2016), tensor factorization (Lan et al., 2019; Winata et al., 2019; Bałazy et al., 2021; Cordonnier et al., 2020) and hybrid approaches (Wang et al., 2019; Mao et al., 2020).

This paper proposes a novel offline factorization-based method for compressing transformer language models. The paper’s main goal is to propose an offline method that produces a competitive language model (compared to the original model perplexity) before any fine-tuning is performed. Similar to Bałazy et al. (2021), we use an autoencoder model (see Figure 1a) to compress different modules’ weights. However, unlike the previous work, our approach is not limited to the token embeddings (a.k.a. word embeddings) and can be applied to other transformer modules as well. We also propose and thoroughly investigate the impact of various enhancements for the approach of obtaining the compressed model. It is worth noting that although the experiments and ablation studies are only done BERT<sub>BASE</sub> model, our

In Section 4 we demonstrate that applying small changes to the autoencoder architecture (e.g., introducing non-linearity to the decoder) and its loss objective results in superior performance to the Singular Value Decomposition (SVD) baseline as measured by model perplexity and its performance on the downstream tasks. Moreover, inspired by the redundancies present across self-attention heads (Cordonnier et al., 2020), in Section 4.2 we show that the compressed models perform in general better when compressing certain modules from different layers together.

Additionally, in Section 4.6 we investigate the effectiveness of a (parameter) sensitivity-based<sup>2</sup> compression by incorporating fisher information (Pascanu and Bengio, 2013) in the loss objective. We later show that incorporating these weights significantly improve the compressed language model performance (i.e., perplexity).

Finally, in Section 4.7 and in Section 4.8 we discuss the performance of our approach in comparison to various offline-compression baselines and demonstrate that our method provides the best or competitive quality of the compressed model.

Our main contribution can be summarized as

<sup>2</sup>We call it sensitivity as it measures how sensitive the model performance is to the reconstruction error of a certain parameter.

follows:

- We propose a novel autoencoder-based framework for low-cost compression of transformer language models and conduct an extensive ablation study on its different aspects.
- We show that enabling collaboration across layers by compressing different layers modules together boosts the performance.
- We demonstrate that our approach significantly outperforms other commonly used offline-compression methods on various NLP downstream tasks.<sup>3</sup>

## 2 Related work

Deep transformer language models have gained increasing attention in recent years since the seminal work of Devlin et al. (2019). Many recent efforts demonstrate that scaling up these language models’ parameters generally results in better performance on a range of downstream tasks (Devlin et al., 2019; Brown et al., 2020; Kaplan et al., 2020). This empirical observation resulted in recent language models having over a thousand times more parameters (Lieber et al., 2021; Rae et al., 2021; Smith et al., 2022) than the BERT<sub>BASE</sub> model (Devlin et al., 2019). Although empirically powerful, these models are becoming harder to use for memory-constrained applications, which led to many efforts toward language model compression in recent literature.

Although many recent efforts for language model compression take advantage of distillation (Hinton et al., 2015) techniques to better imitate the uncompressed model (i.e., the teacher model) behavior, this paper focuses mainly on *offline* compression methods. By *offline*, we refer to approaches that do not need fine-tuning the whole model on a downstream/upstream dataset. In the case of language model compression, these methods aim to output a compressed model without losing too much performance (measured by perplexity) that can then be fine-tuned (with or without distillation) or prompted for a certain downstream task. It is worth noting that these methods can still be combined with distillation techniques, but starting the finetuning from a better compressed

<sup>3</sup>It is worth noting that although the experiments and ablation studies in this paper are only done on the BERT<sub>BASE</sub> model, our proposed approach can potentially be used for any transformer-based architecture.

language model would generally reduce its costs of training (e.g., by improving convergence time).

Offline compression methods, while being diverse, can be roughly categorized into few paradigms, namely weight pruning (See et al., 2016; Li et al., 2016; Han et al., 2015; Fan et al., 2019; Michel et al., 2019; Voita et al., 2019), quantization (Gong et al., 2014; Hubara et al., 2016; Zhou et al., 2016), tensor factorization (Lan et al., 2019; Winata et al., 2019; Bałazy et al., 2021; Cordonnier et al., 2020; Panahi et al., 2021; Ren et al., 2022) and hybrid approaches (Wang et al., 2019; Mao et al., 2020).

This paper primarily focuses on the effectiveness of low-rank factorization-based approaches in recent literature (Lan et al., 2019; Panahi et al., 2021; Ren et al., 2022) as an *offline* compression method. Lan et al. (2019) proposed a SVD-based (Halko et al., 2011) technique to compress the token embedding module. Later works have shown that more complex architectures like autoencoders can result in better compression quality than SVD methods (Lioutas et al., 2020; Bałazy et al., 2021). By taking advantage of the autoencoder, we are able to more easily enforce different properties by either changing its training objective or architecture (e.g., preserving  $l_2$  norm in reconstructed embeddings). Bałazy et al. (2021) emphasized on the importance of *direction* in token embedding compression, and in this work we demonstrate its potential importance for other transformer modules as well in different compression ratios.

Moreover, Cordonnier et al. (2020) showed the significance of redundant information in self-attention heads and compressed different heads (in a certain layer) together to improve compression performance. Following a similar idea, we later show that compressing heads from different layers together would generally further boost the compression quality.

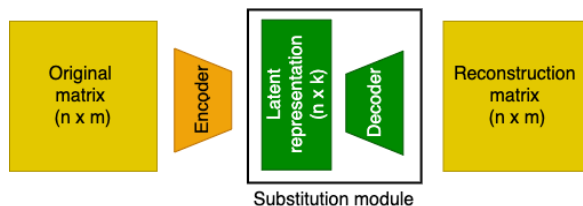
Hsu et al. (2022) proposed a weighted SVD (using Fisher Information) (Pascanu and Bengio, 2013) to outperform the classical SVD. We further investigate the benefits of a non-uniform compression (i.e., a weighted reconstruction loss in the autoencoder loss objective) in Section 4.6 by analyzing different weighting schemes for parameters.

Moreover, Ren et al. (2022) proposed using tensor decomposition techniques to compress language models to relatively high compression ratios while using a two-stage distillation technique.

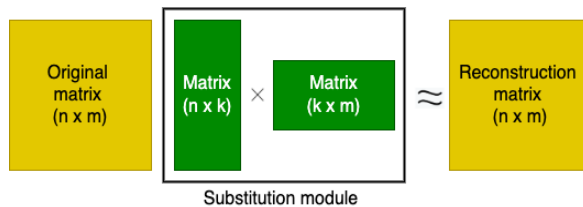
Moreover, Panahi et al. (2021) proposes using the Kronecker product as an alternative for the factorization of transformer modules. Appendix A.7 discusses using Tucker (De Lathauwer et al., 2000) or Kronecker-based methods as an offline approach. It is worth noting that models with relatively high compression ratios become highly dependent on distillation techniques to perform reasonably on downstream tasks. For instance, Ren et al. (2022) claims that even randomly initializing the compressed BERT nearly achieves identical performance compared to tensor decomposition from a pre-trained model.

### 3 Model

Our offline compression approach is based on the an autoencoder neural network architecture, similar to Lioutas et al. (2020) and Bałazy et al. (2021). However, in this work, we focus on compressing all the transformer weight matrices rather than just the token embedding matrix. Furthermore, we are exploring many more compression improvements using autoencoder as well as investigating architecture-independent techniques.



(a) Autoencoder-based compression with customizable objective function. Our approach minimizes the root mean square error (RMSE) and cosine distance between the original and reconstruction matrix. In this setting, the original matrix’s latent representation and the decoder form the substitution module.



(b) Classical matrix-factorisation-based compression with root mean square error (RMSE) minimization objective. Two smaller matrices, approximating the target matrix after multiplication, form the substitution module.

Figure 1: A high-level view of the matrix compression approaches using classical matrix-factorization and autoencoder model that is leveraged in this work. The purpose of the compression is to provide a parameter-efficient substitution module to replace the original matrix in the considered model.

The autoencoder architecture consists of encoder function  $h(\cdot)$  that maps model input  $x \in \mathbb{R}^m$  to some latent representation  $l \in \mathbb{R}^k$ . The second element of the architecture is the decoder function  $g(\cdot)$  responsible for mapping  $l \in \mathbb{R}^k$  into the approximation  $\tilde{x} \in \mathbb{R}^m$  of the input  $x$ .

Let us assume that we want to compress the matrix  $A \in \mathbb{R}^{(n \times m)}$ . Using the gradient-descent algorithm, we train the autoencoder model to produce the appropriate approximation  $\tilde{A}$  of the original matrix. As a compressed module, we understand the hidden representation  $h_\Psi(A) \in \mathbb{R}^{(n \times k)}$  together with the decoding module  $g_\Phi(\cdot)$ . In this setting, the formula for the *compression ratio* of the original module can be expressed as:

$$\frac{n \cdot m}{(n \cdot k) + |\Phi|}, \quad (1)$$

which is the ratio of the original matrix size to the hidden representation size and the number of parameters in the decoder module  $|\Phi|$ . We illustrated the approach of compressing a matrix using the autoencoder model in Figure 1a.

Our approach to offline compression based on the autoencoder offers flexibility in performing the ablation study as we are able to easily modify its elements, for example, decoder module complexity level or the loss function components. By using an autoencoder architecture with a linear decoder and the RMSE cost function, we can obtain the equivalent approximation as provided by a simple matrix factorization. An illustrative comparison of our compression method with the classical matrix factorization approach is shown in Figure 1.

Following Bałazy et al. (2021), we train our autoencoder with the multi-objective cost function consisting of  $l_2$  norm loss and cosine distance loss:

$$\Psi_\beta(X, \tilde{X}) = (1 - \beta) \cdot L_2(X, \tilde{X}) + \beta \cdot CD(X, \tilde{X}), \quad (2)$$

where  $X$  represents the original matrix,  $\tilde{X}$  is the reconstructed matrix,  $L_2(X, \tilde{X})$  represents the root mean square error (RMSE) loss function, and  $CD(X, \tilde{X})$  is the mean cosine distance loss for all pairs of vectors (rows) of the original and reconstructed matrices. The  $\beta$  hyperparameter ( $0 \leq \beta \leq 1$ ) is responsible for determining the weight we would like to assign to the different components of the loss function.

## 4 Experiments

This section describes our motivations and the results of various analyses and experiments that we

conducted to investigate the topic of offline compression thoroughly.

We performed our experiments for different weight matrices in the transformer architecture, as each type of weight matrix may have different characteristics, and a given compression method may or may not be appropriate. The analyses described below are performed for token embedding, self-attention (keys, queries, values), and output-dense weight matrices.

We focus our study on the BERT<sub>BASE</sub> model (Devlin et al., 2019), but the same methods could be applied to other transformer architectures. All experiments are conducted for three compression ratios (3, 10, and 25) to investigate the differences given the different number of available parameters. All experimental settings of the various studies presented in the following sections are included in Appendix A.

We evaluate the quality of our compressed models on the masked (Devlin et al., 2019) language modeling task (using the WikiText-103 test dataset (Merity et al., 2016)) and multiple datasets from the GLUE benchmark (Wang et al., 2018).

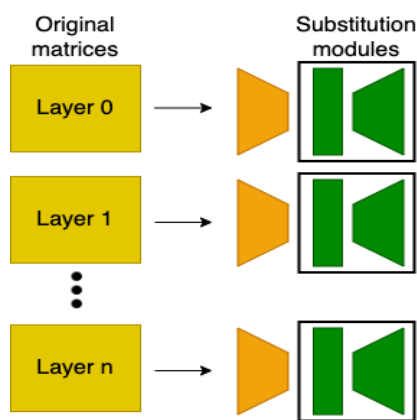
### 4.1 Cosine distance objective

First, we investigate whether including the direction component in the compression objective has a positive effect on the compression of weight matrices other than token embeddings in the transformer model. Bałazy et al. (2021) demonstrated that supplementing the loss function with the cosine distance between pairs of rows of the original and reconstructed matrix produces noticeably better compression results for the token embeddings matrix. Unfortunately, their study does not examine other matrices in the transformer, whereas because of the different nature of these matrices, we believe it is worth investigating.

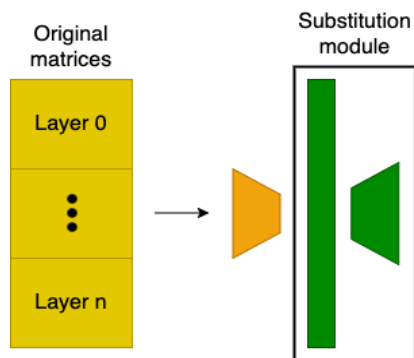
**Results** In Table 3 and Table 4 (in Appendix A.1), we present the effect of adding the cosine distance component to the cost function for the keys and output-dense matrices from the BERT<sub>BASE</sub> model. It seems that for matrices other than token embeddings, considering the direction of vectors (rows in the matrix) in most cases may have a positive impact on the final results of the compressed model. However, the benefits of using this component are not as significant as in the case of the token embeddings matrix. Indeed, there are some examples where minimizing only Euclidean distance or

adding only a small proportion of cosine distance provides the best results. We suppose this behavior is the consequence of the token embeddings matrix nature, where the rows represent specific tokens used to construct the words. It seems that the representation of the part or entire word is largely encoded in the vector direction. This characterization does not necessarily apply to other matrices though often there is a subtle benefit from adding a cosine distance component to the reconstruction objective.

## 4.2 Concatenated and separated weight matrices



(a) Separated matrices compression mode. All substitution modules have separate decoder.



(b) Concatenated matrices compression mode. All substitution modules share the decoder module that enables collaboration between layers.

Figure 2: Separated and concatenated weight matrices compression. We demonstrate that compressing concatenated matrices from all layers provides a better-performing substitution module. In the concatenated setting, the substitution modules share a decoder that allows for cross-layer collaboration and provides the potential to further eliminate redundant information.

This section investigates whether compressing separately each weight matrix from the considered model is the best possible strategy. Our in-

tuition is that compressing together the same type of weight matrices from different layers may bring certain advantages. First, it could allow for minimizing redundant information in the model weights, and second, it could enable collaboration between compressed modules across different layers. Suppose that the neural network model consists of  $n$  layers ( $l_0, l_1, \dots, l_{n-1}$ ). Each layer  $l_i$  encapsulates a particular weight matrix  $W_{l_i}$ . Conventionally, each  $W_{l_i}$  matrix is considered separately during the compression process. In the concatenated mode, we propose compressing a single matrix  $W = [W_{l_0}, W_{l_1}, \dots, W_{l_{n-1}}]$  resulted from concatenating all  $W_{l_i}$  matrices. Given the proposed compression process, the compressed weight matrices share a common decoder as illustrated in Figure 2.

**Results** Experiments discussed in this section demonstrate that compressing concatenated weight matrices performs better than compressing each matrix separately in terms of the compressed model performance as well as the compression process time. Figure 3 presents the performance achieved by models with compressed output-dense matrices in separated and concatenated modes (similar experiments are presented for key, query, and value matrices in Figure 7 in Appendix A.2). We report the initial perplexity and the final score achieved on the MRPC and SST2 downstream tasks for different compression ratios. We observe the apparent dominance of the concatenated mode over the separated mode for both perplexity and downstream task performance. This may indicate that sharing the decoder helps to reduce redundant information and saves parameters for further knowledge encoding.

Furthermore, the separated compression mode is more computationally expensive at the initial stage as we must compress each matrix individually. In the concatenated mode, we perform only a single compression process on the matrices' concatenation.

Considering the better performance and faster training process, we only analyze the concatenated weight matrices compression in the following sections.

## 4.3 Initial perplexity vs downstream tasks performance

Perplexity is a popular measure determining how well the language model predicts a particular se-

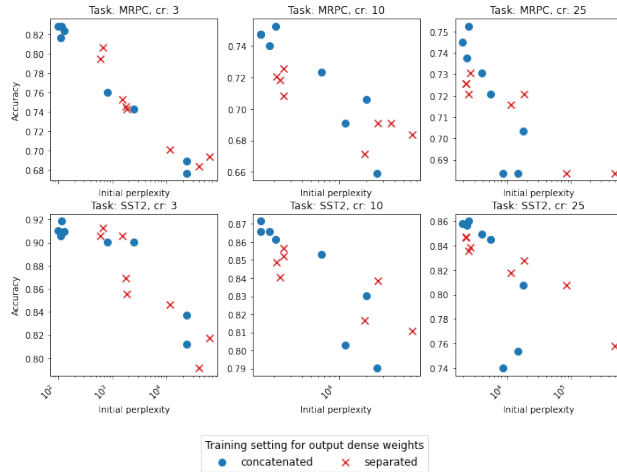


Figure 3: Initial perplexity and downstream tasks performance of output-dense matrices’ compression in separated/concatenated compression modes<sup>4</sup>. Using the concatenated mode generally results in better performance for this module. We observe similar pattern for key, query, and value weight matrices (see Figure 7 in Appendix A.2).

quence of tokens (the lower the perplexity the better). This section introduces that the masked language model perplexity metric may be considered as a low-cost yet effective way to evaluate a compressed module. We show that in most cases the compressed models with the lowest initial perplexity yields the best performance when fine-tuned on a downstream tasks.

**Results** We examined the relation between the initial  $BERT_{BASE}$  perplexity after applying compressed weight matrices and its final performance after fine-tuning on a downstream task. In Figure 8 (in Appendix A.3) we report results for token embeddings matrix, self-attention keys, queries, values and final output-dense matrices. We observe that in most cases models with the lowest initial perplexity result in the best performance on the downstream task (MRPC and SST2). Therefore, we consider the masked language model perplexity metric to be a good low-cost method to preliminarily evaluate the quality of a compressed module.

#### 4.4 Linear and non-linear decoder module

In this section we investigate the effect of using different decoder module in the autoencoder model on the final model’s performance. We experiment with a simple linear layer decoder and two non-linear decoder versions.

**Results** Figure 4 presents initial perplexity and final downstream tasks performance achieved when using linear and non-linear decoder in the autoen-

coder model while compressing token embeddings matrix. We may observe that for the token embeddings better final results are produced when using non-linear decoder. However, as demonstrated in Figure 9 (in Appendix A.4), a different pattern is apparent for key matrices where linear models considerably outperform the non-linear versions in most cases.

#### 4.5 Preserving vector norm

Furthermore, we examine whether preserving the original  $l_2$  vector norms of the vectors representing rows in the reconstructed matrix to be the same as in the original vectors is beneficial for the compression.

**Results** Figure 5 presents initial perplexity and downstream tasks performance when enabling or disabling the preserving vector norm technique for the token embeddings matrix. We may see that in most cases the version with enabled preserving vector norm achieves better results. In addition to token embeddings, in Figure 10 (in Appendix A.5) we also demonstrate the effect of preserving  $l_2$  vector norm during compression of the output-dense matrix.

#### 4.6 Sensitivity

Most offline compression methods focus only on the raw weight matrices taken from the considered pre-trained model. However, we could also leverage the unsupervised upstream dataset to improve the compression quality. Hsu et al. (2022) proposed using additional weights computed on the

<sup>4</sup>Each point represents one hyperparameter setting.

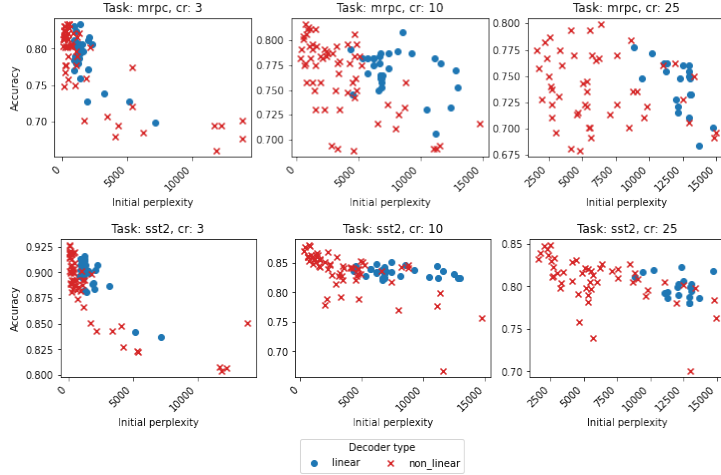


Figure 4: Initial perplexity and downstream tasks performance for the compressed token embeddings matrix when using either a linear or a non-linear decoder module in the autoencoder model. Here we present results for different hyperparameters settings (see Appendix A.4 for further details). A non-linear decoder seems to be a better choice for compressing this module.

entire upstream dataset to enhance the low-rank factorization method. They used the Fisher Information weights  $I$  that measure the amount of the observable information in dataset  $D$  about a single model parameter  $w$ . A feasible approximation  $\hat{I}_w$  of the Fisher Information  $I_w$  for parameter  $w$  may be expressed as:

$$I_w = E[(\frac{\partial}{\partial w} \log P(D|w))^2] \approx \frac{1}{|D|} \sum_{d \in D} (\frac{\partial}{\partial w} L(d; w))^2 = \hat{I}_w \quad (3)$$

where  $L$  is the target pre-training task objective (e.g., cross-entropy or MSE).

For the entire weights matrix  $W$ , Hsu et al. (2022) presented even more simplified and computationally effective row-wise diagonal Fisher Information matrix  $\hat{I}$ , where each diagonal value is the sum of the corresponding row of the Fisher Information approximation matrix  $\hat{I}_W$ :

$$\hat{I} = \text{diag}(\sqrt{\sum_{j=1} \hat{I}_{W_{1j}}}, \dots, \sqrt{\sum_{j=1} \hat{I}_{W_{nj}}}). \quad (4)$$

We present the distributions of the row-wise Fisher Information for the upstream dataset (i.e., masked language modeling on the WikiText-103 dataset) in Figure 6. We notice that each distribution contains some outliers which point to the potentially irrelevant weights in the considered weight matrix. In this section, we demonstrate that the weights’ relevance information may be leveraged in the compression process to improve the quality of compressed modules.

Hsu et al. (2022) used the Fisher Information directly on the original model weights in their Fisher-Weighted SVD (FWSVD) approach:

$$W \approx FWSVD(W) = \hat{I}^{-1} SVD(\hat{I}W). \quad (5)$$

In contrast, our method does not modify the original weight matrices but rather uses the Fisher Information in the loss function to help the model focus more on important weights. Moreover, we apply different transformations on the original Fisher Information values to modify the relative importance of the module weights to reduce the undesirable influence of outliers (Appendix A.6 discusses various transformations we experimented with for different modules).

CR	Method	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)
3	AE	118.41	91.97	<b>88.53 / 84.31</b>
	AE+Fisher	33.27	<b>92.55</b>	88.36 / 83.33
10	AE	712.98	88.07	85.87 / 80.88
	AE+Fisher	250.59	<b>89.33</b>	<b>87.27 / 81.62</b>
25	AE	4926.08	82.80	84.19 / 77.45
	AE+Fisher	2728.41	<b>83.83</b>	<b>84.35 / 77.70</b>

Table 1: The effect of adding the Fisher Information to the autoencoder-based (AE) compression of token embeddings. We report the compressed BERT<sub>BASE</sub> upstream task perplexity (on the WikiText-103 dataset) and the downstream performance over two GLUE tasks. Each AE result represents a median from 3 runs with different seeds.

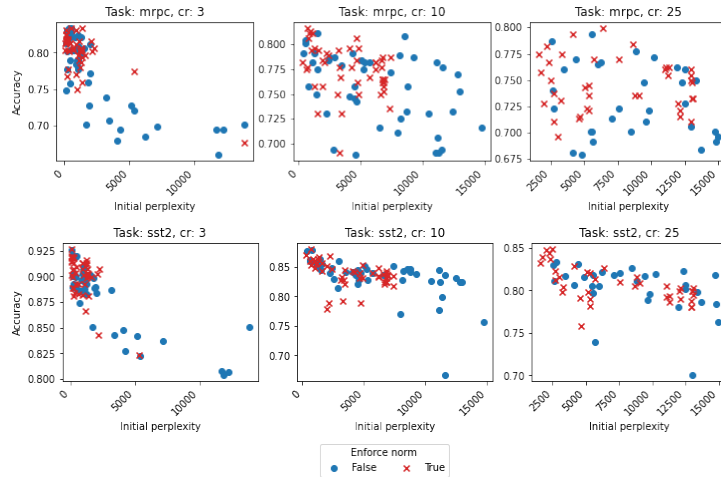


Figure 5: The effect of preserving  $l_2$  vector norm on the perplexity and downstream tasks performance while compressing the token embeddings matrix for different hyperparameters settings. Preserving norm seems to generally improve the compression performance for this module.

CR	Architecture	MRPC (F1/Acc)	SST-2 (Acc)	RTE (Acc)	QNLI (Acc)	QQP (F1/Acc)
1	BERT <sub>BASE</sub>	88.85/84.07	92.32	65.70	90.66	87.49/90.71
3	SVD	84.35/77.45	86.70	62.09	85.61	84.64/88.38
	Our	<b>85.25/77.77</b>	<b>90.25</b>	<b>62.45</b>	<b>88.68</b>	<b>86.07/89.71</b>
10	SVD	78.46/68.38	82.00	52.71	77.48	79.36/83.45
	Our	<b>81.45/71.08</b>	<b>83.94</b>	<b>57.76</b>	<b>81.48</b>	<b>81.57/85.75</b>
25	SVD	77.10/66.42	78.33	53.43	62.66	74.21/79.10
	Our	<b>81.60/71.32</b>	<b>80.05</b>	<b>55.23</b>	<b>72.96</b>	<b>78.45/83.11</b>

Table 2: Final BERT<sub>BASE</sub> model compression (token embedding matrix, all key matrices, and all output-dense matrices). The baseline SVD algorithm compresses each matrix separately. Our autoencoder-based approach, incorporates mechanisms developed in the ablation study presented in this work (see Table 8 for the detailed AE design choices). For each setting, we present the median score from experiments with three different seeds. Our approach consistently outperforms the classical factorization method.

**Results** Table 1 presents the benefits of incorporating sensitivity for compression of token embeddings where both upstream perplexity and downstream task performance is improved. We further demonstrate the positive influence of Fisher information for three transformer modules in Table 5 (in Appendix A.6) for both autoencoder and SVD methods. Additionally, Table 6 (in Appendix A.6) presents the compression performance provided by AE using different Fisher Information transformations. We observe that incorporating the Fisher Information with batch normalization into the compression process considerably improves the model perplexity as well as the downstream task performance.

#### 4.7 Comparison with other offline compression approaches

In Table 5, we compare our approach with the most popular matrix factorization method, namely Singular Value Decomposition (SVD), for the compression of three different transformer modules. We may see that our approach outperforms or is competitive with SVD in most settings. Additionally, in Appendix A.7, we discuss the poor performance of Kronecker Product and Tucker Decomposition (as two other factorization-based methods) in the offline compression setting. We also compare our solution to a non-factorization baseline, namely pruning, and show that our autoencoder-based method also outperforms it in most studied settings (see Table 7 in Appendix A.7).



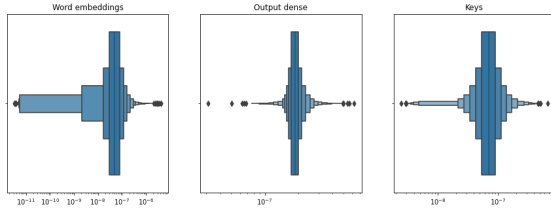


Figure 6: Row-wise Fisher Information distribution for three different modules in the  $BERT_{BASE}$  model. The Fisher information values are passed as importance weights in the autoencoder loss function to help the compression model focus more on the module’s important weights. Appendix A.6 discusses different transformations applied to Fisher information to help the compression model handle outlier Fisher information values.

#### 4.8 Compressing multiple types of transformer modules

For the final experiment, we took into consideration all the analysis insides presented in this work and prepared the concluding experiment on offline compression methods. In this experiment, we compressed multiple  $BERT_{BASE}$  weight matrices at the same time (token embedding matrix, all key matrices and all output-dense matrices). We compared the offline compression quality produced by the autoencoder approach and the baseline SVD factorization matrix method.

For the compression with SVD we classically compressed each matrix separately. We tested three different seeds and various number of iterations for SVD algorithm. For the autoencoder approach, we compressed considered matrices by selecting appropriate mechanisms based on our ablation study. For both SVD and our approach, we report the median of the final scores from the experiments with three different seeds to exclude potential outliers. The compressed matrices with the lowest perplexity were applied into  $BERT_{BASE}$  that was then fine-tuned on various NLP tasks (MRPC, SST-2, RTE, QNLI and QQP). In the resulting table we reported the median of the final scores to exclude potential outliers.

Final experiment results are presented in Table 2. Our approach consistently outperforms the SVD baseline on all tested downstream tasks.<sup>5</sup>

#### 4.9 Compression time

Generally, compressing modules using autoencoder and SVD takes a comparable amount of time. How-

<sup>5</sup>The hyperparameter setting for autoencoder is provided in Table 8 in Appendix A.8

ever, using concatenated mode (as proposed in our paper) speeds up this process significantly. In Appendix A.9, we report compression times for different modules (Table 9) and the compressed models’ inference and fine-tuning time.

## 5 Conclusions

This work comprehensively studies various methods for the offline compression of transformer language models. We analyze various changes in the proposed architecture and its optimization function. We test different input modifications and evaluate the compressed language model performance in each scenario. By analyzing various compression settings, we show that our autoencoder-based approach outperforms classical matrix factorization on various NLP downstream tasks. Furthermore, we believe the techniques analyzed in this study might also be useful for low-cost compression of different weight matrices unrelated to language models.

### Limitations

A limitation of our approach that we may identify is the need to analyze each module type to determine the best mechanisms for its compression. Our module-specific findings could be reflected in corresponding modules in other language models, but this would require further investigation. Additionally, A (reasonably-sized) unsupervised corpus must also be used for computing the Fisher Information for the compression procedure, which is more computationally demanding than other offline approaches suggested in this study.

### Acknowledgements

The work of Klaudia Bałazy was supported by the National Centre of Science (Poland) Grant No. 2020/39/D/ST6/01332. Klaudia Bałazy is affiliated with Doctoral School of Exact and Natural Sciences at the Jagiellonian University. The research of Jacek Tabor was carried out within the research project "Bio-inspired artificial neural network" (grant no. POIR.04.04.00-00-14DE/18-00) within the Team-Net program of the Foundation for Polish Science co-financed by the European Union under the European Regional Development Fund. Artur Kasymov work was supported by the National Centre of Science (Poland) Grant No. 2019/33/B/ST6/00894.

## References

- Klaudia Balaży, Mohammadreza Banaei, Rémi Lebret, Jacek Tabor, and Karl Aberer. 2021. Direction is what you need: Improving word embedding compression in large language models. In *Proceedings of the 6th Workshop on Representation Learning for NLP (RepL4NLP-2021)*, pages 322–330.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Jean-Baptiste Cordonnier, Andreas Loukas, and Martin Jaggi. 2020. Multi-head attention: Collaborate instead of concatenate. *arXiv preprint arXiv:2006.16362*.
- Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. 2000. A multilinear singular value decomposition. *SIAM journal on Matrix Analysis and Applications*, 21(4):1253–1278.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Angela Fan, Edouard Grave, and Armand Joulin. 2019. Reducing transformer depth on demand with structured dropout. *arXiv preprint arXiv:1909.11556*.
- Yunchao Gong, Liu Liu, Ming Yang, and Lubomir Bourdev. 2014. Compressing deep convolutional networks using vector quantization. *arXiv preprint arXiv:1412.6115*.
- Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. 2011. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288.
- Song Han, Jeff Pool, John Tran, and William Dally. 2015. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems*, 28.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Yen-Chang Hsu, Ting Hua, Sungeun Chang, Qian Lou, Yilin Shen, and Hongxia Jin. 2022. [Language model compression with weighted low-rank factorization](#). In *International Conference on Learning Representations*.
- Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2016. Binarized neural networks. *Advances in neural information processing systems*, 29.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2019. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). Cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. 2016. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*.
- Opher Lieber, Or Sharir, Barak Lenz, and Yoav Shoham. 2021. Jurassic-1: Technical details and evaluation. *White Paper. AI21 Labs*.
- Vasileios Lioutas, Ahmad Rashid, Krtin Kumar, Md Akmal Haidar, and Mehdi Rezagholizadeh. 2020. Improving word embedding factorization for compression using distilled nonlinear neural decomposition. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2774–2784.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Andrew L Maas, Awni Y Hannun, Andrew Y Ng, et al. 2013. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, page 3. Atlanta, Georgia, USA.
- Yihuan Mao, Yujing Wang, Chufan Wu, Chen Zhang, Yang Wang, Yaming Yang, Quanlu Zhang, Yunhai Tong, and Jing Bai. 2020. Ladabert: Lightweight adaptation of bert through hybrid model compression. *arXiv preprint arXiv:2004.04124*.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? *Advances in neural information processing systems*, 32.

- Aliakbar Panahi, Seyran Saeedi, and Tom Arodz. 2021. [Shapeshifter: a parameter-efficient transformer using factorized reshaped matrices](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 1337–1350. Curran Associates, Inc.
- Razvan Pascanu and Yoshua Bengio. 2013. Revisiting natural gradient for deep networks. *arXiv preprint arXiv:1301.3584*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. 2021. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*.
- Yuxin Ren, Benyou Wang, Lifeng Shang, Xin Jiang, and Qun Liu. 2022. Exploring extreme parameter compression for pre-trained language models. *arXiv preprint arXiv:2205.10036*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Abigail See, Minh-Thang Luong, and Christopher D Manning. 2016. Compression of neural machine translation models via pruning. *arXiv preprint arXiv:1606.09274*.
- Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. 2022. Using deep-speed and megatron to train megatron-turing nl-g 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990*.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762*.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. *arXiv preprint arXiv:1905.09418*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Ziheng Wang, Jeremy Wohlwend, and Tao Lei. 2019. Structured pruning of large language models. *arXiv preprint arXiv:1910.04732*.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Genta Indra Winata, Andrea Madotto, Jamin Shin, Elham J Barezi, and Pascale Fung. 2019. On the effectiveness of low-rank matrix factorization for lstm model compression. *arXiv preprint arXiv:1908.09982*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. 2016. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*.

## A Experiments

In this section, we describe the general assumptions for the experiments and the specific setting of the hyperparameters for each individual experiment.

We train the autoencoder model using gradient descent procedure and Adam optimizer (Kingma and Ba, 2014). In most of our experiments, resulting compressed modules are inserted into a pre-trained language model and fine-tuned on two different downstream tasks from GLUE benchmark (Wang et al., 2018), MRPC and SST2, with

a default learning rate  $\lambda = 2 \cdot 10^{-5}$  proposed by Hugging Face Transformers<sup>6</sup> (Wolf et al., 2019).

### A.1 Cosine distance objective

In Table 3 and Table 4, we report the effect adding the cosine distance component to the compression objective for the keys and output-dense matrices from the BERT<sub>BASE</sub> model.

**Experimental setup** We study the effect of the cosine distance component in the loss function formulated in Equation (2) on the compression quality of self-attention keys matrices and the fully-connected output-dense weight matrices. We compress each of these matrices from each layer separately and then apply all the compressed matrices from certain type (keys or output-denses) to the transformer model to evaluate the compression quality. We inspect the effect of various ratios between Euclidean distance component ( $L_2$ ) and the cosine distance component ( $CD$ ) in Equation (2), namely 1:0, 10:1, 1:1, 1:10 which corresponds to  $\beta \in \{0.0, 0.0909, 0.5, 0.909\}$ . For each model we test different learning rates  $\lambda \in \{5 \cdot 10^{-3}, 10^{-3}, 5 \cdot 10^{-4}, 10^{-4}\}$ .

Compression ratio	Cosine coefficient	Initial perplexity	MRPC (Acc)	SST-2 (Acc)
3	0.0	22.9	74.5	91.0
	0.0909	22.2	76.2	92.1
	0.5	25.5	77.2	92.1
	0.9091	27.0	76.2	91.4
10	0.0	45.8	72.5	91.5
	0.0909	32.5	72.0	90.7
	0.5	34.4	70.3	91.5
	0.9091	43.2	71.3	91.6
25	0.0	116.9	70.3	90.6
	0.0909	167.6	69.6	90.2
	0.5	115.3	69.1	90.9
	0.9091	507.7	69.3	90.1

Table 3: The impact of adding a component responsible for preserving the direction in the compression of the key matrices from the self-attention block in BERT<sub>BASE</sub> model. We report the best (across different learning rates) initial model’s perplexity after compression as well as it’s performance on a downstream tasks for the different cosine coefficients ( $\beta$  in Equation (2)). Adding a component to the compression objective that aims to preserve the direction of the vectors in reconstructed matrix may bring a slight improvement on the quality of the result.

<sup>6</sup><https://github.com/huggingface/transformers>

Compression ratio	Cosine coefficient	Initial perplexity	MRPC (Acc)	SST-2 (Acc)
3	0.0	327.4	83.1	91.0
	0.0909	332.8	82.3	91.5
	0.5	599.2	79.4	90.5
	0.9091	1510.4	75.2	90.5
10	0.0	1396.5	74.5	88.1
	0.0909	1917.8	74.5	85.9
	0.5	2099.3	72.0	84.9
	0.9091	2326.1	71.8	84.0
25	0.0	1988.0	73.8	84.5
	0.0909	2099.0	73.8	85.3
	0.5	2144.7	72.5	84.7
	0.9091	2148.6	72.5	84.6

Table 4: The impact of adding a component responsible for preserving the direction in the compression of the output-dense matrices in BERT<sub>BASE</sub> model. Different cosine coefficients refer to  $\beta$  in Equation (2). We report the score for the model with the best initial perplexity across different learning rates as well as it’s performance on MRPC and SST2 downstream tasks.

### A.2 Concatenated and separated weight matrices

In Figure 7 we present the performance achieved by models with compressed key, query and value matrices when using separated and concatenated mode in the compression process.

**Experimental setup** With the objective of comparing the compression of separate and concatenated matrices, we analyze the various matrices in the transformer architecture: the key, query, and value matrices from the self-attention module and output-dense matrix (one of the fully connected end matrices). We optimize loss function described previously in Equation (2) with 1:1 and 1:10 ratios for the  $l_2$  norm loss coefficient and cosine distance coefficient, respectively. The decoder in the autoencoder model is a single fully connected layer. The model is trained with different learning rates  $\lambda \in \{5 \cdot 10^{-3}, 10^{-3}, 5 \cdot 10^{-4}, 10^{-4}\}$ .

### A.3 Initial perplexity vs downstream tasks performance

In Figure 8 we present the initial perplexity and the performance on MRPC and SST2 downstream tasks for the language model with various compressed modules.

**Experimental setup** We examine the initial masked language model perplexity and downstream tasks performance relation for token embed-

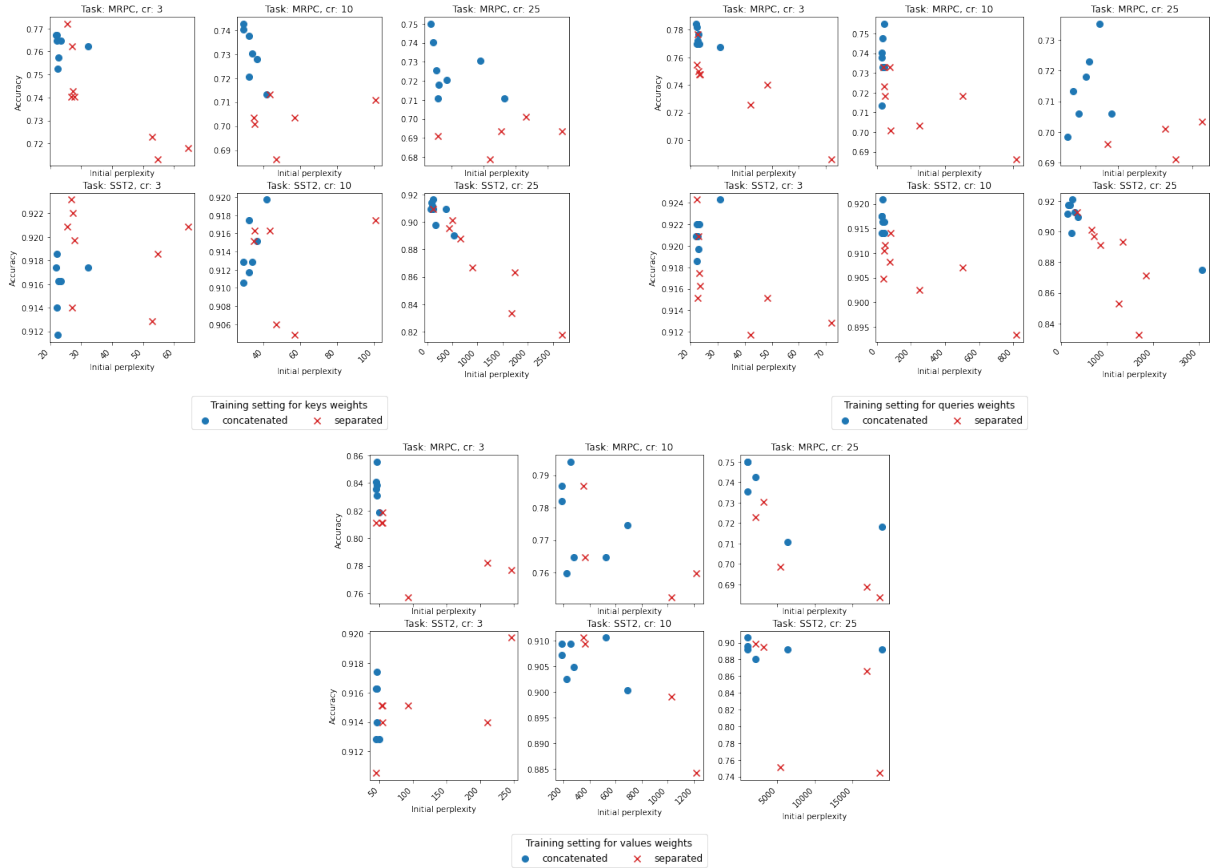


Figure 7: Initial perplexity and downstream tasks performance with separated and concatenated compression mode for output-dense weights matrices, queries weights matrices and values weights matrices in the BERT model.

ding matrix, self-attention matrices: keys, queries, values and final output-dense matrices. For simplicity of the experiment, we use the linear decoder in the autoencoder model. The loss and learning rate combinations for the models’ training are the same as in experiment in previous section. We conducted all experiments with two different seeds to obtain more reliable correlation information.

#### A.4 Linear and non-linear decoder module

In Figure 9 we present the initial perplexity and final downstream tasks performance achieved when using linear and non-linear decoder in the autoencoder model for key matrices compression. Unlike token embeddings for key matrices, the linear models outperform non-linear ones in most scenarios. Similar set of experiments on output-dense matrices also showed that linear models outperform non-linear ones.

**Experimental setup** We conduct experiments with the following settings: linear encoder/decoder and non-linear encoder/decoder. For non-linear encoder/decoder case, we examine architectures

with 1 and 2 hidden layers. As non-linear activation functions we use LeakyReLU (Maas et al., 2013) and Tanh. We investigate the loss configurations (Equation (2)) with 1:0, 1:1, 1:10 and 1:100 ratios of the  $l_2$  norm loss to the cosine distance loss.

#### A.5 Preserving vector norm

In Figure 10 we show the potential benefits of preserving  $l_2$  vector norm during compression of the output-dense matrices for two different downstream tasks .

**Experimental setup** We repeat the same set of experiments as in the previous section (Section 4.4), but each experiment is executed with and without the preserving vector norm option enabled.

#### A.6 Sensitivity

In Table 5 and Table 6 we report the positive impact of adding Fisher Information (with different coefficient transformations) for weight matrix compression for both autoencoder and SVD approaches.

**Experimental setup** We precompute the Fisher Information coefficients for token embeddings, self-

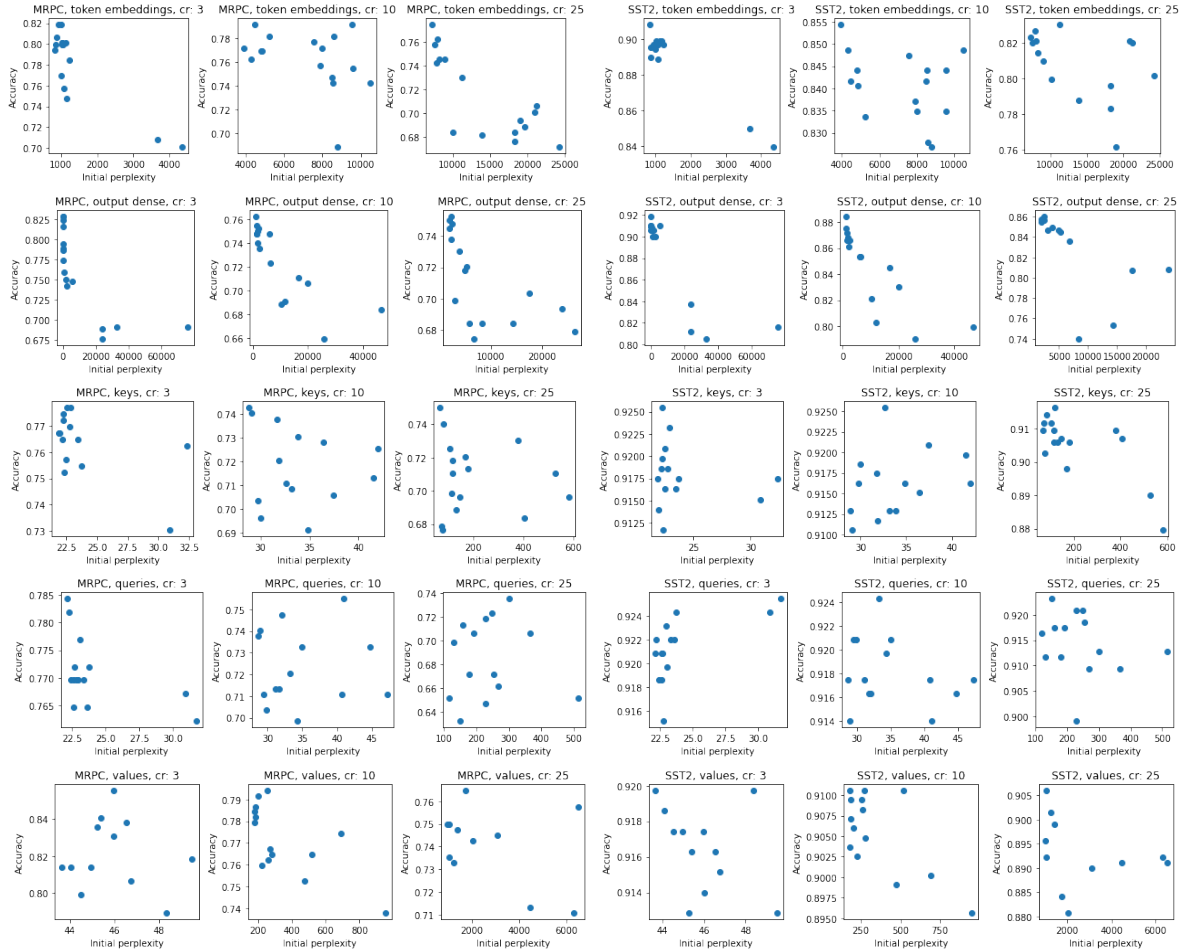


Figure 8: Initial perplexity and final accuracy on MRPC and SST2 downstream tasks for token embeddings, output-dense, keys, queries and values BERT<sub>BASE</sub> matrices compression.

attentions keys and output-dense weight matrices. We apply them the best hyperparameter setting obtained from the previous experiments. We use various transformations for Fisher Information coefficients: exponential transformation  $x^a$  with  $a = 0.1, 0.5, 0.9, 2.0$ , logarithmic transformation  $\log_e(x) + C$  with a  $C$  equals to minimum value so that all  $x$  are positive, logarithmic transformation  $\log_e(x) + C + 10$ , and raw Fisher Information coefficients without any transformation (*Vanilla*). For some transformations, we also add a batch sum normalization (+BN). We also report results without using Fisher Information (*No Fisher*).

### A.7 Comparison to other offline compression methods

**Kronecker Product offline compression** Inspired by a promising results achieved by using a Kronecker product for training the transformer model from scratch (Panahi et al., 2021) we have attempted to produce a compression of the orig-

inal transformer matrices by using a Kronecker product of two matrices approximating the original matrix. We have trained these matrices using the gradient descent algorithm. Unfortunately, the results were unsatisfactory for each of the tested settings. For example, for concatenated key matrices and a compression ratio of 10 the perplexity for the Kronecker product was around 1500, while for the autoencoder perplexity below 50 is achieved in many different settings.

### Tucker decomposition offline compression

Moreover, we also experimented with the Tucker decomposition (De Lathauwer et al., 2000) as an offline compression method. For token embeddings compression, we observed that the compressed language model starts having high perplexities even in low compression ratios. For instance, for CR=3, the model perplexity becomes almost 1500, while the autoencoder model can achieve perplexities below 40 for the same compression ratio. This finding is consistent with the observation of (Ren et al.,

CR	Method	Word embeddings			Output dense			Keys		
		Perplexity	SST-2 (Acc)	MRPC (F1/Acc)	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)
3	SVD	1842.38	89.68	83.71 / 75.49	99.98	<b>91.74</b>	<b>88.77 / 84.07</b>	22.04	92.32	85.30 / 77.45
	SVD+Fisher	20.20	91.86	<b>88.62</b> / 83.82	71.05	<b>91.74</b>	88.74 / 83.82	20.67	91.74	<b>86.22 / 78.92</b>
	AE	118.41	91.97	88.53 / <b>84.31</b>	70.69	<b>91.74</b>	86.17 / 78.92	22.35	<b>92.43</b>	84.48 / 76.23
	AE+Fisher	33.27	<b>92.55</b>	88.36 / 83.33	64.67	91.40	85.58 / 77.94	22.27	92.32	84.44 / 75.98
10	SVD	13196.30	83.37	82.85 / 74.02	989.18	89.56	83.92 / 75.49	30.75	<b>91.97</b>	<b>81.54 / 73.04</b>
	SVD+Fisher	65.17	87.73	85.01 / 77.70	723.61	<b>90.14</b>	<b>84.76 / 76.47</b>	41.49	91.74	81.00 / 72.06
	AE	712.98	88.07	85.87 / 80.88	1197.19	88.42	83.97 / 75.49	29.01	91.40	80.42 / 72.30
	AE+Fisher	250.59	<b>89.33</b>	<b>87.27 / 81.62</b>	1249.62	89.45	84.64 / 75.98	29.30	91.40	81.34 / 72.79
25	SVD	20178.74	77.64	82.33 / 71.81	1603.45	<b>88.19</b>	84.09 / 75.25	52.50	90.71	78.19 / 69.36
	SVD+Fisher	913.23	75.23	82.02 / 73.77	1205.34	87.27	83.88 / 74.75	80.14	91.17	74.36 / 65.69
	AE	4926.08	82.80	84.19 / 77.45	1462.41	85.61	82.02 / 72.06	69.24	<b>91.74</b>	<b>78.78 / 69.36</b>
	AE+Fisher	2728.41	<b>83.83</b>	<b>84.35 / 77.70</b>	1453.44	87.56	<b>84.54 / 75.98</b>	73.28	90.83	78.40 / 69.61

Table 5: The effect of adding the Fisher Information to the SVD-based and autoencoder-based (AE) compression. We report the  $BERT_{BASE}$  upstream task perplexity and the downstream tasks final scores. Each autoencoder result represents a median from 3 runs with different seeds and each SVD score is a result of the best iteration from run with one seed. For autoencoder model compression we selected the Fisher Information transformation for each of the compressed modules based on the results from Table 6 ( $x^{0.5} + BN$  for word embeddings;  $x^{2.0} + BN$  for output-dense matrices;  $\log_e(x) + C + 10$  for key matrices).

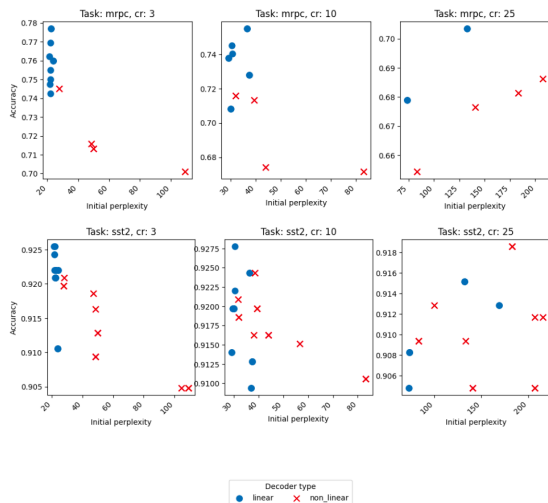


Figure 9: Initial perplexity and downstream tasks performance when using linear and non-linear decoder module in the autoencoder model for the compression of the keys matrix. Using a linear decoder generally appears to be a better choice for this matrix.

2022) that even randomly initializing the factorized tensors perform very close to the models initialized by e.g., Tucker decomposition. Therefore, we also do not find Tucker decomposition an efficient method in the context of offline compression.

**Pruning** We also compare our proposed autoencoder framework with a pruning baseline as another offline compression baseline. The pruning algorithm here is based on PyTorch unstructured L1 pruning (Paszke et al., 2019). For this experiment,

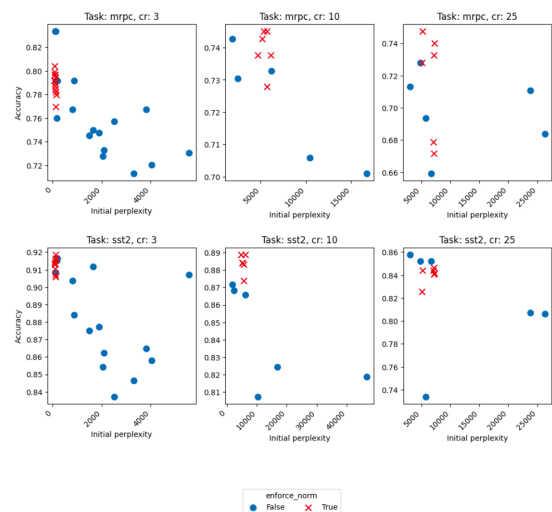


Figure 10: The effect of preserving  $l_2$  vector norm on the perplexity and downstream tasks performance while compressing the output-dense matrices. We can see that enforcing norm for compression of this module generally improves the result.

we compress token embedding, keys, and output-dense matrices using either autoencoder or pruning approaches. The models are evaluated on four GLUE tasks as presented in Table 7. We can see that autoencoder-based compression outperforms pruning baseline in most studies settings, especially when higher compression ratios are studied.

CR=3	Word embeddings			Output-dense			Keys		
	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)
$x^{2.0} + \text{BN}$	30.88	92.32	87.97 / 82.84	64.67	91.40	85.58 / 77.94	22.44	92.20	83.94 / 75.25
$x^{0.9} + \text{BN}$	24.97	<b>92.55</b>	88.32 / 83.09	62.93	91.51	86.13 / 78.92	<b>22.18</b>	92.20	84.04 / 75.49
$x^{0.5} + \text{BN}$	33.27	<b>92.55</b>	88.36 / 83.33	65.07	91.40	85.81 / 78.19	22.32	91.97	83.89 / 75.25
$x^{0.1} + \text{BN}$	58.08	91.86	86.47 / 81.13	68.19	91.40	<b>86.17 / 78.92</b>	22.29	91.97	83.99 / 75.25
Vanilla+BN	<b>24.68</b>	<b>92.55</b>	87.25 / 82.11	<b>62.60</b>	91.63	86.08 / 78.43	22.39	92.20	83.70 / 74.51
Vanilla	564.37	87.96	87.32 / 82.35	66.91	90.77	85.99 / 78.68	22.73	<b>92.55</b>	83.62 / 74.75
$\log_e(x) + C + 10$	68.46	92.43	86.91 / 81.62	68.61	<b>91.86</b>	86.04 / 78.68	22.27	92.32	84.44 / 75.98
$\log_e(x) + C$	236.24	92.09	86.75 / 81.13	68.17	91.74	<b>86.17 / 78.92</b>	22.26	91.97	84.36 / 75.74
No Fisher	118.41	91.97	<b>88.53 / 84.31</b>	70.69	91.74	<b>86.17 / 78.92</b>	22.35	92.43	<b>84.48 / 76.23</b>

CR=10	Word embeddings			Output-dense			Keys		
	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)
$x^{2.0} + \text{BN}$	147.29	86.93	84.85 / 77.94	1249.62	<b>89.45</b>	<b>84.64 / 75.98</b>	30.78	91.40	79.73 / 70.10
$x^{0.9} + \text{BN}$	<b>95.32</b>	88.30	85.22 / 78.92	1251.53	88.19	84.04 / 75.25	30.47	91.28	79.46 / 70.10
$x^{0.5} + \text{BN}$	250.59	<b>89.33</b>	<b>87.27 / 81.62</b>	1195.41	88.53	83.81 / 75.00	29.81	91.40	80.60 / 71.57
$x^{0.1} + \text{BN}$	567.88	87.16	84.73 / 77.21	1148.47	88.42	83.68 / 75.25	29.26	91.17	80.74 / 72.06
Vanilla+BN	136.84	88.99	85.86 / 79.90	1251.65	88.65	83.79 / 75.00	30.53	91.40	78.85 / 69.36
Vanilla	15994.88	79.70	81.22 / 68.38	1547.14	87.16	83.20 / 74.26	31.65	<b>91.63</b>	77.82 / 68.14
$\log_e(x) + C + 10$	806.71	88.30	86.06 / 80.15	<b>1145.25</b>	88.42	83.71 / 75.00	29.30	91.40	<b>81.34 / 72.79</b>
$\log_e(x) + C$	854.19	88.07	86.25 / 80.39	1148.46	88.99	84.09 / 75.25	29.25	91.17	80.14 / 72.06
No Fisher	712.98	88.07	85.87 / 80.88	1197.19	88.42	83.97 / 75.49	<b>29.01</b>	91.40	80.42 / 72.30

CR=25	Word embeddings			Output-dense			Keys		
	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)	Perplexity	SST-2 (Acc)	MRPC (F1/Acc)
$x^{2.0} + \text{BN}$	2779.96	80.73	83.99 / 75.98	1453.44	87.56	<b>84.54 / 75.98</b>	77.04	91.06	77.95 / 68.63
$x^{0.9} + \text{BN}$	<b>1983.19</b>	83.49	82.08 / 73.04	<b>1371.93</b>	<b>87.96</b>	83.31 / 74.26	71.90	90.60	79.66 / 70.34
$x^{0.5} + \text{BN}$	2728.41	<b>83.83</b>	<b>84.35 / 77.70</b>	1438.66	87.73	83.20 / 74.26	74.69	90.71	<b>79.73 / 70.59</b>
$x^{0.1} + \text{BN}$	5502.34	82.45	84.12 / 76.96	1447.31	85.89	82.58 / 73.28	74.34	91.17	78.68 / 69.85
Vanilla+BN	2167.99	80.50	82.60 / 73.77	1384.83	87.10	83.87 / 75.49	71.37	90.71	78.97 / 69.85
Vanilla	14880.66	77.29	80.19 / 70.10	1748.90	86.35	82.26 / 73.28	<b>66.53</b>	91.17	76.92 / 67.65
$\log_e(x) + C + 10$	5426.84	82.45	83.90 / 76.47	1449.46	86.93	82.89 / 73.53	73.28	90.83	78.40 / 69.61
$\log_e(x) + C$	4359.66	82.91	83.89 / 76.47	1446.55	86.58	81.79 / 72.06	76.81	91.06	77.78 / 68.63
No Fisher	4926.08	82.80	84.19 / 77.45	1462.41	85.61	82.02 / 72.06	69.24	<b>91.74</b>	78.78 / 69.36

Table 6: Incorporating Fisher Information coefficients in the autoencoder-based compression process for different compression ratios (CR) on the SST2 and the MRPC tasks. The first column demonstrates the transformation(s) applied to Fisher information before being passed to autoencoder loss function (more details in Appendix A.6). We also report the perplexity of the compressed model on the upstream task. Each score is the median of experiment results with three different seeds.

CR	Architecture	MRPC (F1/Acc)	SST-2 (Acc)	RTE (Acc)	QNLI (Acc)
1	BERT <sub>BASE</sub>	88.85/84.07	92.32	65.70	90.66
3	Pruning	<b>86.25/80.15</b>	<b>90.37</b>	56.47	<b>88.68</b>
	Our	85.25/77.77	90.25	<b>62.45</b>	<b>88.68</b>
10	Pruning	79.61/69.61	79.93	50.21	70.65
	Our	<b>81.45/71.08</b>	<b>83.94</b>	<b>57.76</b>	<b>81.48</b>
25	Pruning	80.86/69.36	75.11	51.97	60.57
	Our	<b>81.6/71.32</b>	<b>80.05</b>	<b>55.23</b>	<b>72.96</b>

Table 7: Comparison of BERT<sub>BASE</sub> model compression using autoencoder (AE) and pruning approaches (compressing token embedding matrix, all key matrices and all output-dense matrices). For AE we present the median score from experiments with 3 different seeds.



## A.8 Hyperparameter setting for the final experiment

This section presents the hyperparameter setting used for the Section 4.8 where multiple modules (token embedding, keys and output-dense) are compressed together. The autoencoder hyperparameters setting for this experiment can be found in Table 8.

## A.9 Compression time

In Table 9 we report compression times for keys, output-denses and token embeddings matrices from BERT<sub>BASE</sub> model using autoencoder and SVD compression approaches. For autoencoder we present times for both separated and concatenated matrices compression (as proposed in our paper) showing the advantage of using the latter approach.

We additionally performed experiments to compare the inference time of uncompressed BERT<sub>BASE</sub> with the extreme case of CR=25 using our approach (AE) and the SVD baseline. Compressed modules are token embeddings, key, and output-dense modules. We observe that inference times are very similar, with a slight increase in inference time when our model is used. In particular, the evaluation times (in seconds) for BERT/SVD/AE were respectively 0.474/0.477/0.493 (for the MRPC dataset) and 0.987/0.997/1.034 (for the SST2 dataset). Moreover, fine-tuning the BERT<sub>BASE</sub> model using compressed modules from autoencoder for the bigger datasets in GLUE, namely QNLI and QQP, takes at most 15% and 25% longer than the SVD baseline, respectively. It is also worth noting that when a linear autoencoder is incorporated, the inference time is the same as the SVD baseline.

Module	CR	Learning rate	Cosine:L <sub>2</sub> coefficients	Decoder	Enforce norm	Fisher transformation
Token embeddings	3	$5 \cdot 10^{-4}$	10:1	Non-linear (1 hidden layer)	Yes	$x^{0.5} + BN$
	10	$5 \cdot 10^{-4}$	10:1	Non-linear (2 hidden layers)	Yes	$x^{0.5} + BN$
	25	$10^{-3}$	10:1	Non-linear (1 hidden layer)	Yes	$x^{0.5} + BN$
Output-denses	3	$10^{-3}$	1:10	Linear	Yes	$x^{2.0} + BN$
	10	$10^{-4}$	1:10	Linear	No	$x^{2.0} + BN$
	25	$5 \cdot 10^{-4}$	0:1	Linear	No	$x^{2.0} + BN$
Keys	3	$5 \cdot 10^{-4}$	0:1	Linear	Yes	$\log_e(x) + C + 10$
	10	$5 \cdot 10^{-4}$	1:1	Linear	No	$\log_e(x) + C + 10$
	25	$10^{-3}$	0:1	Linear	No	$\log_e(x) + C + 10$

Table 8: The best hyperparameters for the BERT<sub>BASE</sub> model compression described in Section 4.8. In this experiment, token embedding matrix, all keys and all output-dense matrices are compressed using our proposed autoencoder-based framework.

Method	Mode	CR	Token embeddings	Keys	Output-denses
SVD	separated	3	~9.5min	(~5.5*12)min	(~6.0*12)min
SVD	separated	10	~8.0min	(~5.5*12)min	(~6.0*12)min
SVD	separated	25	~7.5min	(~5.5*12)min	(~6.0*12)min
AE	separated/concatenated	3	~7.7min	(~5.5*12)min/~6.1min	(~6.0*12)min/~6.5min
AE	separated/concatenated	10	~7.5min	(~5.5*12)min/~6.0min	(~6.0*12)min/~6.1min
AE	separated/concatenated	25	~7.5min	(~5.5*12)min/~5.9min	(~6.0*12)min/~6.1min

Table 9: Training time to retrieve compressed modules (key, output-dense, and token embeddings) of BERT<sub>BASE</sub> model using autoencoder (AE) and SVD approach. For the AE, we provide training times for the separated and concatenated modes to demonstrate another benefit of using the concatenated version, given its much better training time.