

THE-X: Privacy-Preserving Transformer Inference with Homomorphic Encryption

Tianyu Chen^{*♦♦}, Hangbo Bao[♥], Shaohan Huang[♥], Li Dong[♥], Binxing Jiao[♠],
Daxin Jiang[♠], Haoyi Zhou^{♠♦}, Jianxin Li^{♠♦✉}, Furu Wei[♥]

BDBC, Beihang University, China[♠]

SKLSDE, Beihang University, China[♦]

Microsoft Research[♥] NLP Group, Microsoft STCA[♠]

{tianyuc, zhouhy, lijx}@buaa.edu.cn

{t-habao, shaohan, lidong1, binxjia, djiang, fuwei}@microsoft.com

Abstract

As more and more pre-trained language models adopt on-cloud deployment, the privacy issues grow quickly, mainly for the exposure of plain-text user data (e.g., search history, medical record, bank account). Privacy-preserving inference of transformer models is on the demand of cloud service users. To protect privacy, it is an attractive choice to compute only with ciphertext in homomorphic encryption (HE). However, enabling pre-trained models inference on ciphertext data is difficult due to the complex computations in transformer blocks, which are not supported by current HE tools yet. In this work, we introduce *THE-X*, an approximation approach for transformers, which enables privacy-preserving inference of pre-trained models developed by popular frameworks. *THE-X* proposes a workflow to deal with complex computation in transformer networks, including all the non-polynomial functions like GELU, softmax, and Layer-Norm. Experiments reveal our proposed *THE-X* can enable transformer inference on encrypted data for different downstream tasks, all with negligible performance drop but enjoying the theory-guaranteed privacy-preserving advantage.

1 Introduction

Accompanying the revolution of pre-trained models in many NLP applications, such as sentiment analysis (Xu et al., 2019a), question answering (Yang et al., 2019b), information retrieval (Yang et al., 2019c), and text generation (Raffel et al., 2020), many related technologies have been deployed on the cloud to process user data from personal customers, small businesses, and large enterprises by industrial service providers. However, the convenience of the on-cloud pre-training technology also comes with a series of

^{*} Contribution during internship at MSRA. [✉]The corresponding author is Jianxin Li <lijx@buaa.edu.cn>.




Figure 1: An overview of our *THE-X*. The transformer-based model could inference on encrypted data with our *THE-X*, enabling theory-guaranteed privacy protection for users.

privacy challenges due to the sensitive nature of user data. For example, the input text or even text vector representations in user requests can leak private information, which may cause the specific user to be identified (Schwartz and Solove, 2011; Zhu and Han, 2020). This lack of privacy guarantees may impede privacy-conscious users from releasing their data to service providers. Thus, service providers may suffer from the deficiency of evolving models with user data. Besides, unintended data disclosure and other privacy breaches may result in litigation, fines, and reputation damages for service providers. These concerns spark our proposal of *THE-X*, to enable privacy-preserving inference of transformer.

Specifically, we identify two challenges for the privacy-preserving inference of pre-trained models. The first challenge is how to protect users' plain text data from access by third-party service providers. (e.g., the clinic record or shopping history). Prior work has applied Differential Privacy (DP) (Dwork et al., 2006) and its variants to address similar privatization issues - originally for statistical databases and more recently for DL (Abadi et al., 2016) and NLP (Qu et al., 2021; Basu et al., 2021b; Fernandes et al., 2019; Lyu et al., 2020; Basu et al., 2021a). However, this solution may suffer from eavesdropping attackers. A handful of research (Zhu and Han, 2020; Zhao et al., 2020)

demonstrated it possible to recover raw data from gradient leakage. Also, privacy protection could never be theory-guaranteed. The second challenge is the performance concern, recent works like *TextHide* (Huang et al., 2020) and *FedNLP* (Lin et al., 2021) leverages the federated learning (Yang et al., 2019a) to train model on encrypted data, at cost of considerable performance dropping. Focusing on the privacy of training data, they have not fully explored privacy-preserving inference.

To solve the concerns above, we depict one practice of privacy-preserving inference in Figure 1, where a fine-tuned language model could be converted into the cloud service mode with *THE-X*, and process users’ data with its eyes blind. During inference, the content of the user query is anonymous to the transformer model. The results of computation are also ciphertext, which only can be decrypted by the user’s private key.

In addition, we need a theory-guaranteed encryption solution like the homomorphic encryption (HE) (Gentry, 2009) to convince both service providers and users of the privacy security in production scenarios. The semantic security of HE is guaranteed by lattice-based cryptography, and the HE computation results on ciphertext could be decrypted to the same results in plaintext, preventing performance reduction cost. The basic idea of homomorphic encryption is to perform computations on encrypted data without first decrypting it, which could fully ensure privacy in cloud-serving scenarios. It allows user data to be encrypted and out-sourced to commercial cloud environments for processing.

However, due to the complex operations (e.g., GELU activation) in transformer-based models, the popular partially homomorphic encryption solution, which only supports addition or multiplication, can not easily be adapted into scenarios of pre-trained models. Based on HE transformer backend (Boemer et al., 2019b,a, 2020), we designed a series of approximation components to fulfill the whole inference pipeline of the mainstream transformer backbone. We evaluate *THE-X* for BERTiny on the GLUE benchmark (Wang et al., 2019) and the CONLL2003 task (Tjong Kim Sang and De Meulder, 2003). Our results show that *THE-X* can achieve the privacy-preserving inference with the averaged performance reduction of only 1.49%.

Our contributions include:

- We are the first work to explore the privacy-

preserving transformer inference with HE.

- We design a practical and effective approximation workflow for converting transformer-based models into a function that consists of fully HE operations.
- A thorough set of experiments confirms the negligible performance reduction with our proposed *THE-X* approximation.

2 Background

2.1 Security and Privacy Concern of Pre-trained Models

Pre-trained models like BERT (Devlin et al., 2019) and GPT-3 (Brown et al., 2020) rely heavily on the use of plain text data to get human-like performance. Despite the remarkable achievements of pre-trained models, these state-of-the-art models can not directly answer some sensitive use cases, including the medical record (Christoph et al., 2015), search history (Shen et al., 2007) and other personally identifiable information (PII).

To avoid the direct computation on plain-text data, recent works like *TextHide* (Huang et al., 2020) and *DP-finetuning* (Kerrigan et al., 2020) introduce the classical federated learning and differential privacy (DP) to protect the sensitive data. However, *TextHide* (Huang et al., 2020) can only be applied to sentence-level tasks. Due to the mix-up operation, *TextHide* fails to model token-level tasks like named entity recognition or semantic role labelling. *DP-finetuning* would greatly sacrifice the performance of fine-tuned model by 20% perplexity for a generation model like GPT-2.

2.2 Practical Homomorphic Encryption

The classic definition of homomorphic encryption is a form of encryption that permits users to perform computations on its encrypted data without first decrypting it. These computations results are retained in an encrypted form, which could be decrypted into identical output produced by the same computations on the unencrypted data. Let F be a function or the entire pre-trained model, E as an encryption function, D as a decryption function. Then for any allowed plain text input x , we have:

$$F(x) = D(g(E(x))), \quad (1)$$

where g is a constructed function to play the same role of function F , except on encrypted data. Figure 1 shows how a user performs inference using

a cloud-deployed pre-trained model which is not trusted. First, the pre-trained model receives a ciphertext encrypted by the user private key and performs inference function g on the ciphertext. Then, the server will send an encrypted result to the user, which can only be decrypted by the user key. At no point does the cloud service provider gain access to the plain text.

The Intel HE transformer for nGraph (Boemer et al., 2019b,a) is a Homomorphic Encryption backend to the deep learning models. Currently, it supports the CKKS (Cheon et al., 2017) encryption scheme, implemented by the Simple Encrypted Arithmetic Library (SEAL) (SEAL) from Microsoft Research. It is a research tool to demonstrate the feasibility of HE on deep learning.

2.3 Challenges of Transformer Inference with HE

Some HE schemes only support a single algebraic operation, such as addition or multiplication. These are known as "partially homomorphic" schemes (PHE). Other schemes, called "fully homomorphic" (FHE), support two such as addition and multiplication. Note that composing addition and multiplication suffices to construct polynomial functions, and hence polynomial approximations to non-polynomial functions such as GELU (Hendrycks and Gimpel, 2016) or LayerNorm (Xu et al., 2019b). Notably, this limitation prevents the exact computation of any comparison-based operations such as Max, Min, as well as common functions such as exponential or sigmoid. Finally, "leveled homomorphic" schemes (LHE) support addition and multiplication, only up to a fixed computational depth.

3 THE-X: Formal Description

There are two core ideas in *THE-X*. The first one is to incorporate the user device into the HE inference, and the second is using "simplified computation" to approximate the non-polynomial functions.

In the following, we will describe how to enable homomorphic encryption of transformer-based models with *THE-X*.

3.1 Approximation Workflow

First, we present the approximation workflow of *THE-X*, which consists of two stages: Standard Finetuning and LN Distill as depicted in Figure 2. Given a pre-trained model \mathcal{M} and corresponding

Algorithm 1: Approximation Workflow

Data: labeled task data \mathcal{D} .
Input: pre-trained Transformer model \mathcal{M} , softmax estimation model \mathcal{S} .

- 1 $\widehat{\mathcal{M}} \leftarrow \mathcal{M} \odot (\mathcal{S}, ReLU)$.
// replace GELU and Softmax
- 2 **while not done do**
- 3 sample batches (x_i, y_i) from \mathcal{D} ,
- 4 let (x_i, y_i) optimize $\widehat{\mathcal{M}}$ with \mathcal{S} frozen.
- end**
- 5 $\widetilde{\mathcal{M}} \leftarrow \widehat{\mathcal{M}} \oplus \widetilde{\mathcal{N}}$.
// add the layernorm approximation
- 6 **while not done do**
- 7 sample batches (x_i, y_i) from \mathcal{D} ,
- 8 freeze the parameters of $\widetilde{\mathcal{M}}$ except $\widetilde{\mathcal{N}}$.
- 9 compute k -th layernorm output O_k, \widetilde{O}_k .
- 10 compute loss $\ell_k = \text{MSELoss}(O_k, \widetilde{O}_k)$.
- 11 update $\widetilde{\mathcal{N}}$ with loss $\mathcal{L} = \sum_k \ell_k$.
- end**
- 12 $\bar{\mathcal{M}} \leftarrow \widetilde{\mathcal{M}} \ominus \mathcal{N}$.
// discard the origin layernorm

return $\bar{\mathcal{M}}$.

downstream data, we aim to produce a fully HE supported $\bar{\mathcal{M}}$ which is fine-tuned and ready for deployment.

The two-stage optimization of algorithm 1 aims to find the best approximation checkpoint. For computation efficiency, pre-trained models can also be fine-tuned together with the layernorm approximation, and it needs only a single optimization loop. We will discuss the schedule of the different approximation workflow in Sec 4.6. There are three major non-polynomial functions in the transformer block, where we will study in detail.

3.1.1 Gaussian Error Linear Units (GLEU)

With a computation of Gaussian error, Gaussian Error Linear Units (GLEUs) (Hendrycks and Gimpel, 2016) is not suitable to serve as an active function in HE state. The Gaussian kernel includes unsupported functions like exponential. While in the implementation of the transformer, GELU is defined as a fast approximated version, where the \tanh function is still non-polynomial, unsupported by HE.

$$G(x) = 0.5x(1 + \tanh[\sqrt{2/\pi}(x + 0.044715x^3)]). \quad (2)$$




Figure 2: The Approximation Workflow of *THE-X*. To replace the non-polynomial operations, we split the fine-tuning stage into several subphases. Given a pre-trained checkpoint, we drop the pooler of the pre-trained model and replace softmax and GeLU. Afterward, we follow the standard fine-tuning for classification or regression tasks. We add LayerNorm approximation into the fine-tuned model and distill knowledge from original LN layers. After dropping the original LN, we convert the model into fully HE-supported ops with the HE transformer.




Figure 3: The activation results of GELU compared with ReLU. With an input around zero, the activation results are very close. With a larger or smaller input value, the activation results tend to converge.

We illustrate the numerical comparison between GELU and RELU in Figure 3, where the outputs of GELU are very close to RELU. Hence, we propose to replace the GELU layer in the model with a ReLU activation function. Despite the *Max* function in ReLU, other computations are well supported by HE. To enable the computation of *Max*, we implement the first key idea, incorporating the user device into the inference. The server will convey ciphertext input to the user for local *Max* computation. Once received the connection, a user device decrypts the ciphertext input and calls the local *Max* function to get the results and return re-encrypted results to the server. Despite the communication cost, no plaintext is leaked during the TLS connection and semantic security is guaranteed.

3.1.2 Softmax

The second non-polynomial function is softmax, which includes the exponential and division computation.

$$\text{Softmax}(x_i) = \frac{\exp(x_i)}{\sum_j \exp(x_j)}. \quad (3)$$

The first thought to approximate softmax is to find alternatives of softmax operation in transformer, which include Taylor series approximation (Vincent et al., 2015), softmax-free linear attention (Lu et al., 2021). However, both of them have some limitations. The Taylor series approximation can only approximate the exponential operation. Softmax-free linear attention utilizes newton-inverse to approximate division, but the approximation error is unbounded in full-scale attention settings.

For these considerations, we have no choice but to design an estimation network with addition and multiplication.

$$S(x_i) = x_i * T\left(\sum_j \text{ReLU}\left(\frac{x_j}{2} + 1\right)^3\right). \quad (4)$$

Equation 4 is the formal description of our softmax estimation network. Same as the approximation of GELU, ReLU operation here is realized by communication with the client. Instead of a division operation, we approximate reciprocal operation with a three-layer linear neural network denoted as T .

To get a better estimation of softmax, we randomly generate input tensors whose values are between $[-3, 3]$ and use their softmax scores as MSE targets. Then we optimize the T for 100k steps with a learning rate of 1e-3 until the MSE loss drop down to 1e-6.

An under-explored problem here is the *Infinite value of Masked Attention*, where the input of softmax is always the masked attention scores. To prevent the attention of masked tokens, the origin transformer model fills the masked attention scores with negative infinity before softmax. When fed with an infinite value, the softmax estimation model

may face numerical disaster. We will discuss this phenomenon and the corresponding solution in Sec 4.5.

3.1.3 LayerNorm

Recall that the layer normalization (Ba et al., 2016) in transformer is implemented over a mini-batch of inputs, which could be formulated as:

$$y = \frac{x - E[x]}{\sqrt{Var[x] + \epsilon}} * \gamma + \beta. \quad (5)$$

The mean and standard deviation are calculated over division operations where the approximation is needed. γ and β are learnable affine transform parameters. To avoid the introduction of new parameters, we keep the learnable parameters while leaving the mean and standard deviation achieved by regression.

$$\hat{y} = x \circ \gamma + \beta. \quad (6)$$

The new parameter γ predicts the value of standard deviation by regression from origin \hat{y} . We find the simple linear replacement is enough for values with a small scale of bias. Here $\gamma, \beta \in \mathbb{R}$ and \circ denotes the Hadamard product.

The layer normalization will be applied in each multi-head attention block and after the output dense layer. So the approximation error tends to accumulate when the transformer stacks with too many layers.

We treat the layernorm approximation as an individual stage in Figure 1 as *LN-Distill* to learn from origin LN layers. A challenge here is the *Attention Overflow*, where the input attention score before normalization may have an unbounded scale, leading to numerical problems. We will discuss the detail of *Attention Overflow* in Sec 4.5.

3.1.4 Other Practical Replacement

After the approximation workflow, a fine-tuned model consists of only addition and multiplication operations, which is fully compatible with homomorphic encryption. We power the model by HE transformer backend. Since the HE transformer backend could only work for TensorFlow checkpoint, any pre-trained transformers inherited from PyTorch building version need to be converted into TensorFlow format first. There are some other details worth mentioning here.

- For the $\text{softmax}(\frac{QK^T}{\sqrt{d_k}})V$ operation in attention score computation, we absorb the value

of $\frac{1}{\sqrt{d_k}}$ into the weights of query projection layer.

- We use a fully kernel convolution layer instead of linear projection due to the lack of supported dense operation.
- All matrix multiplication will be converted into the element-wise style.
- We drop the pooler layer for the unsupported operation of tanh.

3.2 Privacy-preserving Inference

In this section, we describe the behavior of HE models during privacy-preserving inference. Note that inference is completed by the joint effort of the server and the user device.

Algorithm 2: Inference with HE

Input: user plain text query \mathcal{P}_q , private key \mathcal{K} generated under server protocol, encrypted server model \mathcal{M} .

- 1 client computes embeddings: $\mathcal{E}_q \leftarrow \mathcal{P}_q$.
 - 2 client encrypts query embeddings:
 $\mathcal{C}_q \leftarrow \text{Encrypt}(\mathcal{E}_q, \mathcal{K})$.
 - 3 server forwards the model: $\mathcal{C}_i = \mathcal{M}(\mathcal{C}_q)$.
 - 4 client handles activation: $\mathcal{C}_a = \text{ReLU}(\mathcal{C}_i)$.
 - 5 server continues forwarding: $\mathcal{C}_o = \mathcal{M}(\mathcal{C}_a)$.
 - 6 client decrypts results:
 $\mathcal{P}_o = \text{Decrypt}(\mathcal{C}_o, \mathcal{K})$.
-

In Algorithm 2, notably absent is the support of ReLU operations, where the server exchanges the activation results with the client. However, all the communication between client and server is in ciphertext, ensuring the privacy of user queries and may prevent eavesdropping attackers from recovering private text data.

4 Experiments

In this section, we design both sequence-level and token-level tasks to evaluate the approximation performance of our *THE-X* solution. We also discuss several identified factors which greatly affect approximation workflow.

4.1 Evaluation Tasks

GLUE (Wang et al., 2019), the General Language Understanding Evaluation benchmark, is a collection of tools for evaluating the performance of models across a diverse set of existing NLU tasks.

Table 1: Performance on the GLUE¹ tasks for both baseline (standard finetuning) and *THE-X* with BERT-tiny, measured on the development sets. We report the best results by hyper-parameter search. $|\mathcal{D}|$ denotes the number of training examples. *THE-X* only suffers average utility performance loss: $< 1.5\%$ in most tasks. ‘P/S corr.’ is Pearson/Spearman correlation and ‘m/mm’ denotes the accuracy scores on matched/mismatched set.

Tasks	$ \mathcal{D} $	Type	Metrics	Baseline	ReLU	ReLU-S	ReLU-S-L	HE	Perf ↓
SST-2	67k	Sentiment	Acc.	82.45	82.40	82.34	82.11	82.11	0.34
MRPC	3.7k	Paraphrase	F1/Acc.	81.57/70.10	81.69/70.34	80.81/69.85	79.93/68.87	79.94/68.87	1.63/1.23
STS-B	7k	Similarity	P/S corr.	72.83/73.66	72.89/73.03	74.19/74.27	68.38/70.96	68.39/70.97	4.44/2.69
QQP	364k	Paraphrase	F1	80.28/84.03	79.55/82.89	79.38/83.36	78.28/83.75	78.33/83.63	1.95/0.40
MNLI	393k	NLI	m/mm	69.75/70.75	69.51/70.60	68.61/69.13	68.59/69.41	68.47/69.08	1.28/1.67
QNLI	108k	NLI	ACC.	78.38	78.35	78.33	78.33	78.20	0.18
RTE	2.5k	NLI	ACC.	58.56	58.32	58.27	58.12	58.12	0.44
Average Perf ↓				0.00	0.25	0.34	1.42	1.48	1.48

Table 2: Performance on the CONLL2003 task for both baseline and *THE-X* with BERT-tiny, measured on the development sets. We find that the replacement with ReLU has a slight effect on performance and even gets a better F1 score by 0.12 than original GELU activation.

Metrics	Precision	Recall	F1	Perf ↓
Raw	82.34	84.85	83.57	0
ReLU	82.29	85.13	83.69	-0.12
ReLU-S	82.08	84.73	83.38	0.19
ReLU-S-L	79.65	83.79	81.67	1.90
HE	79.65	83.79	81.67	1.90

We choose a subset of GLUE¹ tasks, which include: MRPC (Dolan and Brockett, 2005), SST-2 (Socher et al., 2013), QQP², STS-B (Cer et al., 2017), MNLI (Williams et al., 2018), QNLI (Rajpurkar et al., 2016), and RTE (Dagan et al., 2005; Haim et al., 2006; Giampiccolo et al., 2007; Bentivogli et al., 2009).

Following previous work (Devlin et al., 2019; Turc et al., 2019), we exclude the WNLI task from the GLUE benchmark. We also use the famous CoNLL-2003 (Tjong Kim Sang and De Meulder, 2003) named entity recognition task as our additional token-level evaluation. In conclusion, we include the most varieties of NLU tasks, covering both sequence-level and sentence-level tasks, in both regression and classification format.

4.2 Experiment Settings

For computation efficiency and energy-saving consideration, we use the released BERT-tiny (Turc

et al., 2019) as our demo model, which is a standard transformer-based language model with only 2 layers and a hidden size of 128. We provide four settings to evaluate different parts of our approximation components.

- **Baseline.** In this setting, we make no replacement or approximation. We use the raw pre-trained checkpoint to fine-tune on downstream tasks.
- **ReLU.** We fine-tune the pre-trained model with all GELU activation replaced with ReLU.
- **ReLU-S.** In addition to ReLU, we fine-tune the model with the softmax operation replaced by the softmax estimation model.
- **ReLU-S-L.** We implement full approximation including a layer normalization replacement.
- **HE.** We convert the fine-tuned checkpoint with HE-transformer and power the inference with SEAL backend.

Implementation. To reduce the variance of results under different settings, we choose hyper-parameters from a fixed set during approximation fine-tuning and HE inference runtime.

- For fine-tuning the approximation components, we choose a batch size from {4, 8, 16, 32, 128} and a learning rate from 1e-4, 3e-4, 3e-5, 5e-5 as mentioned in the initial bert code (Turc et al., 2019). We use an Adam optimizer with weight decay chosen from {0.05, 0.1, 0.2, 0.4, 0.5}
- For HE evaluation, we use the HE-transformer backend, where two parameters are recommended searching by Intel, the poly modules

¹CoLA task is not reported because of the limited capacity of BERT-tiny.

²<https://www.quora.com/profile/Ricky-Riche-2/First-Quora-Dataset-Release-Question-Pairs>

and coeff-modules. We choose the poly modules degree from {1024, 2048, 4096, 8192, 16384} and choose the coeff-modules from {20, 30, 60}.

4.3 Approximation Results

Table 1 shows the results of the baseline and *THE-X* on the GLUE benchmark. The averaged performance reduction of *THE-X* is 1.48% when compared to the baseline model. We observe the most performance reduction comes from the approximation of layernorm, which incurs a reduction of 1.08%. The softmax estimation model contributes the least performance drop among the approximation components, for only 0.09% on average, indicating the softmax function could be well imitated by neural networks. We also find the average performance reduction of HE is quite negligible, where the slight drop may be due to the sequence truncation.

The results of *THE-X* on token-level NER task are reported in Table 2. The replacement of GELU with ReLU even improves the performance of the F1 score. We assume the slight improvement may come from unexpected bias. However, the layernorm approximation incurs the most performance reduction. We assume token-level tasks need a more detailed pattern in attention score. After all, *THE-X* still works well in the token-level task with a merely F1 reduction of 1.9%.

Across different types of tasks, we find our *THE-X* yields the best performance on the classification tasks, including paraphrase, sentiment and NLI. Among the classification tasks, the performance of QNLI drops the least, for only 0.18%. We also find the performance drops most on the regression tasks, such as the similarity task STS-B, for 4.44% pearson correlation and 2.69% spearman correlation. We assume the regression task needs a higher numerical precision than the classification task.

4.4 Negative Infinity

Recall in Equation 4, we replace softmax with a neural estimation model. To prevent the attention of masked tokens, the origin transformer model fills the masked attention scores with negative infinity before softmax, where the numerical disaster occurs in our approximation method. In Figure 4, to solve this problem, we give an empirical study of how "negative" the masked attention scores should be. Despite the indistinguishable F1 score change of raw model fine-tune with different




Figure 4: Performance on CONLL2003 task with different mask values. We find the “Negative Infinity” value of the 0 mask greatly reduces approximation performance. In *THE-X* using a mask value in [-3, -5] might be a default choice.

attention mask values, the approximation method is extremely sensitive to the numerical changes. We assume the softmax estimation model fails to deal with large input values and leads to a credible performance drop. However, when the value of the attention mask becomes too small, it serves as a bias to attention scores, which also leads to a certain performance drop. We recommend using a moderate mask value between -2 and -5.

4.5 Attention Overflow

Another challenge of *THE-X* is the attention score input of layer normalization. In most cases, the scale of multi-head attention output is very dense around $[-1, 1]$. However, before normalization, we also observe the attention scores are sparsely sparse, with some extreme values reaching $1e4$, which is difficult for our LN-distill stage. To prevent the overflow attention scores, we use the **weight decay** of Adam optimizer as regularization.

In Figure 5, we present the attention overflow phenomenon across different tasks. Without any regularization, our approximation method yields uncontrolled attention scores, leading to poor performance. As the weight decay increases, the attention scores tend to converge and benefit better approximation results. We also observe that the larger weight decay may harm the performance on NLI tasks, where the regularization could be seen as trade-off between better approximation results and higher performance upper bound. For the NER task, larger weight decay may even benefit the performance and also boost our approximation method.

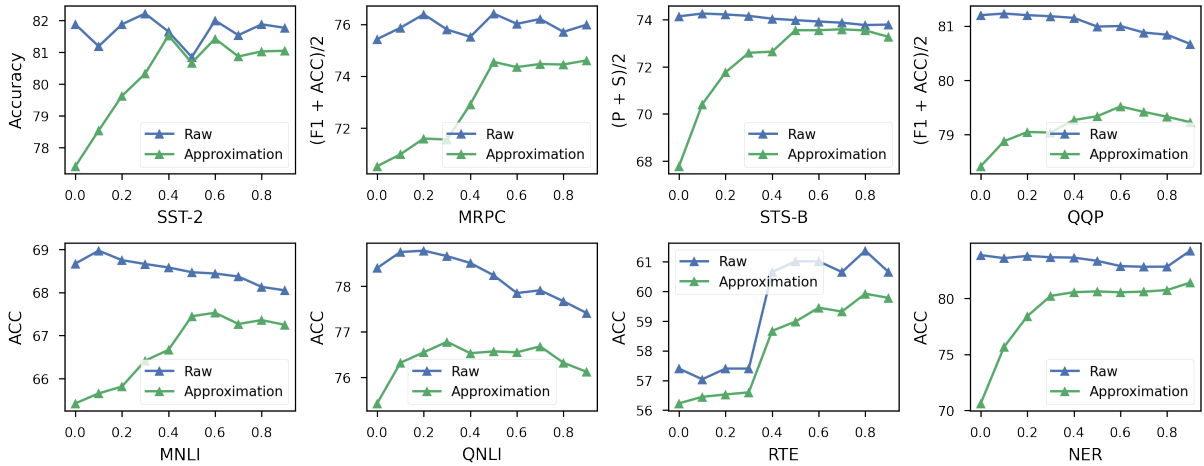


Figure 5: Performance on all tasks with different weight decay values, measured on the development sets. Metrics are marked on the y-axis and weight decay values are marked on the x-axis.

4.6 Schedule of Approximation workflow

There are still doubts about how to organize the several optimization steps for the best approximation performance. We investigate four schedule plans:

- **Two Stages.** Where we freeze the softmax estimation model during standard fine-tuning. We select the best checkpoint to implement the second stage - distill the layer normalization network.
- **Joint FT S.** We optimize the softmax estimation model during standard fine-tuning and apply the LN-distillation after.
- **Joint FT LN.** We apply one-pass optimization with the softmax estimation model frozen but update the other parameters including layer normalization network. No further LN-distill will be implemented.
- **Joint FT S + LN.** A total one-pass optimization with all approximation parameters updated with the model together.

As illustrated in Figure 6, we observe that fine-tuning the different approximation components individually (aka. "Two stages") may be a good default to keep the best performance of approximation. For the regression task STS-B, jointly fine-tuning the softmax estimation model and approximated layernorm even fails to fulfill the approximation pipeline, pulling the performance down to 0.4%. We assume fine-tuning different components may fall into a *bi-level* optimization problem and




Figure 6: Performance on all tasks with different organizations of approximation workflow. Jointly fine-tuning the softmax estimation model or approximated layernorm leads to a performance drop across all tasks.

it is hard to achieve satisfying results. In conclusion, the softmax estimation model and the approximated layernorm are both critical components to the performance of *THE-X*, deserving individual optimization.

5 Conclusions

We present *THE-X*, a practical approach to enable pre-trained transformer models to infer under homomorphic encryption. It requires several approximation components to replace the original operations in the transformer model. It imposes a slight burden in terms of performance cost but enjoys the full advantage of homomorphic encryption - the theory-guaranteed user privacy.

We see this as a first step in combing homomorphic encryption to address emerging privacy issues in pre-trained models. We hope our work motivates further research, including better approximation solutions on different NLP applications.

6 Acknowledgments

This paper is supported in part by the NSFC through grant No.U20B2053. We also thanks the support from Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing.

References

- Martín Abadi, Andy Chu, Ian J. Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. [Deep learning with differential privacy](#). In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 308–318. ACM.
- Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#). *CoRR*, abs/1607.06450.
- Priyam Basu, Tiasa Singha Roy, Rakshit Naidu, and Zumrut Muftuoglu. 2021a. [Privacy enabled financial text classification using differential privacy and federated learning](#). In *Proceedings of the Third Workshop on Economics and Natural Language Processing*, pages 50–55, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Priyam Basu, Tiasa Singha Roy, Rakshit Naidu, Züm-rüt Müftüoglu, Sahib Singh, and Fatemehsadat Mireshghallah. 2021b. [Benchmarking differential privacy and federated learning for BERT models](#). *CoRR*, abs/2106.13973.
- Luisa Bentivogli, Bernardo Magnini, Ido Dagan, Hoa Trang Dang, and Danilo Giampiccolo. 2009. [The fifth PASCAL recognizing textual entailment challenge](#). In *Proceedings of the Second Text Analysis Conference, TAC 2009, Gaithersburg, Maryland, USA, November 16-17, 2009*. NIST.
- Fabian Boemer, Rosario Cammarota, Daniel Demmler, Thomas Schneider, and Hossein Yalame. 2020. [MP2ML: A mixed-protocol machine learning framework for private inference](#). In *PPMLP'20: Proceedings of the 2020 Workshop on Privacy-Preserving Machine Learning in Practice, Virtual Event, USA, November, 2020*, pages 43–45. ACM.
- Fabian Boemer, Anamaria Costache, Rosario Cammarota, and Casimir Wierzynski. 2019a. [ngraph-he2: A high-throughput framework for neural network inference on encrypted data](#). In *Proceedings of the 7th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography, WAHC@CCS 2019, London, UK, November 11-15, 2019*, pages 45–56. ACM.
- Fabian Boemer, Yixing Lao, Rosario Cammarota, and Casimir Wierzynski. 2019b. [ngraph-he: a graph compiler for deep learning on homomorphically encrypted data](#). In *Proceedings of the 16th ACM International Conference on Computing Frontiers, CF 2019, Alghero, Italy, April 30 - May 2, 2019*, pages 3–13. ACM.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Jung Hee Cheon, Andrey Kim, Miran Kim, and Yong Soo Song. 2017. [Homomorphic encryption for arithmetic of approximate numbers](#). In *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I*, volume 10624 of *Lecture Notes in Computer Science*, pages 409–437. Springer.
- J Christoph, L Griebel, I Leb, I Engel, F Köpcke, D Toddenroth, H-U Prokosch, J Laufer, K Marquardt, and M Sedlmayr. 2015. [Secure secondary use of clinical data with cloud-based nlp services](#). *Methods of information in medicine*, 54(03):276–282.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. [The PASCAL recognising textual entailment challenge](#). In *Machine Learning Challenges, Evaluating Predictive Uncertainty, Visual Object Classification and Recognizing Textual Entailment, First PASCAL Machine Learning Challenges Workshop, MLCW 2005, Southampton, UK, April 11-13, 2005, Revised Selected Papers*, volume 3944 of *Lecture Notes in Computer Science*, pages 177–190. Springer.
- 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.

- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing, IWP@IJCNLP 2005, Jeju Island, Korea, October 2005, 2005*. Asian Federation of Natural Language Processing.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. 2006. [Calibrating noise to sensitivity in private data analysis](#). In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer.
- Natasha Fernandes, Mark Dras, and Annabelle McIver. 2019. [Generalised differential privacy for text document processing](#). In *Principles of Security and Trust - 8th International Conference, POST 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings*, volume 11426 of *Lecture Notes in Computer Science*, pages 123–148. Springer.
- Craig Gentry. 2009. [Fully homomorphic encryption using ideal lattices](#). In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 169–178. ACM.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. [The third PASCAL recognizing textual entailment challenge](#). In *Proceedings of the ACL-PASCAL@ACL 2007 Workshop on Textual Entailment and Paraphrasing, Prague, Czech Republic, June 28-29, 2007*, pages 1–9. Association for Computational Linguistics.
- R Bar Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, Bernardo Magnini, and Idan Szpektor. 2006. [The second pascal recognising textual entailment challenge](#). In *Proceedings of the Second PASCAL Challenges Workshop on Recognising Textual Entailment*.
- Dan Hendrycks and Kevin Gimpel. 2016. [Bridging nonlinearities and stochastic regularizers with gaussian error linear units](#). *CoRR*, abs/1606.08415.
- Yangsibo Huang, Zhao Song, Danqi Chen, Kai Li, and Sanjeev Arora. 2020. [TextHide: Tackling data privacy in language understanding tasks](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1368–1382, Online. Association for Computational Linguistics.
- Gavin Kerrigan, Dylan Slack, and Jens Tuyls. 2020. [Differentially private language models benefit from public pre-training](#). In *Proceedings of the Second Workshop on Privacy in NLP*, pages 39–45, Online. Association for Computational Linguistics.
- Bill Yuchen Lin, Chaoyang He, Zihang Zeng, Hulin Wang, Yufen Huang, Mahdi Soltanolkotabi, Xiang Ren, and Salman Avestimehr. 2021. [Fednlp: A research platform for federated learning in natural language processing](#). *CoRR*, abs/2104.08815.
- Jiachen Lu, Jinghan Yao, Junge Zhang, Xi Tian Zhu, Hang Xu, Weiguo Gao, Chunjing Xu, Tao Xiang, and Li Zhang. 2021. [SOFT: softmax-free transformer with linear complexity](#). *CoRR*, abs/2110.11945.
- Lingjuan Lyu, Xuanli He, and Yitong Li. 2020. [Differentially private representation for NLP: formal guarantee and an empirical study on privacy and fairness](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 2355–2365. Association for Computational Linguistics.
- Chen Qu, Weize Kong, Liu Yang, Mingyang Zhang, Michael Bendersky, and Marc Najork. 2021. [Natural language understanding with privacy-preserving bert](#). *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100, 000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2383–2392. The Association for Computational Linguistics.
- Paul M Schwartz and Daniel J Solove. 2011. [The pii problem: Privacy and a new concept of personally identifiable information](#). *NYUL rev.*, 86:1814.
- SEAL. 2020. Microsoft SEAL (release 3.6). <https://github.com/Microsoft/SEAL>. Microsoft Research, Redmond, WA.
- Xuehua Shen, Bin Tan, and ChengXiang Zhai. 2007. [Privacy protection in personalized search](#). *SIGIR Forum*, 41(1):4–17.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1631–1642. ACL.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In

- Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Well-read students learn better: On the importance of pre-training compact models. *arXiv preprint arXiv:1908.08962*.
- Pascal Vincent, Alexandre de Brébisson, and Xavier Bouthillier. 2015. Efficient exact gradient update for training deep networks with very large sparse targets. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1108–1116.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu. 2019a. BERT post-training for review reading comprehension and aspect-based sentiment analysis. 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 2324–2335. Association for Computational Linguistics.
- Jingjing Xu, Xu Sun, Zhiyuan Zhang, Guangxiang Zhao, and Junyang Lin. 2019b. Understanding and improving layer normalization. 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 4383–4393.
- Qiang Yang, Yang Liu, Yong Cheng, Yan Kang, Tianjian Chen, and Han Yu. 2019a. *Federated Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Wei Yang, Yuqing Xie, Aileen Lin, Xingyu Li, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019b. End-to-end open-domain question answering with bertserini. 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 72–77. Association for Computational Linguistics.
- Wei Yang, Haotian Zhang, and Jimmy Lin. 2019c. Simple applications of BERT for ad hoc document retrieval. *CoRR*, abs/1903.10972.
- Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. 2020. idlg: Improved deep leakage from gradients. *CoRR*, abs/2001.02610.
- Ligeng Zhu and Song Han. 2020. Deep leakage from gradients. In Qiang Yang, Lixin Fan, and Han Yu, editors, *Federated Learning - Privacy and Incentive*, volume 12500 of *Lecture Notes in Computer Science*, pages 17–31. Springer.