



<b>Title</b>	A Sentence-level Hierarchical BERT Model for Document Classification with Limited Labelled Data
<b>Authors(s)</b>	Lu, Jinghui, Henchion, Maeve, Bacher, Ivan, MacNamee, Brian
<b>Publication date</b>	2021-10-09
<b>Publication information</b>	Lu, Jinghui, Maeve Henchion, Ivan Bacher, and Brian MacNamee. "A Sentence-Level Hierarchical BERT Model for Document Classification with Limited Labelled Data." Springer, October 9, 2021. <a href="https://doi.org/10.1007/978-3-030-88942-5_18">https://doi.org/10.1007/978-3-030-88942-5_18</a> .
<b>Series</b>	Lecture Notes in Computer Science, 12986
<b>Publisher</b>	Springer
<b>Item record/more information</b>	<a href="http://hdl.handle.net/10197/25404">http://hdl.handle.net/10197/25404</a>
<b>Publisher's statement</b>	The final publication is available at <a href="http://www.springerlink.com">www.springerlink.com</a> .
<b>Publisher's version (DOI)</b>	<a href="https://doi.org/10.1007/978-3-030-88942-5_18">10.1007/978-3-030-88942-5_18</a>

Downloaded 2026-05-01 23:38:19

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd\_oa)



© Some rights reserved. For more information

# A Sentence-level Hierarchical BERT Model for Document Classification with Limited Labelled Data\*

Jinghui Lu<sup>1</sup>[0000-0001-7149-6961], Maeve Henchion<sup>2</sup>[0000-0001-7842-1516], Ivan Bacher<sup>3</sup>, and Brian Mac Namee<sup>1</sup>[0000-0003-2518-0274]

<sup>1</sup> University College Dublin, Dublin, Ireland [Jinghui.Lu@ucdconnect.ie](mailto:Jinghui.Lu@ucdconnect.ie),  
[Brian.MacNamee@ucd.ie](mailto:Brian.MacNamee@ucd.ie)

<sup>2</sup> Teagasc Agriculture and Food Development Authority, Dublin, Ireland  
[Maeve.Henchion@teagasc.ie](mailto:Maeve.Henchion@teagasc.ie)

<sup>3</sup> ADAPT Research Centre, Dublin, Ireland  
[Ivan.Bacher@gmail.com](mailto:Ivan.Bacher@gmail.com)

**Abstract.** Training deep learning models with limited labelled data is an attractive scenario for many NLP tasks, including document classification. While with the recent emergence of BERT, deep learning language models can achieve reasonably good performance in document classification with few labelled instances, there is a lack of evidence in the utility of applying BERT-like models on long document classification. This work introduces a long-text-specific model — the Hierarchical BERT Model (HBM) — that learns sentence-level features of the text and works well in scenarios with limited labelled data. Various evaluation experiments have demonstrated that HBM can achieve higher performance in document classification than the previous state-of-the-art methods with only 50 to 200 labelled instances, especially when documents are long. Also, as an extra benefit of HBM, the salient sentences identified by learned HBM are useful as explanations for labelling documents based on a user study.

**Keywords:** document classification · low-shot learning · BERT.

## 1 Introduction

In many real-world applications, it is not often that people working in NLP can access sufficiently large labelled data to train deep learning models. With the recent advance in Bidirectional Encoder Representations from Transformers (BERT) and its variants [4,9], language models can extract extensive general language features from large corpora like never before and transfer learned knowledge to the target domain where labelled data is limited. Previous work has demonstrated that the pre-training and fine-tuning paradigm based on BERT outperforms the traditional NLP approaches in document classification where people have limited number of labelled examples [15]. However, we argue that

---

\* Supported by organization Teagasc and Insight Centre for Data Analytics.

when texts are long, even BERT can not generate proper results with a small training data due to the text length limitation of BERT [5].

Inspired by *Hierarchical Attention Network* (HAN) [3] that extracts sentence-level information for document classification, in this work, we attempt to further improve the performance of BERT-based models in document classification under low labelled data context by considering the sentence structure information of texts. Specifically, following the intrinsic attention mechanism of the BERT model, which is good at capturing the relationship between words, we propose a *Hierarchical BERT Model* (HBM) which is focused on the sentence-level structure of a document. In the same way that BERT captures the connections between words, the HBM is designed to capture the connections between sentences. We believe that this approach can overcome the loss of information regarding document structure that occurs when simple averaging of word vectors (such as word2vec, BERT-based language models) is used to form the document representations. We expect that the proposed model can increase classification accuracy especially for long documents. Furthermore, due to the sentence attention mechanism in HBM, not all sentences are equally weighted when constructing the document representation. In HBM, information from sentences that receive higher attention scores is more likely to be aggregated to form the representation of the document. We propose that sentences contributing more to the representation of the document are important sentences, and that these sentences can be used as explanatory sentences to make it easier for the human to make a decision about a document when performing the labelling task. Hence, in this paper, we raise the following research questions:

**RQ1:** *Is the proposed sentence-level BERT-based model, the Hierarchical BERT Model (HBM), able to achieve high classification performance in scenarios with small amounts of labelled data for training, especially when documents are long?*

**RQ2:** *Can highlighting important sentences identified by the HBM help the annotator label documents faster?*

The remainder of the paper is organised as follows: Section 2 summarises the related work; Section 3 describes the proposed HBM method in detail; Section 4 describes an evaluation experiment comparing the proposed HBM method with state-of-the-art methods under the low labelled data scenario; Section 5 introduces the inference of important sentences by HBM; Section 6 demonstrates the effectiveness of the HBM method in selecting important sentences to show to human by a user study; and Section 7 draws conclusions.

## 2 Related Work

This work seeks to build on the BERT language models to further explore the effectiveness of taking sentence-level information into account for document classification with limited labelled data. Though BERT and its variants have achieved impressive results in many NLP tasks, including document classification

in larger datasets [4,9], few studies are focused on the performance of BERT in low labelled data scenario over other existing NLP paradigms. As suggested in [6], supervised deep learning models will generally achieve acceptable performance at about 5,000 examples per class, and therefore, deep learning models may not be suitable for scenarios with low labelled examples. Interestingly, Usherwood and Smit [15] demonstrated that when one has only 100-1,000 labelled examples per class, BERT is the best performing approach for document classification as compared to other traditional machine learning approaches, due to the utility of pre-training and fine-tuning paradigm. However, their work is an empirical comparison of BERT and other NLP approaches without any modification in BERT. Also, BERT is incapable of processing long texts [5], the datasets used in [15] are Amazon reviews and Twitter that consist of short texts, lacking evidence of the utility of BERT-based models for long text classification with limited labelled data. Hence, our work proposes using a sentence-level language model to alleviate the long-texts problem suffered by BERT.

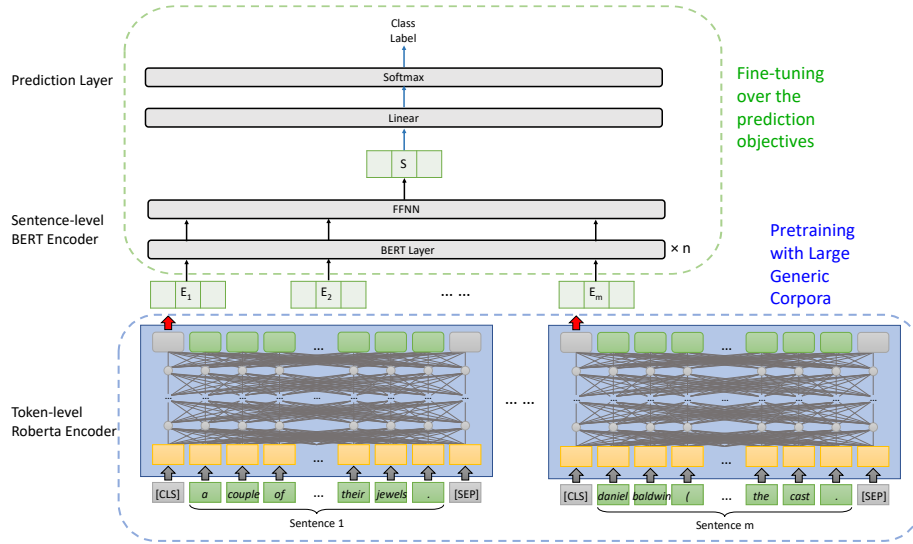
Yang et al. [18] were the first to propose the sentence-level model *Hierarchical Attention Network* (HAN) for document classification, and they adopted the Bi-GRU with attention as their architecture. There are also some concurrent works investigating the effectiveness of sentence-level model based on BERT. For example, Xu et al. [16] used a sentence-level BERT to a query matching task in Chinese (judge whether the two sentences convey the same meaning), Yang et al. [17] proposed a sentence-level BERT integrated with audio features to forecast the price of a financial asset over a certain period, and Zhang et al. [19] designed Hierarchical Bidirectional Encoder Representations from Transformers (HIBERT) for document summarisation. Though all used sentence-level information, these works vary from the concrete design, pre-training/fine-tuning paradigm, the datasets used, the task performed and the domain. Moreover, these studies assume an abundance of labelled data instead of focusing on the limited labelled data scenarios.

The contributions of this paper are: (1) we propose a sentence-level BERT-based model, the *Hierarchical BERT Model*, that captures the sentence structure information in a document and an evaluation experiment that demonstrates that HBM is effective in scenarios with small amounts of labelled data; (2) We demonstrate that the salient sentences identified by the HBM are useful as explanations for document classification, and can help annotators label documents faster based on a user study.

### 3 The Hierarchical BERT Model

Figure 1 summarises the proposed sentence-level Hierarchical BERT Model for document classification which consists of 3 components: (1) the token-level Roberta encoder; (2) the sentence-level BERT encoder; and (3) the prediction layer. To make predictions using this model, raw text is first fed into a token-level Roberta encoder to extract text features and form the vector representation for each sentence. Subsequently, the sentence vectors generated by the token-level

Roberta encoder are used as input for the sentence-level BERT encoder. Finally, the intermediate representation generated by the sentence-level BERT encoder is input into the prediction layer that predicts the document class. The detailed implementations of each component are described in the following sections.



**Fig. 1.** The Hierarchical BERT Model for document classification.

### 3.1 Token-level Roberta Encoder

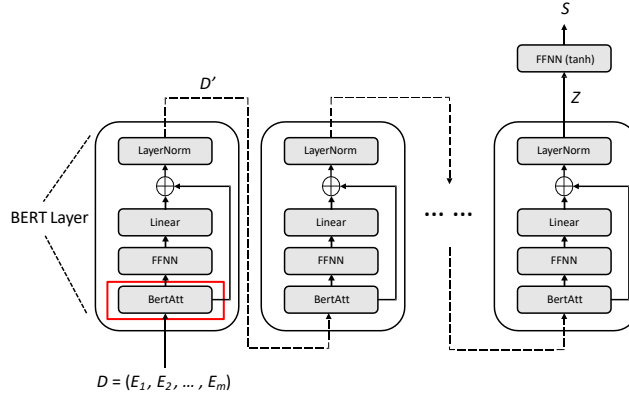
Using the same architecture of BERT but optimising the key hyper-parameters and pre-training with larger datasets, Roberta [9] is recognised as a state-of-the-art approach across many NLP tasks and adopted as the word-level foundation of HBM. Briefly, the multi-head self-attention mechanism in Roberta can effectively capture the semantic and syntactic information between words in a document. In this work, we leverage a pre-trained Roberta model to extract raw text features which act as the input to a sentence-level BERT encoder. The output of this component can be represented as follows:

$$\mathcal{D} = (E_1, E_2, \dots, E_m) \quad (1)$$

Here  $E_i$  denotes the  $i^{th}$  sentence vector (of dimension  $d_e$  by averaging token embeddings contained in each sentence) of document  $\mathcal{D} \in \mathbb{R}^{m \times d_e}$  generated by the token-level Roberta encoder, where  $m$  is the maximum number of sentences set by the user (padding with zeros if the number of sentences is less than  $m$ ).

### 3.2 Sentence-level BERT Encoder

The sentence-level Hierarchical BERT encoder extracts the sentence structure information for generating an intermediate document representation,  $\mathcal{S}$ , based on the sentence vectors,  $\mathcal{D}$ , output by the token-level Roberta encoder. As shown in Figure 2, the sentence-level BERT encoder consists of several identical BERT layers and one feedforward neural network with a tanh activation. Initially, sentence vectors  $\mathcal{D}$  from the token-level Roberta encoder are input together into a BERT layer, more specifically, the BertAtt layer where the multi-head self-attention mechanism is applied (denoted by a red rectangle in Figure 2).



**Fig. 2.** The BERT layers in the sentence-level BERT encoder.

The output of the BertAtt layer is calculated as follows:

$$\text{BertAtt}(\mathcal{D}) = \text{LayerNorm}(\mathcal{D} + \text{MultiHead}(\mathcal{D})) \quad (2a)$$

$$\text{MultiHead}(\mathcal{D}) = \text{Concat}(\text{head}_1, \text{head}_2, \dots, \text{head}_h) \times W^O \quad (2b)$$

where  $h$  is the number of heads in the BertAtt layer; LayerNorm is layer normalization;  $W^O \in \mathbb{R}^{hd_e \times d_e}$  is the weight matrix for dimension transformation; and  $\text{head}_i$  is the attention of the  $i^{\text{th}}$  head. Each  $\text{head}_i$  value is calculated as follows:

$$\text{head}_i = \text{Attention}(Q_i, K_i, V_i) = \text{Softmax}(Q_i \times K_i^\top / \sqrt{d_e}) \times V_i \quad (3a)$$

$$Q_i = \mathcal{D} \times W_i^Q \quad K_i = \mathcal{D} \times W_i^K \quad V_i = \mathcal{D} \times W_i^V \quad (3b)$$

where  $W_i^Q, W_i^K, W_i^V \in \mathbb{R}^{d_e \times d_e}$  are weight matrices for the  $i^{\text{th}}$  head; and,  $\text{Softmax}(Q_i \times K_i^\top / \sqrt{d_e})$  is an  $m \times m$  matrix in which the entry at the  $a^{\text{th}}$  row and  $j^{\text{th}}$  column denotes the attention weight that the  $a^{\text{th}}$  sentence pays to the  $j^{\text{th}}$  sentence. Here,  $V_i$  contains information from sentences to pass through the

following layers while the attention weights matrix,  $\text{Softmax}(Q_i \times K_i^\top / \sqrt{d_e})$ , acts as a gate to control how much information can be passed through (i.e. after the multiplication, it is hard for sentences with low attention score to further propagate their information).

Then, the output of BertAtt is passed through a standard feedforward neural network with the residual mechanism and layer normalisation which is computed as:

$$\mathcal{D}' = \text{LayerNorm}(\text{BertAtt}(\mathcal{D}) + \text{Relu}(\text{BertAtt}(\mathcal{D}) \times W^r) \times W^S) \quad (4)$$

where  $\text{BertAtt}(\mathcal{D}) \in \mathbb{R}^{m \times d_e}$  is the output of the BertAtt layer calculated by Equation 2a;  $W^r \in \mathbb{R}^{d_e \times n d_e}$  is a weight matrix that transforms the dimension of  $\text{BertAtt}(\mathcal{D})$  from  $d_e$  to  $n d_e$  (we set  $n$  to 4 in our work); and  $W^S \in \mathbb{R}^{n d_e \times d_e}$  is the weight matrix with dropout to transform the dimension back to  $d_e$ .  $\mathcal{D}'$  then goes into multiple identical BERT layers to form the matrix  $\mathcal{Z}$  which is used for computing the intermediate document representation  $\mathcal{S}$ .

The intermediate document representation  $\mathcal{S} \in \mathbb{R}^{1 \times d_e}$  output by the sentence-level BERT encoder can be computed as (see Figure 2):

$$\mathcal{S} = \text{Tanh}(\text{Avg}(\mathcal{Z}) \times W^t) \quad (5)$$

where  $W^t \in \mathbb{R}^{d_e \times d_e}$  is a weight matrix for linear transformation,  $\mathcal{Z} \in \mathbb{R}^{m \times d_e}$  is the output of the multiple BERT layers, Avg denotes the mean pooling layer (i.e.  $\text{Avg}(\mathcal{Z}) \in \mathbb{R}^{1 \times d_e}$ ).

Finally, the output of the sentence-level BERT Encoder,  $\mathcal{S}$ , is passed to the prediction layer.

### 3.3 Prediction Layer

The prediction layer is appended to the sentence-level BERT encoder for final prediction as shown in Figure 1. The output,  $\mathcal{S}$ , from the sentence-level BERT encoder is used to compute the raw score for each class, the process can be shown as:

$$[t_0, t_1, \dots, t_y] = \mathcal{S} \times W \quad (6)$$

where  $W \in \mathbb{R}^{d_e \times y}$  is a weight matrix and  $t_0, t_1, \dots, t_y$  represents the raw score of the document for the  $0^{th}$  to  $y^{th}$  class respectively which is fed into a softmax head later for computing loss.

### 3.4 Training a Hierarchical BERT Model (HBM)

It is common for BERT-like models to exploit the pre-training + fine-tuning procedure to train the model [4,9]. More specifically, in pre-training, the model is trained via several self-supervised methods with various, large, unlabelled datasets. Then in fine-tuning, a task-specific head is added to the model to

further update all parameters of the model by supervised learning over small local datasets.

However, the pre-training and fine-tuning is a little different in the proposed HBM approach. Since the Roberta encoder used in HBM is pre-trained with large generic corpora, we assume that the Roberta encoder can encode the token-level information well. Hence, we do not further adjust the pre-trained Roberta model during fine-tuning. In other words, in fine-tuning, we freeze the weights of the Roberta encoder while weights in the sentence-level BERT encoder, as well as in the prediction layer, are randomly initialised and updated based on our prediction objectives. This is illustrated in Figure 1 through the dashed green and blue rectangles.

The objective function optimised during training is a weighted cross entropy:

$$\text{loss}(t, y) = \text{weight}[y](-\log(\exp(t_y) / \sum_j^z \exp(t_j))) \quad (7)$$

where  $y \in [0, 1, \dots, z]$  is the index of the ground truth class,  $t_y$  denotes the raw, unnormalised scores for the  $y^{\text{th}}$  class computed by Equation 6 and  $\text{weight}[y]$  implies the factor given to the  $y^{\text{th}}$  class proportional to the class ratio.

## 4 Experiment 1: Evaluating the Effectiveness of HBM

This section describes the design of an experiment performed to evaluate the performance of HBM for document classification, especially when training datasets are small, by comparing it to other state-of-the-art methods. This addresses the first research question, **RQ1**. The following subsections describe the datasets used in the experiments, the set up of several baselines as well as HBM, and the experimental results.

### 4.1 Datasets

We evaluate the effectiveness of our model on six fully-labelled binary document classification datasets that primarily consist of long documents (from news feeds, blogs, etc.), we also include datasets that consist of short documents to investigate the performance of HBM on short texts:

- **Movie Reviews:** 1,000 *positive* sentiment and 1,000 *negative* sentiment processed movie reviews collected from *rottentomatoes.com* [12]. The average number of sentences per document is 33.12 and the maximum number of sentences is 114.
- **Multi-domain Customer Review:** 3,999 *negative* sentiment and 4,000 *positive* sentiment processed customer reviews collected from *Amazon* from the books, DVDs, electronics and kitchen & housewares sections [2]. The average number of sentences per document is 3.78, these are not long documents, and the maximum number of sentences is 181.

- **Blog Author Gender:** 1,678 *male* blogs and 1,548 *female* blogs [11]. The average number of sentences per document is 22.83 and the maximum number of sentences is 511.
- **Guardian 2013:** 843 *politics* news and 843 *business* news items collected from the Guardian across 2013 [1]. The average number of sentences per document is 27.88 and the maximum number of sentences is 724.
- **Reuters:** 2,369 *acq* news items and 2,369 *earn* news items. The average number of sentences per document is 7.40, these are not long documents, and the maximum number of sentences is 95.
- **20 Newsgroups:** 578 *comp.sys.mac.hardware* newsgroup documents and 590 *comp.sys.pc.hardware* newsgroup documents. The average number of sentences per document is 10.61 and the maximum number of sentences is 367.

## 4.2 Baselines and Setup

Partially following [15], we compare HBM with several baseline methods. We selected an SVM model using a document representation based on FastText, an SVM model using text representations based on a pre-trained Roberta model and a fine-tuned Roberta model as baselines due to their strong performance. We also used another sentence-level model, Hierarchical Attention Network (HAN) [18], as a baseline.

**FastText + SVM:** We use 300-dimensional word vectors constructed by a FastText language model pre-trained with the Wikipedia corpus [7]. Averaged word embeddings are used as the representation of the document. For preprocessing, all text is converted to lowercase and we remove all punctuation and stop words. SVM is used as the classifier. We tune the hyper-parameters of the SVM classifier using a grid-search based on 5-fold cross-validation performed on the training set, after that, we re-train the classifier with optimised hyper-parameters. This hyper-parameter tuning method is applied in Roberta + SVM as well.

**Roberta + SVM:** We use 768-dimensional word vectors generated by a pre-trained Roberta language model [9]. We do not fine-tune the pre-trained language model and use the averaged word vectors as the representation of the document. Since all BERT-based models are configured to take as input a maximum of 512 tokens, we divided the long documents with  $W$  words into  $k = W/511$  fractions, which is then fed into the model to infer the representation of each fraction (each fraction has a “[CLS]” token in front of 511 tokens, so, 512 tokens in total). Based on the approach of [14], the vector of each fraction is the average embeddings of words in that fraction, and the representation of the whole text sequence is the mean of all  $k$  fraction vectors. For preprocessing, the only operation performed is to convert all tokens to lowercase. SVM is used as the classifier.

**Fine-tuned Roberta:** For the document classification task, fine-tuning Roberta means adding a softmax layer on top of the Roberta encoder output and fine-tuning all parameters in the model. In this experiment, we fine-tune the same 768-dimensional pre-trained Roberta model with a small training set. The settings of all hyper-parameters follow [9]. we set the learning rate to  $1 \times 10^{-4}$  and the batch size to 4, and use the Adam optimizer [8] with  $\epsilon$  equals to  $1 \times 10^{-8}$  through

hyperparameter tuning. However, since we assume that the amount of labelled data available for training is small, we do not have the luxury of a hold out validation set to use to implement early stopping during model fine tuning. Instead, after training for 15 epochs we roll back to the model with the lowest loss based on the training dataset. This rollback strategy is also applied to HAN and HBM due to the limited number of instances in training sets. For preprocessing, the only operation performed is to convert all tokens to lowercase.

**Hierarchical Attention Network:** Following [18], we apply two levels of Bi-GRU with attention mechanism for document classification. All words are first converted to word vectors using GloVe [13] (300 dimension version pre-trained using the wiki gigaword corpus) and fed into a word-level Bi-GRU with attention mechanism to form sentence vectors. After that, a sentence vector along with its context sentence vectors are input into sentence-level Bi-GRU with attention mechanism to form the document representation which is then passed to a softmax layer for final prediction. For preprocessing, the only operation performed is to convert all tokens to lowercase, and separate documents into sentences.<sup>4</sup>

**Hierarchical BERT Model:** For HBM, we set the number of BERT layers to 4, and the maximum number of sentences to 114, 64, 128, 128, 100, and 64 for the *Movie Review*, *Multi-domain Customer Review*, *Blog Author Gender*, *Guardian 2013*, *Reuters* and *20 Newsgroups* datasets respectively, these values are based on the length of documents in these datasets. After some preliminary experiments, we set the attention head to 1, the learning rate to  $2 \times 10^{-5}$ , dropout probability to 0.01, used 50 epochs, set the batch size to 4 and used the Adam optimizer with  $\epsilon$  equals to  $1 \times 10^{-8}$ . The only text preprocessing operation performed is to convert all tokens to lowercase and split documents into sentences.

To simulate the scenarios of limited training data, we randomly subsample a small set of data from the fully labelled dataset as the training data. For each dataset, we subsample 50, 100, 150 and 200 instances respectively and we always use the full dataset minus 200 instances as the test set. For each experiment, the training is repeated 10 times with training datasets subsampled by 10 different random seeds and we report the averaged result of 10 repetitions. We use area under the ROC curve (AUC) to measure the classification performance.

### 4.3 Results and Analysis

Table 1 summarises the results of the evaluation experiments performed. In Table 1, the leftmost column denotes the modelling approach, the topmost row denotes the datasets used in the experiment, and  $n$  is the number of instances in the training set. The numbers in brackets stand for the ranking of each method when compared to the performance of other approaches for a dataset with a specific number of labelled instances. The Avg Rank column denotes the average ranking of each model for a specific dataset. The best performance in each column is highlighted in bold. It is clear from Table 1 that the performance of each method improves as the number of instances in the training set increases. The performance of the proposed HBM method consistently surpasses the performance

<sup>4</sup> We apply Python NLTK `sent_tokenize` function to split documents into sentences.

**Table 1.** Performance measured using AUC score of different methods that are trained with limited labelled data across six datasets.

Dataset	Movie Review				
#Instances	n = 50	n = 100	n = 150	n = 200	Avg Rank
FastText	0.6653±0.171 (5)	0.7942±0.020 (4)	0.8040±0.018 (4)	0.8260±0.010 (4)	4.25
Roberta	0.8743±0.032 (3)	0.9132±0.025 (3)	0.9397±0.010 (3)	0.9449±0.008 (3)	3
Fine-tuned Roberta	0.8878±0.018 (2)	0.9298±0.013 (2)	0.9451±0.007 (2)	0.9536±0.005 (2)	2
HAN	0.7013±0.096 (4)	0.7789±0.014 (5)	0.7504±0.082 (5)	0.8128±0.011 (5)	4.75
HBM	<b>0.9139±0.020 (1)</b>	<b>0.9420±0.011 (1)</b>	<b>0.9572±0.007 (1)</b>	<b>0.9638±0.006 (1)</b>	<b>1</b>
Dataset	Blog Author Gender				
#Instances	n = 50	n = 100	n = 150	n = 200	Avg Rank
FastText	0.6400±0.043 (4)	0.6669±0.021 (5)	0.6899±0.009 (4)	0.6861±0.023 (5)	4.5
Roberta	0.5538±0.153 (5)	0.6783±0.025 (3)	0.7058±0.017 (3)	0.7213±0.012 (3)	3.5
Fine-tuned Roberta	0.6462±0.036 (2)	0.6892±0.021 (2)	0.7177±0.019 (2)	0.7295±0.024 (2)	2
HAN	0.6402±0.017 (3)	0.6670±0.010 (4)	0.6845±0.008 (5)	0.6876±0.014 (4)	4
HBM	<b>0.6820±0.025 (1)</b>	<b>0.7150±0.031 (1)</b>	<b>0.7371±0.007 (1)</b>	<b>0.7488±0.013 (1)</b>	<b>1</b>
Dataset	Reuters				
#Instances	n = 50	n = 100	n = 150	n = 200	Avg Rank
FastText	0.9757±0.010 (4)	0.9795±0.005 (5)	0.9851±0.004 (5)	0.9862±0.003 (5)	4.75
Roberta	0.9838±0.007 (2)	0.9890±0.003 (3)	0.9931±0.001 (3)	0.9930±0.001 (3)	2.75
Fine-tuned Roberta	<b>0.9885±0.005 (1)</b>	<b>0.9933±0.002 (1)</b>	0.9953±0.001 (2)	0.9955±0.001 (2)	<b>1.5</b>
HAN	0.9270±0.038 (5)	0.9804±0.005 (4)	0.9865±0.003 (4)	0.9897±0.002 (4)	4.25
HBM	0.9825±0.008 (3)	0.9917±0.004 (2)	<b>0.9980±0.003 (1)</b>	<b>0.9990±0.001 (1)</b>	1.75
Dataset	Multi-domain Customer Review				
#Instances	n = 50	n = 100	n = 150	n = 200	Avg Rank
FastText	0.6694±0.047 (4)	0.6927±0.030 (4)	0.7226±0.030 (5)	0.7471±0.018 (5)	4.5
Roberta	0.8317±0.020 (2)	0.8558±0.036 (2)	0.8976±0.019 (2)	0.9190±0.006 (2)	2
Fine-tuned Roberta	<b>0.9110±0.036 (1)</b>	<b>0.9437±0.007 (1)</b>	<b>0.9534±0.003 (1)</b>	<b>0.9565±0.004 (1)</b>	<b>1</b>
HAN	0.6497±0.021 (5)	0.6907±0.011 (5)	0.7312±0.014 (4)	0.7739±0.034 (4)	4.5
HBM	0.7669±0.024 (3)	0.8342±0.014 (3)	0.8615±0.010 (3)	0.8913±0.004 (3)	3
Dataset	Gurdian 2013				
#Instances	n = 50	n = 100	n = 150	n = 200	Avg Rank
FastText	0.9720±0.003 (3)	0.9789±0.005 (5)	0.9794±0.004 (5)	0.9789±0.005 (5)	4.5
Roberta	0.9694±0.011 (4)	0.9814±0.003 (3)	0.9860±0.003 (2)	0.9852±0.002 (3)	3
Fine-tuned Roberta	0.9727±0.010 (2)	0.9848±0.002 (2)	0.9854±0.003 (3)	0.9864±0.002 (2)	2.25
HAN	0.9684±0.005 (5)	0.9794±0.001 (4)	0.9850±0.002 (4)	0.9849±0.002 (4)	4.25
HBM	<b>0.9740±0.013 (1)</b>	<b>0.9862±0.007 (1)</b>	<b>0.9904±0.003 (1)</b>	<b>0.9925±0.001 (1)</b>	<b>1</b>
Dataset	20Newsgroups				
#Instances	n = 50	n = 100	n = 150	n = 200	Avg Rank
FastText	<b>0.7052±0.030 (1)</b>	0.7516±0.033 (3)	0.7827±0.012 (3)	0.8094±0.013 (3)	2.5
Roberta	0.5969±0.100 (5)	0.6988±0.024 (5)	0.7117±0.030 (5)	0.7436±0.011 (4)	4.75
Fine-tuned Roberta	0.6098±0.051 (4)	0.7555±0.041 (2)	0.8576±0.060 (2)	0.8838±0.045 (2)	2.5
HAN	0.6296±0.030 (3)	0.7142±0.005 (4)	0.7295±0.012 (4)	0.7417±0.029 (5)	4
HBM	0.6883±0.094 (2)	<b>0.8168±0.024 (1)</b>	<b>0.8579±0.021 (1)</b>	<b>0.9158±0.018 (1)</b>	<b>1.25</b>

of the other state-of-the-art methods in the majority of experiments. This clearly demonstrates the effectiveness of the HBM approach under the scenario of limited labelled data.

The performance of HBM is not good for the *Multi-domain Customer Review* and *Reuters* datasets (even lower than that of the approach based on the pre-trained Roberta model for the former). We hypothesise that this is caused by the low number of sentences per document in these datasets. If we refer to Section 4.1, we see that the average number of sentences per document in the *Multi-domain Customer Review* and *Reuters* dataset is 3.78 and 7.40 respectively, which is much lower than that of the other datasets and this leads to a situation in which HBM cannot effectively extract the sentence structure information.

We can conclude that compared to several baseline methods, the proposed Hierarchical BERT Model gives superior performance across the datasets containing long documents, which is evidenced by the higher AUC scores. This clearly demonstrates the effectiveness of the proposed model when limited labelled data

is available, but also highlights that the approach only works well in domains with multiple sentences per document.

## 5 Extracting Salient Sentences Using a HBM Model

As mentioned previously, one advantage of the HBM approach is that the attention scores assigned to sentences in a document can be used as an indication of the importance, or saliency, of those sentences in determining the class the document belongs to. The value of this could be that, the low attention sentences can be ignored and the high attention sentences can be selected to highlight to the user to reduce the effort required to annotate a document. In other words, highlighting the salient sentences in a document can be used to facilitate an explanation of the classifications that the model produces. In this section, we will introduce how the HBM infer salient sentences.

### 5.1 How to Infer Salient Sentences

As described in Equation 3,  $\text{Softmax}(Q_i \times K_i^\top / \sqrt{d_e})$  is a matrix of attention weights used to compute the weighted average of the self-attention heads in the BertAtt layer of HBM. In this matrix, entry  $e^{aj}$  (at the  $a^{\text{th}}$  row and  $j^{\text{th}}$  column) is the attention weight that the  $a^{\text{th}}$  sentence pays to the  $j^{\text{th}}$  sentence. In other words, the  $j^{\text{th}}$  column represents all the attention weights the  $j^{\text{th}}$  sentence received from other sentences and we adopt the sum of all entries in the  $j^{\text{th}}$  column as the saliency score of the  $j^{\text{th}}$  sentence. Therefore, we can compute the saliency score of each sentence in a document by summing up all entries in the corresponding column of the attention weights matrix.

After obtaining the saliency score of all sentences in a document, we select the salient sentences based on the ratio of the sentence saliency score to the highest saliency score. We regard all sentences with a ratio bigger than 0.9 as salient sentences in a document and disregard other sentences. We believe that, besides the sentence with the highest saliency score, those with a saliency score close to the highest saliency score should also contribute intensively to the document representation and should be highlighted. In the next subsection, we will show the effectiveness of highlighting salient sentences through a user study.

## 6 User Study

It is interesting to understand to what extent we can take advantage of important sentences identified by the HBM to help the human with the annotation task. Therefore, the objective of this study is to investigate whether or not highlighting the important sentences in a text makes it easier for labelers to provide these category labels, thus facilitating labelling of large numbers of texts in a short time. This addresses the second research question **RQ2**.

In this user study, we use the ten documents with highlighted sentences inferred by HBM from the *Movie Review* and *Multi-domain Customer Review datasets*

and another ten documents with the same text shown without highlighting.<sup>5</sup> Figures 3 and 4 shows a pair of examples of one of these.

Instructions

The following article is a product/movie review from Amazon/IMDB.

The highlighted sentences (shaded by blue) represent the important sentences of the review.

Please classify the review as a positive or a negative review, then click on the "positive" button at the bottom if you think this is a positive review or click on the "negative" button at the bottom if you think this is a negative review.

10/10

Okay, I bought for the price. And it is effective enough. **But it is very clumsy and unergonomic, so that after a couple minutes of use it is difficult to handle any longer.** The main problems are: Balance of weight, and the switch.

First, the switch: Because it is a cheap push-in button, which must be constantly depressed to keep the knife on, you lose optimum control quickly. The handle is badly placed for this, as your thumb must always be on TOP of the handle and forward toward the front of it, but you can't hold your hand up there because your wrist angle will be completely wrong for cutting. In addition, keeping that button depressed is difficult not only because it does require a constant application of pressure at a different direction from both the grasp and the application of the knife, but because the button sinks deep into the handle, so your thumb is inside the handle as you're trying to hold and control the knife: Quick cramp. Moreover, as you're pushing it in, if the food you're cutting is moist (like a turkey), juices may run onto the handle, your thumb--and into the hole into which you're pressing the button. A couple times I've wondered whether that shooting pain in my thumb was from nerve torture or electric shock; either way, it wasn't good.

**Then the weight distribution of the knife is not very good.** Look at the huge bulk of the motor--and then look at where the handle is located, and how it is designed. There's no way to hold this at a proper angle with the handle shaped as it is. Add in the need to be depressing the power button as you're holding it awkwardly.

The review is:

Positive	Negative
----------	----------

**Fig. 3.** A screenshot of one page of the user study website showing a document with highlighting.

Instructions

The following article is a product/movie review from Amazon/IMDB.

Please classify the review as a positive or a negative review, then click on the "positive" button at the bottom if you think this is a positive review or click on the "negative" button at the bottom if you think this is a negative review.

1/10

Okay, I bought for the price. And it is effective enough. But it is very clumsy and unergonomic, so that after a couple minutes of use it is difficult to handle any longer. The main problems are: Balance of weight, and the switch.

First, the switch: Because it is a cheap push-in button, which must be constantly depressed to keep the knife on, you lose optimum control quickly. The handle is badly placed for this, as your thumb must always be on TOP of the handle and forward toward the front of it, but you can't hold your hand up there because your wrist angle will be completely wrong for cutting. In addition, keeping that button depressed is difficult not only because it does require a constant application of pressure at a different direction from both the grasp and the application of the knife, but because the button sinks deep into the handle, so your thumb is inside the handle as you're trying to hold and control the knife: Quick cramp. Moreover, as you're pushing it in, if the food you're cutting is moist (like a turkey), juices may run onto the handle, your thumb--and into the hole into which you're pressing the button. A couple times I've wondered whether that shooting pain in my thumb was from nerve torture or electric shock; either way, it wasn't good.

Then the weight distribution of the knife is not very good. Look at the huge bulk of the motor--and then look at where the handle is located, and how it is designed. There's no way to hold this at a proper angle with the handle shaped as it is. Add in the need to be depressing the power button as you're holding it awkwardly.

The review is:

Positive	Negative
----------	----------

**Fig. 4.** A screenshot of one page of the user study website showing a document with no highlighting.

<sup>5</sup> The user study platform information and all these ten documents with highlighted sentences are available at <https://github.com/GeorgeLuImmortal/Hierarchical-BERT-Model-with-Limited-Labelled-Data>

## 6.1 Experiment Setup

The user study is run on an online platform and data collected is stored anonymously. We recruited participants by circulating an advertisement among social media such as Twitter and Facebook, and internal email lists of two Irish research centres.<sup>6</sup> In this study, all participants were required to be fluent English speakers. Participants are asked to state the categories to which a series of texts (i.e. all ten documents with highlighting or all ten documents without highlighting) belong. They are shown a text and asked to indicate which category, from a list of options (i.e. positive sentiment or negative sentiment) provided, best describes it. Ground truth categories for the texts used in the study exist, and these were used to measure the accuracy of the category labels provided by participants. The time taken for participants to select a category for each document were also collected to measure the annotation efficiency.

A between-subject design was used for the experiment. Half of the participants were shown texts in which important sentences have been highlighted (i.e. highlighting condition group) to make it easier for them to come to a judgement on the category that best describes the text. Half of the participants were shown texts without highlighting (no-highlighting condition group). All documents were presented to the participants in a random sequence. The accuracy of labels provided and the time taken to provide them were compared across the two groups to measure the impact of providing highlighting.


## 6.2 Results and Analysis

We had 75 participants in total: 38 participants in the highlighting condition group and 37 participants in the no-highlighting condition group. To mitigate against participants who didn't pay careful attention to the tasks in the study, we removed the records of participants who achieved the accuracy less than 0.6 or spent less than 90 seconds (i.e. a minute and a half) in total on all ten questions. Similarly, to mitigate against participants who left the study for a long period of time, we also removed data from participants who spent more than 420 seconds (i.e. 7 minutes) on providing the answer for a single document. After this cleaning, we finally have 57 valid participants, among these participants 32 are in highlighting condition group and 25 are in no-highlighting condition group.

Figure 5 shows a box plot of the time spent per user in providing category labels for each document, where red boxes denote the performance of the highlighting condition group, and blue boxes denote the performance of the no-highlighting condition group. It is obvious that except for documents CR-2 and MR-3, the time cost in the highlighting condition group is much less than that in the no-highlighting condition group.

We also analyse the results from the user perspective, the average total time cost in labelling all ten documents per user in the highlighting condition

<sup>6</sup> The Insight Centre for Data Analytics (<https://www.insight-centre.org/>) and ADAPT Centre (<https://www.adaptcentre.ie/>).



**Fig. 5.** A box plot of the time cost per user on each document. The vertical axis is the time cost per document (unit second); the horizontal axis represents the documents.

group is 375.50 seconds, and in the no-highlighting condition group is 645.54 seconds. Following the advice in [10] we analyse these results using a Mann Whitney U test in which the  $p$ -value is 0.01939, which indicates that we can reject the null hypothesis at the 0.05 confidence level and consider a significant difference between the time cost of the two groups. The average accuracy achieved by participants in the highlighting condition group is 0.9575, while the average accuracy achieved by participants in the no-highlighting condition group is 0.9448. The  $p$ -value of a Mann-Whitney U test on average accuracy is 0.3187, which indicates no significant difference between the average accuracy of these two groups exists.

Based on the previous analysis, we conclude that highlighting important sentences identified by the HBM can significantly improve the human’s annotation efficiency while still maintaining a high annotation quality.

## 7 Conclusions

In this work, we have proposed a novel approach, the Hierarchical BERT Model (HBM), which has been demonstrated to perform well with limited labelled data especially when documents are long. Several experiments have been carried out to demonstrate the superior performance of the HBM as compared to other state-of-the-art methods under the assumption that training sets are small. We also demonstrated that using a HBM facilitates selecting salient sentences from a document and showed through a user study that highlighting these sentences leads to more efficient labelling.

## References

1. Belford, M., Mac Namee, B., Greene, D.: Stability of topic modeling via matrix factorization. *Expert Systems with Applications* **91**, 159–169 (2018)

2. Blitzer, J., Dredze, M., Pereira, F.: Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In: 45th annual meeting of the association of computational linguistics. pp. 440–447 (2007)
3. Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. arXiv preprint arXiv:1406.1078 (2014)
4. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
5. Ding, M., Zhou, C., Yang, H., Tang, J.: Cogltx: Applying bert to long texts. *Advances in Neural Information Processing Systems* **33** (2020)
6. Goodfellow, I., Bengio, Y., Courville, A., Bengio, Y.: *Deep learning*, vol. 1 (2016)
7. Joulin, A., Grave, E., Bojanowski, P., Douze, M., Jégou, H., Mikolov, T.: Fasttext.zip: Compressing text classification models. arXiv preprint arXiv:1612.03651 (2016)
8. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
9. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019)
10. Marusteri, M., Bacarea, V.: Comparing groups for statistical differences: how to choose the right statistical test? *Biochemia medica: Biochemia medica* **20**(1), 15–32 (2010)
11. Mukherjee, A., Liu, B.: Improving gender classification of blog authors. In: 2010 conference on Empirical Methods in natural Language Processing. pp. 207–217 (2010)
12. Pang, B., Lee, L.: A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In: 42nd annual meeting on Association for Computational Linguistics. p. 271 (2004)
13. Pennington, J., Socher, R., Manning, C.: Glove: Global vectors for word representation. In: Conference on empirical methods in natural language processing (EMNLP). pp. 1532–1543 (2014)
14. Sun, C., Qiu, X., Xu, Y., Huang, X.: How to fine-tune bert for text classification? In: China National Conference on Chinese Computational Linguistics. pp. 194–206 (2019)
15. Usherwood, P., Smit, S.: Low-shot classification: A comparison of classical and deep transfer machine learning approaches. arXiv preprint arXiv:1907.07543 (2019)
16. Xu, Y., Liu, Q., Zhang, D., Li, S., Zhou, G.: Many vs. many query matching with hierarchical bert and transformer. In: CCF International Conference on Natural Language Processing and Chinese Computing. pp. 155–167 (2019)
17. Yang, L., Ng, T.L.J., Smyth, B., Dong, R.: Hhtml: Hierarchical transformer-based multi-task learning for volatility prediction. In: The Web Conference 2020. pp. 441–451 (2020)
18. Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., Hovy, E.: Hierarchical attention networks for document classification. In: Conference of the North American chapter of the association for computational linguistics: human language technologies. pp. 1480–1489 (2016)
19. Zhang, X., Wei, F., Zhou, M.: Hibert: Document level pre-training of hierarchical bidirectional transformers for document summarization. arXiv preprint arXiv:1905.06566 (2019)