



# Supervised Contrastive Learning for Interpretable Long-Form Document Matching

AKSHITA JHA, Virginia Tech, Arlington, VA

VINEETH RAKESH and JAIDEEP CHANDRASHEKAR, InterDigital, CA

ADITHYA SAMAVEDHI and CHANDAN K. REDDY, Virginia Tech, Arlington, VA

Recent advancements in deep learning techniques have transformed the area of semantic text matching (STM). However, most state-of-the-art models are designed to operate with *short* documents such as tweets, user reviews, comments, and so on. These models have fundamental limitations when applied to long-form documents such as scientific papers, legal documents, and patents. When handling such long documents, there are three primary challenges: (i) the presence of different contexts for the same word throughout the document, (ii) small sections of contextually similar text between two documents, but dissimilar text in the remaining parts (this defies the basic understanding of “similarity”), and (iii) the coarse nature of a single global similarity measure which fails to capture the heterogeneity of the document content. In this article, we describe **CoLDE: Contrastive Long Document Encoder**—a transformer-based framework that addresses these challenges and allows for interpretable comparisons of long documents. CoLDE uses unique positional embeddings and a multi-headed chunkwise attention layer in conjunction with a supervised contrastive learning framework to capture similarity at three different levels: (i) high-level similarity scores between a pair of documents, (ii) similarity scores between different sections within and across documents, and (iii) similarity scores between different *chunks* in the same document and across other documents. These fine-grained similarity scores aid in better interpretability. We evaluate CoLDE on three long document datasets namely, ACL Anthology publications, Wikipedia articles, and USPTO patents. Besides outperforming the state-of-the-art methods on the document matching task, CoLDE is also robust to changes in document length and text perturbations and provides interpretable results. The code for the proposed model is publicly available at <https://github.com/InterDigitalInc/CoLDE>.

CCS Concepts: • **Information Systems** → **Document Representation**; *Document Structure*; • **Computing Methodologies** → Information Extraction;

Additional Key Words and Phrases: Semantic text matching, long documents, contrastive learning, attention, embeddings, interpretability, transformer, BERT

## ACM Reference format:

Akshita Jha, Vineeth Rakesh, Jaideep Chandrashekar, Adithya Samavedhi, and Chandan K. Reddy. 2023. Supervised Contrastive Learning for Interpretable Long-Form Document Matching. *ACM Trans. Knowl. Discov. Data.* 17, 2, Article 27 (February 2023), 17 pages. <https://doi.org/10.1145/3542822>

This work was supported in part by the US National Science Foundation grant IIS-1838730.

Authors’ addresses: A. Jha, Virginia Tech, 900 N Glebe Rd, Arlington, VA 22203; email: akshitajha@vt.edu; V. Rakesh, InterDigital, 4410 El Camino Real Suite 120, Los Altos, CA 94022; email: vineethrakesh@gmail.com; J. Chandrashekar, InterDigital, 4410 El Camino Real Suite 120, Los Altos, CA 94022; email: jaideep.chandrashekar@interdigital.com; A. Samavedhi, Virginia Tech, 900 N Glebe Rd, Arlington, VA 22203; email: adithyas@vt.edu; C. K. Reddy, Virginia Tech, 900 N Glebe Rd, Arlington, VA 22203; email: reddy@cs.vt.edu.



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

© 2023 Copyright held by the owner/author(s).

1556-4681/2023/02-ART27

<https://doi.org/10.1145/3542822>

## 1 INTRODUCTION

**Semantic text matching (STM)** is an important and challenging problem in the field of information retrieval. Given a *query*, which could range from a few words to a few pages, the objective of STM is to retrieve a set of documents related to the query. While recent advancements in language modeling have certainly produced promising results, their effectiveness is limited to the following tasks: (a) matching short-queries to short-documents, or (b) matching short-queries to long-documents. A few examples of short-document matching include contextual text similarity [23], sentence matching [18], and question ranking [3] where one or both documents (i.e., the query and the target document) have limited textual content. However, *STM between two long documents*, which is a more challenging problem, has not been thoroughly studied in the literature. Some important use cases for long-document matching include ranking of research articles, comparison of patent documents, and clustering of Wikipedia articles. In this article, we exclusively focus on *STM for long documents*. There are three main challenges that arise while developing a model for long document matching.

**Challenge 1.** It is infeasible to learn the semantics of individual words, phrases, and sentences by considering only a few lines of text, as longer documents span several pages. In other words, learning a representative embedding that encapsulates the context of an entire document is extremely challenging. Further evidence of this is provided by the state-of-the-art Transformer models, such as BERT [10], which cannot handle more than 512 tokens (or words) during a single feed-forward pass. Current solutions for addressing these limitations include various pooling techniques which have already proven to be ineffective [35].

**Challenge 2.** Accounting for various levels of (dis)similarity between different sections of text is another challenge. To understand this better, let us consider a scenario where the end-user is interested in finding long documents most similar to the given document (*Query*). For the sake of simplicity, we consider only two possible target long documents (*Target 1* and *Target 2*) to choose from (see Figure 1). In the above three documents (*Query*, *Target 1*, and *Target 2*), there are three levels of similarity: (a) A similar boiler plate text about the popularity of deep learning that does not contribute to the distinguishing content of the documents, (b) Semantically similar words – *Query* has a “semantic match” with both *Target 1* and *Target 2* (“reducing the size” is equivalent to the word “compression”), and (c) Contextually similar content—*Query* is contextually similar only to *Target 2*. The context is provided by statements such as “compressing neural networks” unlike *Target 1* which discusses “image compression”. Here, *Query* is most similar to *Target 2* but in a long-form document, such contextual information cannot be readily learned merely from the initial few lines (or paragraphs) as the distinguishing information might be presented later on in the long-form document.

**Challenge 3.** The final challenge is interpretability. When comparing documents, it is not sufficient to just provide a score that quantifies the amount of similarity. It is also important to explain what makes the documents (dis)similar.

**Model Overview.** To tackle the aforementioned challenges, we propose **CoLDE: Contrastive Long Document Encoder** for matching long documents in an interpretable manner. Intuitively, while learning document representations, *sections within the same document should be closer to each other and sections across different documents should be farther apart in the latent embedding space*. Additionally, *representations of documents that belong to the same class should be closer to each other compared to the documents that belong to a completely different class*. With this intuition, we leverage the structure of a long-form document along with its class label and build a supervised contrastive learning framework [24] for long document matching. CoLDE divides a long

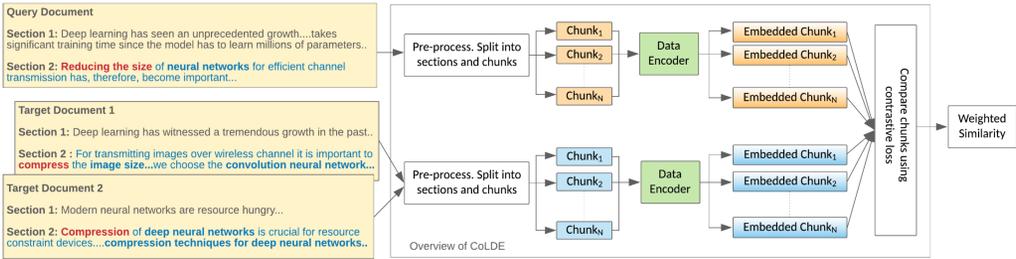


Fig. 1. An example illustrating the challenges in long document matching and an overview of the proposed CoLDE framework. The *Query* document is compared against two target documents *Target 1* and *Target 2*. All documents homogeneously discuss *compression* (albeit the context varies). CoLDE provides a weighted similarity score between different chunks and sections within and across the documents.

document into different sections and uses unique positional embeddings to capture the long document structure for added interpretability. A key challenge in document matching is the coarse nature of a single global similarity score. While some works do demonstrate how attention scores can be used to provide insights into attention-based models [13, 39, 42], we define interpretability as the weighted similarity scores between different sections and different chunks of text in a document. These fine-grained similarity scores can aid the end-users in understanding the reason as to why the two documents are (dis)similar. We provide three levels of similarity scores: (i) similarity score between documents, (ii) similarity scores between sections within and across different documents, and (iii) similarity scores between different text chunks within and across documents. We use a combination of contrastive loss and a multi-headed attention layer for different text chunks to get this weighted similarity score. Figure 1 presents the model overview. We demonstrate the performance of our model using three different semi-structured long document datasets: (i) ACL Anthology Network papers, (ii) Wikipedia articles, and (iii) Patent documents.

Recently, a few attempts have been made to encode long sequences of text using transformers [4, 11, 44]. Our work is significantly different from these approaches since we primarily focus on *end-to-end long document matching*. To the best of our knowledge, our work is the first to utilize the long document structure for contrastive learning in text and provide fine-grained analysis of similarity scores within and across documents.

The primary contributions of this article are as follows:

- Propose a simple yet effective novel supervised contrastive learning model, CoLDE: Contrastive Long Document Encoder, to compare long documents like research articles, patent documents, and Wikipedia articles.
- Develop new data augmentation techniques for text along with unique positional embeddings to capture the long document structure and consequently learn better representations.
- Demonstrate the interpretability provided by the CoLDE model—by providing high-level similarity score between documents, in addition to fine-grained similarity score between different sections, and between different text chunks within and across documents.

The rest of this article is organized as follows. Section 2 describes the related work. Section 3 presents the proposed framework, CoLDE, in detail. Our experimental methodology is discussed in Sections 4 and 5 concludes our work.

## 2 RELATED WORK

**Long Document Matching.** Guo et al. [16] propose a joint deep learning based architecture for ad-hoc retrieval when comparing documents. Several works have also used convolutional

networks [18, 30, 43], with weighting mechanism [40] to generate a final query-document score. Mitra et al. [29] propose a combination model that uses local representations and distributed representations for text matching. Yang et al. [42] propose a hierarchical attention network for document classification whereas Adhikari et al. [1] use BERT for document classification. Jian et al. [21] propose a multi-depth attention-based hierarchical recurrent neural network (SMASH) for long-document matching. However, Yang et al. [41] pre-train a transformer based hierarchical model (SMITH) for text matching that outperforms them across multiple datasets. While our proposed CoLDE model is also based on transformers, we provide fine-grained similarity scores unlike the previous models that only classify whether a document is relevant or irrelevant based on a (coarse) single document level similarity score.

**Long Document Encoding.** Earlier works used Doc2Vec [5, 26] for long document representation. A growing body of literature is now examining the use of transformer based models for long document encoding [7, 17, 25, 27, 32]. Longformer [4], Poolingformer [45], Performers [8], and Linear Transformers [22] adapt transformers to use an attention mechanism to aggregate information for long sequences. Big Bird [44] proposes a sparse attention mechanism that reduces BERT’s quadratic dependency on the sequence length to linear. CogLTX [11] presents a framework that uses multi-step reasoning over key sentences to overcome the insufficient long-range attentions in BERT by using text blocks for rehearsal and decay. Transformer-XL [9] and Compressive Transformers [34] compress the transformers to use attentive sequence over long text. Our work is significantly different from these approaches since we primarily focus on *end-to-end long document matching*.

**Interpretability for Text.** Researchers have devised approaches to measure the feature importance in text [36, 37]. There is an increased interest in using causal frameworks for understanding black-box predictions [2]. Some researchers claim that attention is not an explanation [20]. However, Wiegrefe et al. [39], Yang et al. [42], and Ghaeni et al. [13] demonstrate how attention scores can be used to provide insights into attention-based models. We use multi-headed chunkwise attention in conjunction with supervised contrastive loss to provide fine-grained similarity scores that help in interpreting the final decision made by the model.

**Contrastive Learning.** Contrastive learning has seen tremendous success in recent times. SimCLR [6] and Supervised SimCLR [24] use self-supervised and supervised contrastive learning frameworks, respectively, to learn meaningful image representations that help in downstream tasks. DeCLUTR [14] uses a self-supervised loss function for learning universal sentence embeddings without labels. SimCSE [12] uses dropout as a form of data augmentation for learning sentence representations in a self-supervised as well as a supervised manner. Luo et al. [28] use different semantic paraphrasing techniques for contrastive augmentation to learn better sentence representations. All these methods either focus on images or short-text. To the best of our knowledge, we are the first ones to apply contrastive learning on the task of long-form document matching.

### 3 THE PROPOSED COLDE FRAMEWORK

Inspired by the recent success of contrastive learning algorithms in the computer vision domain [6, 24], we build a simple yet effective contrastive learning framework for long-document matching. Intuitively, CoLDE maximizes the similarity between latent representations of sections within the same document and minimizes the similarity between representations of sections from different documents. Additionally, it ensures that representations of documents belonging to the same class are closer to each other in the latent embedding space when compared to the documents from a different class. CoLDE consists of three primary components: (i) Data Augmentation, (ii) Data Encoder, and (iii) Contrastive Loss Function. Figure 2 presents the complete architecture of the proposed model.

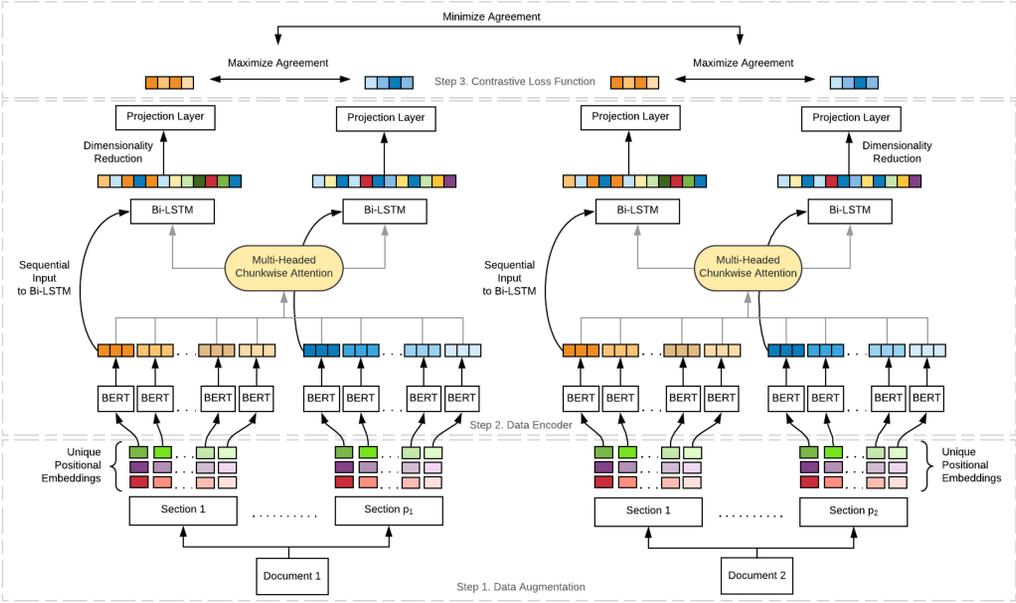


Fig. 2. Model architecture for CoLDE. The contrastive learning framework consists of three primary components: (i) Data Augmentation, (ii) Data Encoder, and (iii) Contrastive Loss Function. It takes as input long documents and divides them into several sections. Each section is further split into chunks of 512 tokens and enhanced with unique positional embeddings (Figure 3) before being given to the data encoder module. The contrastive loss function in conjunction with multi-headed chunkwise attention provide fine-grained similarity scores within and across sections and chunks.

**Problem Statement.** The problem being tackled in this article is defined as follows. Given a query document  $s$ , and a set of target documents  $D_T$ , the goal is to estimate the semantic similarity  $sim(\cdot)$  between  $s$  and a target document  $t \in D_T$  at three different levels: (i) Document level  $\hat{y}_d = sim(s, t)$  for every document pair, (ii) Section level  $\hat{y}_p = sim(s_i, t_j)$ , where  $s_i \in s$  and  $t_j \in t$  for every section pair, and (iii) Chunk level  $\hat{y}_c = sim(c_{q_s}, c_{q_t})$ , where  $c_{q_s} \in s_i$  and  $c_{q_t} \in t_j$  for every chunk pair. The semantically similar target documents, sections, and chunks will have a higher similarity score. The notation used throughout the article and their definitions have been given in Table 1.

### 3.1 Data Augmentation

Long documents, unlike short-text, follow an inherent structure and are organized into sections and subsections (e.g., introduction, related work, and methodology). The task of learning meaningful representations for long documents, can be broken down into learning better representations for their sections and subsections. To achieve this, we divide a long document into different sections as follows:  $\tilde{x}_i, \tilde{x}_j = Aug(x)$  where  $Aug(\cdot)$  is the augmentation module that takes a long document  $x$  as input and extracts different sections  $\tilde{x}_i$  and  $\tilde{x}_j$  from the document. It should be highlighted that data augmentation is different from data preprocessing as it takes a long document and creates more samples from it. This is represented as Step 1 in Figure 2. These sections contain a subset of the text and represent a subset of information present in the original long document. One of the major limitations of BERT is its inability to handle more than 512 tokens at a time, thus making it ineffective for encoding long documents. To overcome this limitation, we divide a document section into multiple chunks of non-overlapping 512 tokens. Additionally, we enhance

Table 1. Table of Notations

| Notation     | Definitions               |
|--------------|---------------------------|
| $r$          | Token ID                  |
| $q$          | Chunk ID                  |
| $p$          | Section ID                |
| $x$          | Long document             |
| $s$          | Query document            |
| $t$          | Target document           |
| $f(\cdot)$   | Encoder module            |
| $sim(\cdot)$ | Similarity function       |
| $Aug(\cdot)$ | Data augmentation module  |
| $\hat{y}_d$  | Document level similarity |
| $\hat{y}_p$  | Section level similarity  |
| $\hat{y}_c$  | Chunk level similarity    |
| $T_r$        | Embedding for token $r$   |
| $C_p$        | Embedding for chunk $q$   |
| $S_p$        | Embedding for section $p$ |

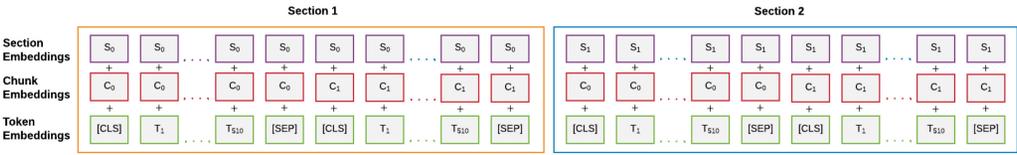


Fig. 3. Visualization of the proposed unique positional embeddings for a long-document with two sections. In this example, each section has two chunks of 512 tokens each. The embeddings are summed up together to form the augmented token embeddings for BERT. They capture the structural information present in the long document.

these input section chunks using unique positional embeddings (described below) to encode the long document structure.

**Unique Positional Embeddings.** The token embeddings from BERT are combined with long document structure aware positional embeddings before providing them as input to the BERT model. These positional embeddings are explained below (shown in Figure 3).

- **Section Embeddings:** Section embeddings  $S_p$  are unique values given to each section where  $p$  represents the section number. Since we divide a long document into two sections, these embeddings take two possible values: 0 and 1. Embeddings for section  $\tilde{x}_i$  take the value of 0 and that for section  $\tilde{x}_j$  take the value of 1.
- **Chunk Embeddings:** Since BERT can only handle 512 tokens at a time, each section is further divided into several chunks of 512 tokens. Each of these chunks is given a unique ID  $C_q$  where  $q$  represents the chunk number.
- **Token Embeddings:** Token embeddings  $T_r$  are the index  $r$  of the token in the input sequence. They are the same as the ones seen in the standard BERT model.

These additional embeddings are summed up together and form the augmented token embeddings  $\tilde{T}_{pqr}$  given by

$$\tilde{T}_{pqr} = \sum_{p=0}^{n_p} \sum_{q=0}^{n_c} \sum_{r=0}^{512} (S_p + C_{pq} + T_{pqr}),$$

where  $n_p$  is the total number of sections and  $n_c$  is the total number of chunks in a section. There are 512 tokens in each chunk. These augmented token embeddings are given to the data encoder for further processing.

### 3.2 Data Encoder

CoLDE's data encoder  $f(\cdot)$  is built over the BERT model [10]. We use a pre-trained BERT model that is fine-tuned during training. The augmented input chunks of 512 tokens are given as input to BERT. Zero-padding is done for chunks smaller than 512 tokens. The encoded BERT chunk embeddings are given to a **bidirectional-LSTM (Bi-LSTM)** [15] layer for aggregation which uses a unique 'multi-headed chunkwise attention' component (described below).

**Multi-headed Chunkwise Attention.** To get chunk level similarity scores, we introduce a multi-headed chunkwise attention layer. This is a self-attention layer [38] that computes attention between the BERT embeddings of different chunks of 512 tokens within and across sections. As a consequence, we can obtain a fine-grained understanding of the chunk that attends the most to other text chunks in the document. The chunk with the highest attention weight w.r.t. a *Query* chunk plays the most important part in computing the similarity score. We compute the multi-headed chunkwise attention between sections by treating each chunk in section  $p$  as a *Query* ( $Q$ ). The chunks of the successive section  $p + 1$  are treated as Keys ( $K$ ) against which the attention is computed. The chunks of section  $p + 1$  are also used in calculating the *Values* ( $V$ ). These vectors are of dimension  $d$ . The weights for calculating the query  $Q$ , keys  $K$ , and values  $V$  are learned during training. The matrix of the outputs is computed as follows:  $Attention(Q, K, V) = softmax(\frac{QK^T}{\sqrt{d}})V$ .

**Dimensionality Reduction.** The encoded BERT output of different chunks, along with the multi-headed chunkwise attention, is given sequentially to a Bi-LSTM layer. The Bi-LSTM layer is required for aggregation of the segmented chunk representations in order to obtain an intermediate high-dimensional section representation  $\tilde{z}_l$  corresponding to section  $\tilde{x}_l$ , where  $l$  is the section number. This intermediate section representation  $\tilde{z}_l = f(\tilde{x}_l)$  is further given to a projection layer  $Proj(\cdot)$  for reducing the dimensionality which results in a final section-level representation  $z_l = Proj(\tilde{z}_l)$ . The final section representation  $z_l$  is used to compute the contrastive loss between sections across query and target documents. The loss is then back-propagated through the network during training.

### 3.3 Contrastive Loss Function

**Supervised Contrastive Loss.** Using the contrastive learning framework, we bring representations of sections within the same documents closer and obtain a similarity score between sections within and across documents. CoLDE also provides document level similarity scores and ensures that documents belonging to the same class are closer to each other in the latent embedding space. To achieve this, we use supervised contrastive loss [24] as our objective function. We randomly sample a set of  $N$  long documents and their corresponding labels  $\{x_k, y_k\}_{k=1}^N$ . This results in  $2N$  data points since each document is divided into two sections for simplicity. In this article, we refer to the set of  $N$  documents as a "batch" and the set of  $2N$  documents as an "augmented batch". The supervised contrastive loss function is described as follows:

$$\mathcal{L}^{sup} = \sum_{i=1}^{2N} \sum_{j=1}^{2N} \frac{-1}{2N y_i - 1} \mathbb{1}_{i \neq j} \cdot \mathbb{1}_{\tilde{y}_i = \tilde{y}_j} \cdot \log \frac{\exp(z_i \cdot z_j / \tau)}{\sum_{k=1}^{2N} \mathbb{1}_{i \neq k} \cdot \exp(z_i \cdot z_k / \tau)}.$$

In the above equation,  $N$  is the batch size;  $z_l = Proj(f(\tilde{x}_l))$  where  $\tilde{x}_l$  is a section;  $f(\cdot)$  is the encoder; and  $Proj(\cdot)$  is the projection layer. The three indicator functions: (i)  $\mathbb{1}_{i \neq j} \in \{0, 1\}$  evaluate

to 1 iff  $i \neq j$ ; (ii)  $\mathbb{1}_{\hat{y}_i=\hat{y}_j} \in \{0, 1\}$  evaluate to 1 when the labels for two sections are the same; and (iii)  $\mathbb{1}_{i \neq k} \in \{0, 1\}$  evaluate to 1 iff  $i \neq k$ . The symbol  $\cdot$  indicates inner (dot) product;  $\tau$  denotes a temperature parameter. The final loss is summed across all the samples. The triplet loss, one of the widely used losses for supervised training, is a special case of the contrastive loss when the batch size  $N = 2$  and contains only one positive and one negative sample.

The numerator incorporates all positive sections in an augmented batch, i.e., every section with the same label in the augmented batch is treated as a positive sample. The denominator, on the other hand, performs a summation over the negative samples. To understand the positive and negative samples better, let us consider a simple scenario where  $N = 3$ . The batch has three documents:  $D_1$ ,  $D_2$ , and  $D_3$  with labels 1, 1, and 0, respectively. Each document is split into two sections which results in an augmented batch of  $2N = 6$ . With respect to the source document  $D_1$ , sections of  $D_1$  and  $D_2$  are considered to be positive samples for each other. Sections from  $D_3$  are considered to be negative samples. Our loss function encourages the encoder to output section representations for  $D_1$  and  $D_2$  to be closer to each other in the latent embedding space. Although the above example only describes the binary class scenario, the loss also generalizes well to a multi-class setting. In such a multi-class scenario, the positive and negative samples are computed w.r.t. to the label of the source document.

**Negative Samples.** For a batch size of  $N$ , all documents having the same label are treated as positive samples and the others are considered to be negative samples w.r.t. a source document. The negative samples belong to one of the three categories: (i) hard negatives, (ii) semi-hard negatives, and (iii) easy negatives. These three categories represent the distance between the representation of the source document and the negative samples in the latent embedding space, ranging from very close for hard negatives to very far for easy negatives. Most models explicitly require these different kinds of negatives as inputs. This is not the case for CoLDE. When using the supervised contrastive loss, hard negative mining is implicitly performed by the model [24]. The discriminatory power of the encoder increases with the increase in the number of negative samples w.r.t. a source document, i.e., with increase in batch size, the number of negative samples increase.

## 4 EXPERIMENTS

In this section, we empirically evaluate the proposed CoLDE framework by studying the following research questions:

- **RQ1:** Does CoLDE perform better than existing methods for downstream tasks such as long document matching?
- **RQ2:** How does the supervised contrastive loss function and the proposed multi-headed attention layer aid in interpretability?
- **RQ3:** How robust is the model to changes in batch size, document length, and text perturbation?
- **RQ4:** What is the impact of different components on the performance of CoLDE?

### 4.1 Methodology

**4.1.1 Datasets.** For our experiments, we consider three different long document datasets: (i) ACL Anthology Network Corpus (AAN), (ii) Wikipedia (WIKI), and (iii) Patent (PAT) datasets.

- **ACL Anthology Network Corpus (AAN)**<sup>1</sup>: The AAN corpus [33] consists of 23,766 articles written by 18,862 authors in 373 venues related to NLP and forms a citation network.

<sup>1</sup><http://aan.how/download/#aanNetworkCorpus>.

Each article is represented by a node with directed edges connecting an article (the parent node) to all its cited articles (children nodes). Articles that have been cited by the parent article are treated as similar samples. We consider this broad notion of similarity which is in line with the previous works in this area [21, 41]. For every similar sample, an irrelevant article is randomly chosen to create a balanced dataset. Sets of similar articles are given the same labels. To prevent leakage of information and make the task more difficult, the references and the abstract sections are removed. Articles without any content are also removed. The dataset consists of 12,665 randomly selected research articles split into two sections. On average, each section has six chunks.

- **Wikipedia (WIKI)**<sup>2</sup>: We use the Wikipedia dump, and adopt a methodology similar to the one proposed by Jiang et al. [21] to process this data. We create a dataset of similar Wikipedia articles by assuming that similar articles have similar outgoing links. The Jaccard similarity between the outgoing links of the source and the target articles is calculated. If the Jaccard similarity  $> 0.5$ , the documents are assumed to be similar, otherwise they are considered dissimilar. Only articles with two or more similar articles are selected. Sets of similar articles are given the same labels. The dataset consists of 59,911 Wikipedia articles split into two sections. On average, each section has three chunks.
- **Patent (PAT)**<sup>3</sup>: This dataset consists of 3,000 documents sampled from the publicly available USPTO database. A patent document is extremely long and primarily consists of (i) Abstract, (ii) Claims, and (iii) Description sections. We make use of the Claims and the Description sections for our experiments. On average, the Claims section has three chunks, and the Description section has 27 chunks. Three human annotators were given pairs of documents and were asked to label them as similar or not based on topics of relevance. They referred to the Abstract, Claims, and CPC Codes<sup>4</sup> of the patents to measure the similarity. The final document content similarity label was based on majority vote.

We report the results of the baseline models on the above datasets.<sup>5</sup>

**4.1.2 Implementation Details.** The proposed CoLDE model was implemented in PyTorch. For the purpose of our experiments, we use the pre-trained BERT<sub>BASE</sub> model provided by the Huggingface library.<sup>6</sup> It has 12 Transformer blocks, a hidden size of 768 and 12 attention blocks. The BERT model is fine-tuned during training. The Bi-LSTM output dimension is 512. The projection layer reduces the dimension to 256. We use a scaled-dot product attention function [38] with four heads as our multi-headed chunkwise attention function. The learning rate is set to  $5e-5$ . We use an Adam optimizer with a weight decay of 0.01. The value for  $\tau$  is fixed to 0.5 for all the experiments. The three datasets are split into 80-10-10 for train, validation, and test sets, respectively. The best performance on the validation set is achieved after training the model for 25 epochs on AAN, after 35 epochs on WIKI, and after 200 epochs on the Patent dataset. We experiment with different batch sizes of 50, 75, 100, and 200. The model training on a batch size of 50 and 75 is done on one 16GB Tesla V100 GPU. For a batch size of 100 and 200, we used one 48GB RTX-8000 GPU.

**4.1.3 Baseline Experiments.** We primarily compare our model with the following long document comparison baselines.

<sup>2</sup><https://dumps.wikimedia.org/enwiki/latest/enwiki-latest-pages-articles.xml.bz2>.

<sup>3</sup><https://github.com/google/patents-public-data>.

<sup>4</sup><https://www.uspto.gov/web/patents/classification/cpc/html/cpc.html>.

<sup>5</sup>Since the pre-processed dataset used in [21, 41, 42] is not public, our results are not exactly aligned with theirs. We use their publicly available code and fine-tune the model on our dataset.

<sup>6</sup>[https://huggingface.co/transformers/model\\_doc/bert.html](https://huggingface.co/transformers/model_doc/bert.html).

- **DSSM [19]**: A siamese deep learning model with three fully-connected feed-forward layers with cosine similarity as its objective function.
- **ARC-I [18]**: A siamese based matching model with three layers of 1D convolutional neural network using cosine similarity.
- **Hierarchical Attention Networks (HAN) [42]**: A hierarchical structure model consisting of two layers of RNN with attention mechanism using cosine similarity.
- **Siamese-BERT (S-BERT) [10]**: This is a siamese matching model with BERT. To handle long document inputs, S-BERT only uses the first 512 tokens of each document. It uses a pre-trained BERT model which is fine-tuned during training and cosine similarity.
- **Siamese Multi-depth Transformed based Hierarchical Encoder (SMITH) [41]**: The current SOTA model for long-form document matching tasks. The model uses a transformer-based hierarchical encoder to obtain sentence-level and document-level representations for long-form documents.
- **Siamese Longformer (S-LONG) [4]**: A transformer-based model for long sequences with an attention mechanism that scales linearly with sequence length. It takes a maximum input length of 4,096 tokens.

## 4.2 RQ1: Evaluation on Long Document Matching

We evaluate the performance of our model on the downstream task of long document matching. The document matching problem can be treated as a classification task.<sup>7</sup> During evaluation, the model is given a pair of documents and the task is to predict whether the two input documents are similar or not. The training is done using the model architecture described in Figure 2. However, during testing, we do not split the document into sections and use only the Data Encoder module to get a single representation for the query and the target documents, respectively. We then use cosine distance to compute the similarity between these representations. If the document level similarity score is above a certain threshold, the documents are considered to be relevant; otherwise, they are considered to be irrelevant. We report precision (P), recall (R), F1-score (F1), and Accuracy (Acc) for this task in Table 2. The utility of the model would depend on the specific use-case, but we observe that CoLDE outperforms all the methods and achieves the best F1-scores and Accuracy on the document matching task on all three datasets. We attribute this superior performance to our new mechanism of exploiting the structure of long documents in conjunction with the contrastive learning framework.

## 4.3 RQ2: Case Study on Model Interpretation

To demonstrate the interpretability provided by CoLDE, we select three documents,  $D_1$ ,  $D_2$ , and  $D_3$  with the following properties.  $D_1$  cites  $D_2$ ; hence, they are similar.  $D_3$  is not cited by either  $D_1$  or  $D_2$ ; hence,  $D_3$  is different from the other two documents. Figure 4 shows the three levels of interpretability (in terms of similarity scores) offered by CoLDE using the architecture described in Figure 2. Qualitative analysis using the text from these documents is shown in Figure 5.

- **Level 1-Document Level Similarity**: We obtain the high-level similarity score using CoLDE. Figure 4(a) illustrates the level of similarity between the three documents  $D_1$ ,  $D_2$ , and  $D_3$ . A document is most similar to itself. Documents  $D_1$  and  $D_2$  have a higher similarity score when compared to document  $D_3$ . Qualitatively,  $D_1$  and  $D_2$  are about “automatic segmentation of speech” whereas  $D_3$  is about “question answering” (see Figure 5).

<sup>7</sup>We use classification metrics instead of information-retrieval metrics due to the limitations of the dataset which has very few positive samples (2–3) for every document.

Table 2. Comparison of Different Methods on Various Long-document Datasets: (i) AAN, (ii) WIKI, and (iii) PAT

| Model                  | AAN          |              |              |              | WIKI         |              |              |              | PAT          |              |              |              |
|------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
|                        | P            | R            | F1           | Acc          | P            | R            | F1           | Acc          | P            | R            | F1           | Acc          |
| DSSM                   | 0.676        | 0.716        | <u>0.711</u> | 0.686        | 0.736        | 0.830        | 0.780        | <u>0.788</u> | 0.538        | 0.538        | 0.538        | 0.521        |
| ARC-I                  | 0.647        | 0.683        | 0.665        | 0.655        | 0.799        | 0.666        | 0.727        | 0.764        | 0.782        | 0.695        | 0.736        | 0.708        |
| HAN                    | 0.641        | <u>0.743</u> | 0.688        | 0.664        | 0.676        | <u>0.954</u> | <u>0.791</u> | 0.748        | 0.521        | 0.521        | 0.521        | 0.538        |
| S-BERT                 | 0.665        | 0.666        | 0.665        | 0.566        | 0.665        | 0.667        | 0.666        | 0.499        | 0.536        | 0.495        | 0.515        | 0.520        |
| SMITH                  | <b>0.732</b> | 0.603        | 0.662        | 0.691        | 0.640        | <b>0.980</b> | 0.774        | 0.714        | 0.524        | 0.485        | 0.502        | 0.509        |
| S-LONG                 | 0.633        | 0.723        | 0.675        | 0.699        | 0.794        | 0.733        | 0.762        | 0.752        | <b>0.848</b> | 0.533        | 0.655        | 0.521        |
| CoLDE                  | <u>0.719</u> | <b>0.756</b> | <b>0.734</b> | <b>0.742</b> | <u>0.860</u> | 0.767        | <b>0.809</b> | <b>0.801</b> | <u>0.829</u> | <b>0.849</b> | <b>0.839</b> | <b>0.833</b> |
| CoLDE <sub>-Aug</sub>  | 0.623        | 0.629        | 0.623        | 0.616        | 0.606        | 0.580        | 0.592        | 0.606        | 0.778        | 0.809        | 0.790        | 0.757        |
| CoLDE <sub>-LSTM</sub> | 0.718        | 0.701        | 0.709        | <u>0.706</u> | 0.780        | 0.754        | 0.766        | 0.786        | 0.795        | <u>0.843</u> | <u>0.818</u> | <u>0.792</u> |
| CoLDE <sub>-CL</sub>   | 0.599        | 0.499        | 0.545        | 0.595        | <b>0.943</b> | 0.510        | 0.658        | 0.500        | 0.614        | 0.584        | 0.595        | 0.586        |

Different metrics such as Precision (P), Recall (R), F1-Score (F1), and Accuracy (Acc) for document matching task are reported. The best result is highlighted in boldface and the second best is underlined. CoLDE's improvements over other methods are statistically significant with  $p < 0.05$  as measured by Student's t-test (whenever CoLDE's performance is the best). We also present the results for the ablated variants of CoLDE.

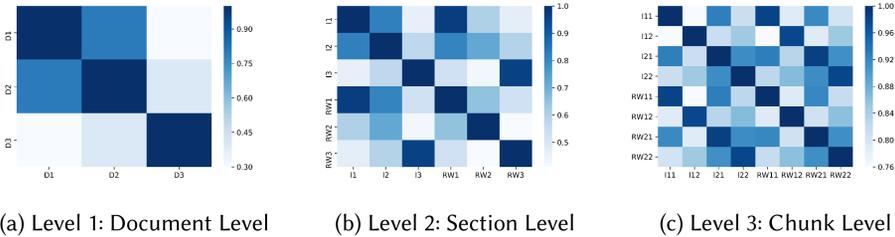


Fig. 4. Three levels of interpretability provided by CoLDE: (a) Document Level Similarity, (b) Section Level Similarity, and (c) Chunk Level Similarity. We consider three research articles:  $D_1$ ,  $D_2$ , and  $D_3$ .  $D_1$  and  $D_2$  are similar to each other. Document  $D_3$  is different from the other two.  $I_1 - I_3$  correspond to the “Introduction”;  $RW_1 - RW_3$  correspond to the “Related Work” sections of documents  $D_1 - D_3$ , respectively. Qualitative analysis for Figure 4(c) is shown in Figure 5.

- **Level 2-Section Level Similarity:** While document-level similarity is certainly useful to get a global overview, it does not demonstrate the reason behind the (dis)similarity. Therefore, CoLDE goes a level deeper by providing section level similarity. To explain this outcome, we split each document  $\{D_k\}_1^3$  into two sections  $\{(I_k, RW_k)\}_1^3$ , where  $I_k$  and  $RW_k$  indicate the Introduction and Related Work of  $k$ th document, respectively. Figure 4(b) shows how these are (dis)similar. First, sections from the same document have a high similarity score ( $I_1$  and  $RW_1$ ;  $I_2$  and  $RW_2$ ;  $I_3$  and  $RW_3$ ). Second, since documents  $D_1$  and  $D_2$  are similar, their introductions  $I_1$  and  $I_2$  are more similar compared to  $I_3$  from  $D_3$ . This is also true for sections  $RW_1$  and  $RW_2$ , which are more similar to each other compared to  $RW_3$ . Figure 5 shows some of the content from  $I_1$ ,  $I_2$ , and  $RW_2$ . The introductions  $I_1$  and  $I_2$  discuss “topic/text segmentation in dialogues and meetings”, and  $RW_2$  presents the related work in the area.
- **Level 3-Chunk Level Similarity:** Besides providing the section-level similarity, CoLDE also explains which segment (or chunk) of a section is important. To illustrate this outcome, we divide  $(I_k, RW_k)$  into  $j$  chunks, which is denoted by  $(I_{jk}, RW_{jk})$ . In our experiments, we set  $j = 2$ ; in other words, each document is divided into four chunks. Figure 4(c) shows some interesting traits of chunk-level interpretability between documents  $D_1$  and  $D_2$ . The introduction chunk  $I_{11}$  from document  $D_1$  is similar to the first introduction chunk  $I_{21}$  from

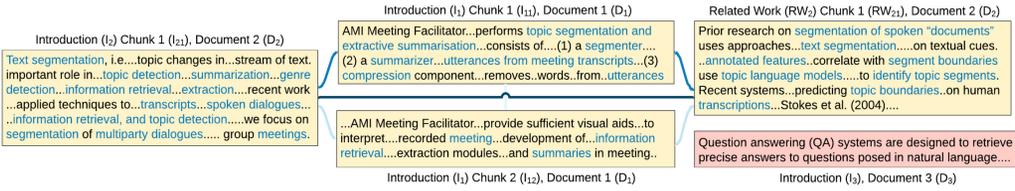


Fig. 5. Qualitative Analysis: Documents  $D_1$  and  $D_2$  are similar. The introduction chunk  $I_{11}$  from  $D_1$  is about “topic segmentation and extractive summarisation”. The introduction chunk  $I_{21}$  from  $D_2$  is about “text segmentation”; the related work chunk  $RW_{21}$  from  $D_2$  presents prior work in the area of “text segmentation”. The introduction chunk  $I_{12}$  from  $D_1$  summarizes the author contributions. Document  $D_3$  is on question answering and is different from the other two documents. Darker colored lines indicate higher similarity.

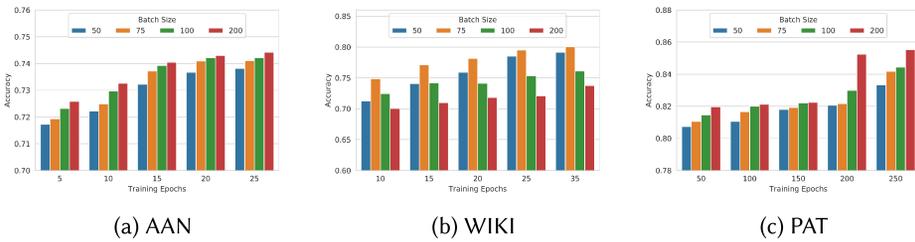


Fig. 6. Effect of batch size on CoLDE’s accuracy for document matching task on the three datasets.

$D_2$ . In Figure 5, we see that  $I_{11}$  broadly discusses “the system components that perform topic segmentation and extractive summarisation”, and  $I_{21}$  discusses “text segmentation and its applications”. Additionally, the related work chunk  $RW_{21}$  from  $D_2$  is very similar to both  $I_{11}$  and  $I_{21}$  from documents  $D_1$  and  $D_2$ , respectively. Qualitatively,  $RW_{21}$  presents prior work in the area of “text segmentation”. However, we observe that the second introduction chunk  $I_{12}$  from document  $D_1$  has very low similarity score with the above chunks ( $I_{11}$ ,  $I_{21}$ , and  $RW_{21}$ ) and interestingly,  $I_{12}$  is about “visual aids for interpretability and a summary of the author contributions”.

CoLDE, therefore, provides fine-grained *inter-document* and *intra-document* (dis)similarity scores at three different levels which makes the model more interpretable.

#### 4.4 RQ3: Robustness Analysis of CoLDE

**4.4.1 Effect of Batch Size on CoLDE.** Recent works on contrastive learning in the computer vision domain [6, 24, 31] have shown that the discriminating ability of the model increases with increase in batch size since there will be an increase in the number of informative negative samples. We empirically validate this hypothesis and demonstrate that this holds true for the long document matching task. Figure 6 shows the impact of batch size on model accuracy when it is trained for the same number of epochs on three datasets: (i) AAN, (ii) WIKI, and (iii) PAT. For the same number of epochs, larger batch-size performs better for AAN and PAT datasets. However, because of the presence of hard as well as easy negatives in the WIKI dataset, a larger batch size will have many more easy negatives making the task more challenging. Although CoLDE does eventually converge to the optimal accuracy of  $\sim 80\%$  for all batch sizes on the WIKI dataset, it takes longer for bigger batch sizes than for smaller ones. Irrespective of the dataset, the best performance of CoLDE steadily increases with the number of training epochs.

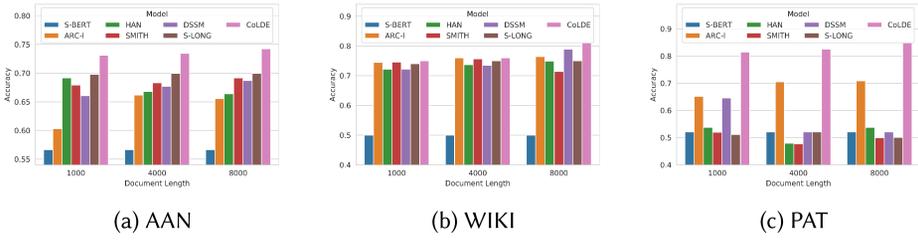


Fig. 7. Effect of document length on CoLDE's accuracy for document matching task on the three datasets.

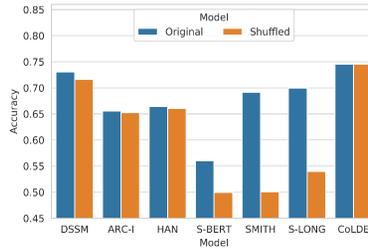


Fig. 8. Effect of text perturbation on AAN dataset for document matching task.

**4.4.2 Effect of Document Length on CoLDE.** In order to study CoLDE's robustness to varying document lengths, we limit the maximum number of tokens in each section during training and observe the final test accuracy on the document matching task. The model is given the entire long document during evaluation. Figure 7 shows the model accuracy for the maximum lengths of 1,000, 4,000 and 8,000 tokens. We compare CoLDE's performance with other methods for similar document lengths. We observe that CoLDE outperforms the baselines and its performance steadily increases with increase in document length, emphasizing the need for explicit long document matching models.

**4.4.3 Effect of Text Perturbation on CoLDE.** We study the effect of text perturbation on the model performance for the AAN dataset. We randomly shuffle the sections before applying different models and measure their accuracy on the document comparison task. Figure 8 shows the model performance for all the baseline methods. We observe a significant drop in the model performance for BERT-based models (S-BERT, S-LONG, and SMITH), whereas CoLDE's performance remains unchanged. There is no significant change in the accuracy for DSSM, ARC-I, and HAN. We observe that, although other transformer-based models are sensitive to text perturbation, CoLDE is robust. The proposed unique positional embeddings used to enhance the tokens during the data augmentation phase are long document aware, and hence, even after shuffling the text sections, they retain the original long document structure information, thus preserving the accuracy on the downstream task.

## 4.5 RQ4: Ablation Study of CoLDE

We study the effect of different components of CoLDE during training by conducting ablation study experiments. The results of the ablation study on the document matching task are shown in Table 2.

**4.5.1 Effect of Data Augmentation.** We remove the Data Augmentation module to study its impact on the overall model performance. This model variant is termed as CoLDE<sub>-Aug</sub>. We do not split

a long document into different sections. We also do not add the proposed positional embeddings to the input document. Instead, the entire long document is split into chunks of 512 tokens and given to the data encoder module. The objective function used is supervised contrastive loss. We observe that there is a significant drop in the performance for document matching on all three datasets. The data augmentation module helps in learning better document representations by bringing the representations of sections within a document closer. Moreover, sections of documents belonging to the same class label are also brought closer to each other in the latent embedding space when used in conjunction with the supervised contrastive loss. In its absence, we lose out on learning meaningful document representations, in addition to section level similarity.

**4.5.2 Effect of LSTM and Multi-headed Chunkwise Attention.** BERT can only handle 512 tokens at a time. To overcome this limitation and get section representations, we use a Bi-LSTM layer for aggregation. We study its effectiveness by removing only the Bi-LSTM and the multi-headed chunkwise attention layer from the data encoder module. We term this variant as CoLDE<sub>-LSTM</sub>. The section representation is the “average” of the BERT output representations for all chunks of 512 tokens within a section. This is then given to the projection layer for dimensionality reduction. The objective function used is supervised contrastive loss. We observe that there is a performance drop on the document matching task when not using the Bi-LSTM layer for aggregation. *Bi-LSTM layer effectively captures the sequential nature of the chunks within a section which is lost when computing the average of the BERT output representations for text chunks.*

**4.5.3 Effect of Using Contrastive Loss.** We conduct an ablation study by replacing the supervised contrastive loss with cosine similarity. Since, we do not require documents to be split into sections for computing the cosine similarity between them, removing the contrastive loss function entails disabling the data augmentation module as well. This variant is termed as CoLDE<sub>-CL</sub>. The input documents are split into multiple chunks of 512 tokens and given to the data encoder module. The final representations from this module are used to compute the cosine similarity. From Table 2, we observe that there is a significant drop in performance when using cosine similarity as opposed to using the supervised contrastive loss as our objective function. Supervised contrastive loss helps in providing fine-grained similarity scores and ensures that the latent representations of sections within the same document are closer to each other. Additionally, representations of documents belonging to the same class are also closer in the latent embedding space.

**4.5.4 Effect of Number of Input Sections.** We study the effect of number of sections on the performance of CoLDE by splitting the source and the target documents into sections ranging from 2 to 5, before being given as input to the model. CoLDE was robust to these changes and the performance was consistent with the ones shown in Table 2 for different input sections for all datasets.

## 5 CONCLUSION

In this work, we introduced a contrastive learning framework CoLDE for long-form document matching. CoLDE leverages the long document structure in conjunction with supervised contrastive loss to provide fine-grained similarity scores within and across documents at three different levels: (i) Document Level, (ii) Section Level, and (iii) Chunk Level. The overall framework consists of three primary components: (i) Data Augmentation, (ii) Data Encoder, and (iii) Contrastive Loss Function. We showed that CoLDE outperforms the existing state-of-the-art models for long document matching and presented a detailed case study for model interpretation. We also empirically demonstrated CoLDE’s robustness to batch size, document length, and text perturbation compared to other methods. Moreover, we conducted an ablation study to examine the

effectiveness of various components of CoLDE. In summary, we demonstrated that CoLDE's contrastive learning approach is effective for interpretable long-form document matching.

## REFERENCES

- [1] Ashutosh Adhikari, Achyudh Ram, Raphael Tang, and Jimmy Lin. 2019. Docbert: Bert for document classification. arXiv:1904.08398. Retrieved from <https://arxiv.org/abs/1904.08398>.
- [2] David Alvarez-Melis and Tommi Jaakkola. 2017. A causal framework for explaining the predictions of black-box sequence-to-sequence models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 412–421.
- [3] Hadi Amiri, Philip Resnik, Jordan Boyd-Graber, and Hal Daumé III. 2016. Learning text pair similarity with context-sensitive autoencoders. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*. 1882–1892.
- [4] Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. arXiv:2004.05150. Retrieved from <https://arxiv.org/abs/2004.05150>.
- [5] Minmin Chen. 2017. Efficient vector representation for documents through corruption. In *Proceedings of the ICLR (Poster)*.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *Proceedings of the International Conference on Machine Learning*. 1597–1607.
- [7] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. arXiv:1904.10509. Retrieved from <https://arxiv.org/abs/1904.10509>.
- [8] Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, David Belanger, Lucy Colwell, and Adrian Weller. 2020. Re-thinking attention with performers. In *Proceedings of the International Conference on Learning Representations*.
- [9] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 2978–2988.
- [10] 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*. 4171–4186.
- [11] Ming Ding, Chang Zhou, Hongxia Yang, and Jie Tang. 2020. CogLTX: Applying BERT to Long Texts. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.). Vol. 33. Curran Associates, Inc., 12792–12804. <https://proceedings.neurips.cc/paper/2020/file/96671501524948bc3937b4b30d0e57b9Paper.pdf>.
- [12] Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 6894–6910. <https://doi.org/10.18653/v1/2021.emnlp-main.552>.
- [13] Reza Ghaeini, Xiaoli Fern, and Prasad Tadepalli. 2018. Interpreting recurrent and attention-based neural models: A case study on natural language inference. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 4952–4957.
- [14] John Giorgi, Osvald Nitski, Bo Wang, and Gary Bader. 2021. DeCLUTR: Deep contrastive learning for unsupervised textual representations. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*. 879–895.
- [15] Alex Graves and Jürgen Schmidhuber. 2005. Framewise phoneme classification with bidirectional LSTM networks. In *Proceedings of the 2005 IEEE International Joint Conference on Neural Networks*. IEEE, 2047–2052.
- [16] Jiafeng Guo, Yixing Fan, Qingyao Ai, and W. Bruce Croft. 2016. A deep relevance matching model for ad-hoc retrieval. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*. 55–64.
- [17] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. 2019. Axial attention in multidimensional transformers. arXiv:1912.12180. Retrieved from <https://arxiv.org/abs/1912.12180>.
- [18] Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional Neural Network Architectures for Matching Natural Language Sentences. In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (Eds.). Vol. 27. Curran Associates, Inc. 2042–2050. <https://proceedings.neurips.cc/paper/2014/file/b9d487a30398d42ecff55c228ed5652b-Paper.pdf>.
- [19] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *Proceedings of the 22nd ACM International Conference on Information and Knowledge Management*. 2333–2338.

- [20] Sarthak Jain and Byron C. Wallace. 2019. Attention is not explanation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 3543–3556.
- [21] Jyun-Yu Jiang, Mingyang Zhang, Cheng Li, Michael Bendersky, Nadav Golbandi, and Marc Najork. 2019. Semantic text matching for long-form documents. In *Proceedings of the World Wide Web Conference*. 795–806.
- [22] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. 2020. Transformers are RNNs: Fast autoregressive transformers with linear attention. In *Proceedings of the International Conference on Machine Learning*. 5156–5165.
- [23] Tom Kenter and Maarten De Rijke. 2015. Short text similarity with word embeddings. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*. 1411–1420.
- [24] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. Supervised Contrastive Learning. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.). Vol. 33. Curran Associates, Inc., 18661–18673. <https://proceedings.neurips.cc/paper/2020/file/d89a66c7c80a29b1bdbab0f2a1a94af8-Paper.pdf>.
- [25] Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2019. Reformer: The efficient transformer. In *Proceedings of the International Conference on Learning Representations*.
- [26] Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proceedings of the International Conference on Machine Learning*. 1188–1196.
- [27] Yang Liu and Mirella Lapata. 2019. Hierarchical transformers for multi-document summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Florence, Italy, 5070–5081. DOI: <https://doi.org/10.18653/v1/P19-1500>
- [28] Dongsheng Luo, Wei Cheng, Jingchao Ni, Wenchao Yu, Xuchao Zhang, Bo Zong, Yanchi Liu, Zhengzhang Chen, Dongjin Song, Haifeng Chen, and Xiang Zhang. 2021. Unsupervised document embedding via contrastive augmentation. arXiv:2103.14542. Retrieved from <https://arxiv.org/abs/2103.14542>.
- [29] Bhaskar Mitra, Fernando Diaz, and Nick Craswell. 2017. Learning to match using local and distributed representations of text for web search. In *Proceedings of the 26th International Conference on World Wide Web*. 1291–1299.
- [30] Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. Text matching as image recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [31] Taesung Park, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. 2020. Contrastive learning for unpaired image-to-image translation. In *Proceedings of the European Conference on Computer Vision*. Springer, 319–345.
- [32] Jiezhong Qiu, Hao Ma, Omer Levy, Wen-tau Yih, Sinong Wang, and Jie Tang. 2020. Blockwise self-attention for long document understanding. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*. 2555–2565.
- [33] Dragomir R. Radev, Pradeep Muthukrishnan, Vahed Qazvinian, and Amjad Abu-Jbara. 2013. The ACL anthology network corpus. *Language Resources and Evaluation* 47, 4 (Jan. 2013), 909–944. <https://doi.org/10.1007/s10579-012-9211-2>
- [34] Jack W. Rae, Anna Potapenko, Siddhant M. Jayakumar, Chloe Hillier, and Timothy P. Lillicrap. 2019. Compressive transformers for long-range sequence modelling. In *Proceedings of the International Conference on Learning Representations*.
- [35] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*. 3973–3983.
- [36] Andrew Slavin Ross, Michael C. Hughes, and Finale Doshi-Velez. 2017. Right for the right reasons: Training differentiable models by constraining their explanations. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. 2662–2670.
- [37] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *Proceedings of the International Conference on Machine Learning*. 3319–3328.
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 6000–6010.
- [39] Sarah Wiegrefe and Yuval Pinter. 2019. Attention is not not explanation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*. 11–20.
- [40] Liu Yang, Qingyao Ai, Jiafeng Guo, and W. Bruce Croft. 2016. aNMM: Ranking short answer texts with attention-based neural matching model. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. 287–296.
- [41] Liu Yang, Mingyang Zhang, Cheng Li, Michael Bendersky, and Marc Najork. 2020. Beyond 512 tokens: Siamese multi-depth transformer-based hierarchical encoder for long-form document matching. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*. 1725–1734.

- [42] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, San Diego, California, 1480–1489. DOI : <https://doi.org/10.18653/v1/N16-1174>
- [43] Jianfei Yu, Minghui Qiu, Jing Jiang, Jun Huang, Shuangyong Song, Wei Chu, and Haiqing Chen. 2018. Modelling domain relationships for transfer learning on retrieval-based question answering systems in e-commerce. In *Proceedings of the 11th ACM International Conference on Web Search and Data Mining*. 682–690.
- [44] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. Big Bird: Transformers for Longer Sequences. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.). Vol. 33. Curran Associates, Inc., 17283–17297. <https://proceedings.neurips.cc/paper/2020/file/c8512d142a2d849725f31a9a7a361ab9-Paper.pdf>.
- [45] Hang Zhang, Yeyun Gong, Yelong Shen, Weisheng Li, Jiancheng Lv, Nan Duan, and Weizhu Chen. 2021. Pooling-former: Long document modeling with pooling attention. In *Proceedings of the 38th International Conference on Machine Learning*, Marina Meila and Tong Zhang (Eds.). PMLR, 12437–12446. Retrieved from <https://proceedings.mlr.press/v139/zhang21h.html>.

Received 30 October 2021; accepted 8 May 2022