NewsEmbed: Modeling News through Pre-trained Document Representations

Jialu Liu*, Tianqi Liu*, Cong Yu
Google Research
{jialu,tianqiliu,congyu}@google.com

ABSTRACT

Effectively modeling text-rich fresh content such as news articles at document-level is a challenging problem. To ensure a content-based model generalize well to a broad range of applications, it is critical to have a training dataset that is large beyond the scale of human labels while achieving desired quality. In this work, we address those two challenges by proposing a novel approach to mine semantically-relevant fresh documents, and their topic labels, with little human supervision. Meanwhile, we design a multitask model called NewsEmbed that alternatively trains a contrastive learning with a multi-label classification to derive a universal document encoder. We show that the proposed approach can provide billions of high quality organic training examples and can be naturally extended to multilingual setting where texts in different languages are encoded in the same semantic space. We experimentally demonstrate NewsEmbed's competitive performance across multiple natural language understanding tasks, both supervised and unsupervised.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning algorithms**; **Natural language processing**.

KEYWORDS

Deep learning; Data mining; Cross-lingual; Representation learning; Contrastive learning; Weak supervision; News

ACM Reference Format:

Jialu Liu*, Tianqi Liu*, Cong Yu. 2021. NewsEmbed: Modeling News through Pre-trained Document Representations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21)*, August 14–18, 2021, Virtual Event, Singapore. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3447548.3467392>

1 INTRODUCTION

Effectively modeling text-rich news content on the public web can enhance a broad set of applications in information retrieval and recommendation, and is especially crucial for long standing news tasks such as story clustering and topic classification. It is,

however, a challenging problem due to out-of-vocabulary concepts and evolving relations between entities.

A typical way to represent general text data is projecting it into a sparse or dense vector space. If the subsequent task is unsupervised, e.g., clustering, one can compute similarity between texts by applying predefined functions such as inner product. If the subsequent task requires additional training, e.g., classification, an additional task-specific fine-tuning can be applied. Our work follows this paradigm and specifically focuses on text-rich news content. We demonstrate the performance of our embedding model through comprehensive evaluations in both unsupervised and supervised setups.

Modeling text data has a rich history. Classical approaches involve statistical modelling of words to represent the whole document, such as bag-of-words [31], term frequency-inverse document frequency [32], and topic modeling [4]. There is no notion of similarity between words in those systems and texts are represented by sparse word vectors. To address this, models such as Word2Vec [40], GloVe [47], and FastText [5] have been developed to capture the co-occurrences of characters or words where there can be dense vector representations. With the development of deep learning, convolutional neural networks [62], recurrent neural networks [33, 64], and transformer [55] are used to represent contextual information [39, 48]. More recently, with the development of self-supervised, unsupervised, and transfer learning, pre-training techniques such as BERT [19], T5 [49], and GPT3 [7] are proposed to leverage the free text on the web and achieve the state-of-the-art performance on various natural language processing (NLP) tasks without the need for huge amount of labelled data.

Although above deep learning approaches show great advantage in learning *token-level* contextual representations from self-supervision, they still require a good amount of labelled fine-tuning data to produce high quality *document-level* embeddings, which are essential for news tasks such as clustering. Naive approaches have been shown to be even less competitive to averaged GloVe [47] in terms of semantic similarity [53]. The high cost of annotating supervised training data impedes the development of high quality models in most research or even industry settings in that direction. At *sentence-level*, several studies have been leveraging existing large scale datasets. Universal Sentence Encoder (USE) [9] leverages open Web data sources such as Wikipedia and uses unsupervised learning augmented with additional training on Stanford Natural Language Inference (SNLI) corpus [6]. Sentence-BERT (SBERT) [53] finetunes siamese and triplet network structures on SNLI and the MultiGenre NLI dataset (MNLI) [57] to obtain competitive performance on Semantic Textual Similarity (STS) tasks [8].

* Contribute to this work equally. Correspondence to: Jialu Liu.

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

Both USE and SBERT are developed for sentence-level representation and it is questionable whether they can capture semantics given longer context. Moreover, while both show great potential by transfer learning from NLI to STS, it is still challenging to adapt them to news due to the many out-of-vocabulary concepts, entities, and events. Recent studies show that domain specific training data can usually improve tasks through domain adaptation [3, 28]. It is also important for a news content representation model to be cross-lingual so that semantically-similar texts in different languages can be aligned in the same vector space. Recent works [14, 60] show that the knowledge can be transferred from high resource languages to low resource languages effectively by forcing alignments between different languages.

Inspired by those prior works, our goal is to leverage large scale open data in a smart and multi-lingual way to construct a document-level representation model specifically for the news domain, as part of our mission to understand news content better and improve a broad range of news applications.

We propose a system called NewsEmbed that trains an encoder from weakly supervised data. We first introduce an approach to obtain billions of cross-lingual document *triplets* (anchor, positive, negative) where anchor and positive documents are semantically close, while anchor and negative documents are related but less similar. In addition, we describe an effective heuristic to mine large-scale document-topic associations with high precision. Regarding the modeling, we adopt triplet neural structure and co-train a BERT-initialized encoder from the previously mentioned two datasets, with InfoNCE loss [45] as the contrastive learning objective and BCE loss [43] as the classification objective respectively. Our contributions of this work are three-fold:

- We propose a scalable approach to obtain multilingual weak supervision data with the purpose of pre-training document representations in news domain.
- We develop an effective neural architecture to co-train contrastive learning with multi-label classification, which demonstrates strong performance in later unsupervised and supervised applications, some of which are even out-of-domain.
- As a side product, we show that a BERT model pre-trained from news corpus alone performs competitive on generic cross-lingual benchmarks [30].

The rest of the paper is organized as follows. Section 2 will introduce our pipeline of collecting pre-training and weakly supervised datasets. Section 3 will cover the model and learning approach. In Section 4, we will compare our approach with other works and evaluate them with public benchmark datasets. Ablation studies will be provided to analyze alternative designs. In Section 5, related works will be discussed, followed by the conclusion in Section 6.

2 DATA COLLECTION

NewsEmbed trains a transformer encoder initialized from a BERT pre-training checkpoint. In this section, we first introduce the pre-training dataset and then describe in details how to mine weak supervision for the next stage of training.

Figure 1: Example of a document triplet.

2.1 Corpus for Pre-training

We have two types of data: monolingual documents and bilingual translation pairs fed to a BERT masked language model.

2.1.1 Monolingual Documents. We collect ~1.5B newsy articles based on an in-house crawler that extracts structured information including title, body, author, byline date and anchors. Documents without author or byline date are filtered out. To ensure the document is text-rich, we discard any document with fewer than 100 words. Furthermore, we apply similar data cleaning heuristics in [49] to deduplicate the dataset and remove low-quality content.

It is worth noting that there exist open sourced, easy-to-use crawlers, such as news-please¹, that can parse and extract structured news articles. They recursively follow internal hyperlinks and RSS feeds to fetch both recent and also old, archived documents.

2.1.2 Bilingual Translation Pairs. To better align the cross-lingual embedding space and transfer knowledge from high resource languages to low resource languages, a translation corpus is constructed from the web pages using a bitext mining system described in [23]. This corpus contains ~15B translation sentence pairs.

2.2 Corpus for Weak Supervision

Recall that the goal of this work is to train a *document* encoder such that both in-domain supervised and unsupervised learning can benefit from it. Typical pre-trained models, such as BERT variants, aren't effective without domain adaption including the modeling of both long context and different components in a document. This subsection provides details how we fill the gap by mining knowledge from the in-domain corpus.

2.2.1 Document Triplets. Inspired by contrastive learning initiated in computer vision [36], we believe text representations can be learned by enforcing similar elements to be equal and dissimilar elements to be different. When it comes to learning document representation, it naturally requires document triplets. Without loss of generality, we denote a triplet as (anchor, positive, negative),

¹<https://github.com/fhamborg/news-please>

Figure 2: Data pipeline to collect document triplets for contrastive learning.

or (a, p, n) in short, where $\langle a, p \rangle$ document pair is more semantically similar than $\langle a, n \rangle$. An example is illustrated in Figure 1.

Due to the special property of news, it is feasible to mine such document triplets based on the following two observations:

OBSERVATION 1. (Redundancy) *News events or stories, especially the most popular ones, are likely to be reported and discussed from multiple publishers.*

OBSERVATION 2. (Time-Sensitivity) *Lifetime of a news event or story is usually short while long running ones will evolve over time.*

The first observation ensures the abundance of organic $\langle a, p \rangle$ pairs if we have an effective approach to find these relevant documents. Since these texts are published by different publishers, they likely include paraphrasing of the same facts but expression of different opinions. From this perspective, by enforcing similarity between these documents, we are learning a semantic space where texts are close if they are targeting the same event or story. Also note that this source is different from typical text augmentation tricks, such as back translation and sentence reordering, because the paraphrasing is more natural and thus the quality is better.

The second, time-sensitivity, observation provides an effective signal to filter out noisy document triplets with high precision. To achieve this, assuming there exist some orthogonal signals to generate candidate triplets, we can apply aggressive filtering to ensure $\langle a, p \rangle$ is close in byline date while $\langle a, n \rangle$ is not. We concede that such cleaning approach causes coverage loss, but empirically it leads to model improvement, implying that the quality of the weak supervision is more important than the quantity.

The overall process for collecting document triplets is illustrated in Figure 2. We first encode each document using some auxiliary embedding signals such as entity and image embeddings. ScaNN [26] is then used to extract approximate nearest neighbors for each document w.r.t these auxiliary embeddings. Based on the nearest neighbors, we generate (a, p, n) document triplets and apply the aforementioned byline date to improve the data quality. Lastly, these triplets are augmented to model both short- and long-context texts in the final dataset. In the rest of this subsection, we discuss these steps in details.

Auxiliary embeddings for retrieving related documents. In Figure 2, the first and second steps are to extract semantically-related documents from the pre-training corpus in large scale. One can resort to various document embedding algorithms proposed in the literature. And there is a high level of tolerance to the noise in these embedding methods because we have a denoising step later.

Concretely speaking, in NewsEmbed we compute three auxiliary embeddings for each document and use each of them to fetch related documents independently, including:

- **Entity embedding:** Topical knowledge base entities are identified and their embeddings [61] are averaged to represent a document.
- **Image embedding:** Thumbnail embedding is derived using wavelet hashing [56].
- **Text embedding:** Token embeddings in the title and keywords are averaged to represent a document [10].

In particular, the former two embeddings are cross-lingual as their embedding sources are language-agnostic.

Triplet generation and filtering. After applying ScaNN [26] to search nearest neighbors for each document in each auxiliary embedding space, we obtain large quantities of document pairs. The next step is transforming them to triplets by finding high quality $\langle a, p \rangle$ and $\langle a, n \rangle$ pairs. To achieve these, we merge all top- K nearest neighbors for a document from all the auxiliary embedding types, and check byline date in the neighbor documents to decide whether the pair should be positive or negative. Empirically we set one day as maximum time difference for $\langle a, p \rangle$ and one year as minimum time difference for $\langle a, n \rangle$. Since a and n are close in at least one of the auxiliary embedding space, we consider them as hard negatives so that later contrastive learning could be more effective.

Nevertheless, there is a shortcoming of heavily relying on byline date to denoise the triplets. We notice many of the false negatives from the above approach are similar in nature, i.e., document content is evergreen. For example, topics like “minimum wage” are mentioned frequently over time and thus are not time-sensitive as described in Observation 2. In this regard, we train an auxiliary *byline date prediction model* specifically to predict byline date based on document title and body, with the motivation that the publication time of evergreen content is more difficult to predict compared to time-sensitive news events and stories. A logistic regression model is then deployed to remove false negatives by considering predicted date distribution together with the embedding similarities mentioned previously. Since the number of features used here is very small, only a few hundred training examples are needed.

Augmented triplets. After mining document triplets, we augment them in various ways as shown in Table 1 before feeding the training data into the triplet neural network. The rationale of the augmentation is to make the model learn a better embedding on both short context (i.e., title or anchor text) and long context (i.e., body alone or title concatenated with body). It is worth noting that augmented positive does not necessarily come from positive document, and we intentionally ensure that the augmented positive and negative are always of the same type, to prevent biases of content length during the optimization of contrastive objective.

To encourage better cross-lingual alignment, for augmented triplets that are entirely not in English, we translate positive and

Augmented Anchor	Augmented Positive	Augmented Negative
Title (a)	Body (a)	Body (n)
Anchor Text* (a)	Title + Body (a)	Title + Body (n)
Title (a)	Title + Body (p)	Title + Body (n)
Title + Body (a)	Title + Body (p)	Title + Body (n)

*We consider sometimes the text of incoming anchor summarizes the document. To denoise, one can add a similarity check using the auxiliary text embedding.

Table 1: Augmented triplets from document triplet (a, p, n). Symbol + denotes segment concatenation.

negative into English. This helps the model to transfer knowledge in a multilingual setup.

2.2.2 Document-Topic Associations. Document triplets capture semantic similarity at the granularity of events or stories. To that extent, it lacks knowledge of coarse-grained topics to teach an encoder whether a document belongs to the category of “investment”, “lifestyle” or some others. Adding this missing piece in the dataset can further enhance the representation power of the embedding.

To mine these document-topic associations in large-scale, we need to rely on an observation described as below:

OBSERVATION 3. (Aggregation) *Publishers choose to aggregate documents with shared topics for ease of information consumption.*

Typical examples of these “hubs” include sub-directories and RSS feeds of a site as shown in Figure 3. However, not all “hubs” on the web are useful for our use case, e.g., many are based on author or publishing time. Realizing this, we apply entity linking to hub page titles and derive a list of the most popular candidates. Then we select a subset that can be considered as coarse-grained topics and ignore the rest. In addition to regular broad topics, we notice some hub entities can be mapped to journalist types such as “interview” and “opinion”. They are also included in our list to help learn the story styles.

A challenge in processing this dataset is that the mined document-topic associations only provide positive labels. In other words, it is not yet clear whether any unobserved <document, topic> pairs are negative. While different publishers tend to use different aggregation strategies consistently over time, we can treat unobserved <document, topic> pairs as negative if (and only if) that topic appear on the hub page of the same publisher. Additional sampling is applied to ensure positives and negatives are balanced.

3 MODELS

In this section, we present the model learning details in pre-training and weak supervision phases respectively.

3.1 Language Model Pre-training

Following the multilingual pre-training recipe, we train a BERT-style transformer encoder by Masked Language Modeling [19] and Translation Language Modeling [35] on whole word masking [16]. Considering the page limit and novelty of this part, we omit some of the model learning details. Interested readers can refer to existing work for the language model formulation.

One difference from the previous work is that our translation sentence pairs are much shorter than monolingual documents. In this regard, we greedily pack multiple pairs into one sequence and

Figure 3: Left: A sub-directory page about automotive; Right: RSS feed for lifestyle.

reset position ids at each pair’s head. Special treatment on self-attention layer is needed to prevent cross-attention among packed texts. Our solution is to correct the attention matrix by masking out cross-pair entries.

3.2 Multitask Weakly Supervised Learning

Recall that we have two datasets collected as weak supervision: augmented document triplets and document-topic associations. At a high level, we treat them as training sources for two separate tasks with different loss functions, but sharing the text encoder parameters during training.

3.2.1 Contrastive Learning. Let’s first introduce the training objective proposed for document triplets. In the framework of contrastive learning, we are expecting a well-trained encoder captures meaningful similarity in the embedding space such that given an *augmented* document triplet, we can get

$$\text{similarity}(\mathbf{a}, \mathbf{p}) > \text{similarity}(\mathbf{a}, \mathbf{n})$$

where $\mathbf{a}, \mathbf{p}, \mathbf{n}$ represent the semantic embedding of components in an augmented triplet. An effective contrastive loss function, called InfoNCE [45], is used to achieve this goal:

$$\text{InfoNCE}(\mathbf{a}, \mathbf{p}, \mathbf{n}, Z; \tau) = \frac{\exp(\mathbf{a} \cdot \mathbf{p} / \tau)}{\sum_{\mathbf{z} \in Z \cup \{\mathbf{p}, \mathbf{n}\}} \exp(\mathbf{a} \cdot \mathbf{z} / \tau)}, \quad (1)$$

where τ is a temperature hyper-parameter and Z is the set of negative samples in the same training batch of triplets. Intuitively, this loss is the log loss of a $(|Z|+2)$ -way softmax-based classifier. In-batch negatives Z is commonly used in the literature to prevent model overfitting and help future retrieval application if needed [11, 23].

3.2.2 Multi-label Classification. In the document-topic association dataset, each document can be mapped to multiple topics and the distribution is skewed when we consider how many documents each topic is associated with.

In this case, Binary Cross-Entropy (BCE) loss can be utilized to model multi-label setups assuming that given a document, its associations with different topics are independent. Its mathematical

Figure 4: NewsEmbed model structure.

formulation is given below,

$$BCE(d, T^d) = \sum_{t \in T^d} y_t^d \log(p_t^d) + (1 - y_t^d) \log(1 - p_t^d) \quad (2)$$

where d indicates a document and T^d represents the set of topics positively or negatively associated with d . Meanwhile, we use y_t^d to represent label value for a document-topic pair. Accordingly, p_t^d denotes the probability that document d is predicted to be associated with t , computed as

$$p_t^d = \text{sigmoid}(\mathbf{w}_t \cdot \mathbf{d})$$

where $\mathbf{w}_t$ is a learnable parameter for topic t and $\mathbf{d}$ is the document's [CLS] embedding.

3.2.3 Model Structure. The model structure is illustrated in Figure 4. We adopt a neural structure with three towers for modeling triplets. Specifically, a pre-trained transformer serves as the text encoder. On top of that, [CLS] embedding from the top transformer layer is connected with two dense layers, corresponding to the semantic embedding and classification logits as described in previous paragraphs, respectively. Similar to other triplet networks, transformer encoders and the two dense layers share parameters across towers.

As you might notice, BCE loss is defined for single document while InfoNCE loss requires triplet input. In order to effectively co-train the two tasks together, we pack three documents in the multi-label classification dataset so that its input format is consistent with contrastive learning.

One missing details in the figure is the stop-gradient operation connected with the transformer output when the input text is after translation. As introduced at the bottom of Section 2.2.1, part of the augmented triplets are augmented with translation for better language alignment. Considering the translated text could be noisy, applying stop-gradient ensures the encoder won't be affected by possibly poor translation, while still allows the knowledge transfer from high-resource to low-resource languages.

3.2.4 Training. The training for NewsEmbed needs to consider both multilingual and multitask where the data scales vary by language and task. Similar to how multilingual BERT handles language imbalance at pre-training stage, we apply the same exponential smoothing to specify the weight for each (language, task) dataset. Its formulation is described in Appendix B in details.

At each training step, we sample one dataset according to the dataset size and feed a random batch to the triplet network depicted

	BERT (B)	XLM (B)	XLM-R (L)	NewsEmbed
XNLI	65.4	69.1	79.2	76.4
PAWS-X	81.9	80.9	86.4	88.3
UDPOS	71.5	71.3	73.8	76.5
PANX	62.2	61.2	65.4	62.2
XQuAD-F1	64.5	59.8	76.6	74.3
MLQA	61.4	48.5	71.6	70.6
TyDiQA	59.7	43.6	65.1	68.7
Tatoeba	38.7	32.6	57.3	80.1
Average	63.0	58.2	71.8	74.6

Table 2: XTREME benchmarks on test dataset. We use B and L to indicate base and large model sizes.

Model	STS 12	STS 13	STS 14	STS 15	STS 16	Avg.
SBERT	63.8	69.3	72.9	75.2	73.3	71.1
USE	65.6	68.0	71.5	80.8	78.7	72.0
mUSE	68.1	70.4	73.5	81.6	79.7	74.2
Laser	62.3	51.6	67.0	75.4	72.3	65.6
LaBSE	66.8	67.7	69.9	78.8	74.2	72.1
NewsEmbed	68.8	76.1	78.7	80.3	78.2	77.2

Table 3: Spearman rank correlation between embedding similarity and gold labels on STS 12-16 datasets [unsupervised].

in Figure 4. This is also called alternate training in the literature, which turns out to be more effective than mixing all datasets together as analyzed in later experiment.

4 EVALUATION

In this section, we first show performance of our pre-trained BERT model developed from news corpus. Then we evaluate NewsEmbed, our document encoder trained with weak supervision for news, in various down-stream applications. Lastly, we conduct ablation studies to investigate alternative designs. Without specific notation, we assume the evaluated model has the same configuration as BERT-base. Reproduction details are provided in Appendix B.

4.1 Pre-trained Model

As previously introduced in Section 2.1, our pre-training corpus comprises monolingual documents and bilingual translation pairs. In particular, monolingual documents only includes news, which can be treated as a special slice of the web data. It is interesting to study how a model pre-trained from this dataset performs, especially when we compare it on the open-domain tasks with other models trained from the entire web without domain filtering.

We evaluate NewsEmbed with XTREME [30], a comprehensive benchmark for cross-lingual transfer learning on a diverse set of languages and tasks. Results are reported in Table 2. NewsEmbed outperforms other base or slightly larger models, including multilingual BERT [19] and XLM [35]. Surprisingly, as model size of NewsEmbed is the same as BERT-base, it performs highly comparable or even better than XLM-R large model [14]. This implies that modeling news together with translation data is suitable or even better for generic natural language understanding tasks and cross-lingual transfer. We suspect data quality of news corpus together with the adoption of large translation corpus contribute to the good XTREME performance.

Model	News	Image Caption	Forum
SBERT	72.2	86.7	64.0
USE	67.3	86.8	67.5
mUSE	71.7	88.8	69.0
Laser	65.0	76.9	63.9
LaBSE	75.3	76.3	62.9
NewsEmbed	80.9	88.9	70.8

Table 4: Spearman rank correlation between embedding similarity and gold labels on in-domain and out-of-domain slices of STS benchmark dataset with train, dev and test merged [unsupervised].

4.2 Weakly Supervised Model

NewsEmbed is trained to be a generic document encoder for news. Moreover, the encoder is supposed to be robust towards text length since the training data is augmented with both short- and long-context. Bearing this in mind, we evaluate NewsEmbed from three perspectives, including 1) whether a task is in-domain or out-of-domain, 2) whether a task is supervised or unsupervised, and 3) whether the text is short or long. To demonstrate the effectiveness of NewsEmbed, we compare it with the following encoders:

- Sentence-BERT (SBERT) [53]²: uses transformer-based siamese and triplet network structures to derive semantic text representation. The work proposes to initialize the model from BERT pre-trained checkpoint and then trains on natural language inference corpus.
- Universal Sentence Encoder (USE) [9]³: USE trains a transformer network on web sources including news, and then augments it with SNLI dataset [6].
- multilingual Universal Sentence Encoder (mUSE) [12]⁴: mUSE extends USE to support 16 languages by training a multi-task siamese network using translation based bridge tasks.
- Laser [2]⁵: Laser represents 93 languages via Bi-LSTM encoder trained on publicly available parallel corpora.
- LaBSE [23]⁶: LaBSE embeds text sentences in a language agnostic way through training the encoder with transformer-based siamese network. The model is first pre-trained on web corpus together with translation sentence pairs, and then is fine-tuned with translation corpus alone.

4.2.1 Semantic Textual Similarity. Inspired by the good performance of pre-trained model in XTREME, we first evaluate the *open-domain* performance on Semantic Textual Similarity (STS) tasks⁷ [8]. Specifically, these tasks require determining how semantically similar a pair of English sentences are.

Unsupervised setup. We use SentEval [13] tool box to evaluate models on STS tasks 2012-2016 [8] without using any STS specific training data. These datasets provide labels between 0 and 5 as semantic relatedness of sentence pairs. Following Reimers and Gurevych [53], we compute the Spearman’s rank correlation ρ

²<https://huggingface.co/sentence-transformers/bert-base-nli-mean-tokens>

³<https://tfhub.dev/google/universal-sentence-encoder/4>

⁴<https://tfhub.dev/google/universal-sentence-encoder-multilingual/3>

⁵<https://github.com/yannnvg/laserembeddings>

⁶<https://tfhub.dev/google/LaBSE/1>

⁷STS dataset is collected from image captions, news headlines and user forums. Thus two out of the three genres are out-of-domain.

Model	MRPC-acc	MRPC-f1	SICK-R	STS-B	Avg.
SBERT	75.1	82.6	79.3	73.6	77.6
USE	72.6	81.8	78.9	79.1	78.1
mUSE	74.6	82.9	80.2	81.5	79.8
Laser	74.1	82.5	79.0	77.8	78.4
LaBSE	74.4	82.3	79.1	77.6	78.4
NewsEmbed	73.5	81.9	79.1	82.3	79.2

Table 5: Evaluation on semantic similarity with head trained on top of frozen embeddings [supervised]. Accuracy/F1 are used for MRPC. Spearman rank correlation is used for SICK-R/STS-B.

between the cosine-similarity of [CLS] embeddings and the gold labels. Performance is reported by convention as $\rho \times 100$.

As shown in Table 3, although trained with news data, NewsEmbed has competitive performance comparing to existing works that are optimized for generic sentence embeddings. Since STS datasets are partially collected from news domain, we further investigate performance variations over data domains. Results are reported in Table 4 for STS benchmark dataset, which is a subset of STS 12-17. One can observe that NewsEmbed outperforms other approaches by a large margin on the news slice, and is still competitive when evaluated on out-of-domain data.

Supervised setup. Supervised learning is another dimension to evaluate text representation quality. The datasets in this experiment include the Microsoft Research Paraphrase Corpus (MRPC) [20], which is a paraphrase identification dataset aiming to identify if two sentences are paraphrases of each other. The SICK Relatedness (SICK-R) dataset [38] contain labels in the range between 1 and 5 to indicate the relatedness of two sentences. The Semantic Textual Similarity Benchmark (STS-B) [8] follows similar label representation but the range is between 0 and 5. On top of encoder’s output, we learn a classification head for MRPC and regression head for SICK-R and STS-B using SentEval.

Results are shown in Table 5. NewsEmbed again shows strong performance as it is only slightly worse than mUSE. The quality drop is likely because of out-of-domain dataset such as SICK-R, and the fact that NewsEmbed has multiple objectives to optimize, affecting the model capacity attributed to short-text similarity.

Based on the above experimental results, we can conclude that NewsEmbed behaves generally well at the sentence granularity regardless of the domain.

4.2.2 News Clustering and Retrieval. A good document encoder is supposed to work well on long contexts. We first evaluate this in an unsupervised setup, and specifically rely on document clustering and retrieval as the proxy tasks. The rationale is that both these tasks depend on a good document representation.

The first dataset used in this experiment is Multi-News [21], which contains journalist-written news summaries and their corresponding source documents. For the purpose of clustering and retrieval, we consider each summary defines a cluster and the similarity within its linked documents should be higher than those not linked. To avoid too many small clusters, we filter the dataset on data points with at least 8 documents, which leaves us with 380 clusters and 3216 source documents.

Model	Multi-News		Stream-News
	ARI	mAP@8	ARI
SBERT	9.7	39.2	32.8
USE	27.6	70.1	35.6
mUSE	26.6	68.3	44.9
Laser	1.7	16.6	7.6
LaBSE	16.7	53.7	39.8
NewsEmbed	29.0	73.1	51.9

Table 6: Adjusted Rand Index for clustering on two datasets and mean Average Precision on Multi-News dataset for retrieving articles from summary [unsupervised].

The second dataset is Multilingual Stream News Cluster dataset (Stream-News for short) [41], proposed here for studying the multilingual performance. It consists of news documents from English, German, and Spanish. We filter out events with size smaller than 50, yielding 11757 evaluation articles with 97 event ids.

For clustering, we apply k -means on document embeddings assuming the cluster size is known and report Adjusted Rand Index [50]. For retrieval, we use Multi-News dataset with the summary as the retrieving query and report mean average precision.

Results are summarized in Table 6. NewsEmbed continues to perform strong when text is long. Since the data is in-domain, it outperforms all other approaches in both clustering and retrieval metrics. Besides NewsEmbed, we notice USE and mUSE are notably competitive. We suspect heavily pre-training on news data is critical as both approaches have this component. In comparison, LaBSE is pre-trained on web data but then heavily tuned on translation sentence pairs, which biases the model to excel at detecting meaning equivalence in cross-lingual environment. The rest two approaches, i.e., SBERT and Laser, lack domain adaptation for news.

4.2.3 News Classification. In this experiment, we evaluate NewsEmbed in the setting of document classification. Similar to previous experiment in Section 4.2.1, we train a logistic regression on top of the [CLS] token embedding. We consider four popularly used datasets that cover a variety of publishers and news styles: BBC [24], 20 Newsgroups⁸, AG News [63], and MIND-small [59], with statistics outlined in Table 7.

To evaluate the effectiveness of the embeddings, we only use 10% training data of AG News. The dataset is down-sampled because 1) logistic regression doesn't require a lot of training data, and 2) extracting embeddings is time costly for baselines. The classification accuracy scores are reported in Table 8. Similar to previous unsupervised experiments, NewsEmbed continues to be the best, followed by mUSE. LaBSE ranks as the 3rd best method, and performs relatively better in supervised than unsupervised settings. This is consistent with the behaviour in the STS experiments.

Dataset	# Classes	# Train	# Test
BBC	5	1780	445
20 Newsgroups	20	11314	7532
10% AG News	4	12000	7600
MIND-small	18	51421	12856

Table 7: Dataset statistics for document classification.

⁸<http://qwone.com/~jason/20Newsgroups/>

Model	BBC	20 Newsgroup	AG News	MIND	Avg.
SBERT	96.0	64.5	86.5	75.2	80.5
USE	97.3	66.4	85.8	75.0	81.1
mUSE	97.3	71.4	86.8	75.8	82.8
LaBSE	96.6	70.3	86.3	73.9	81.8
Laser	92.1	63.5	81.4	63.1	75.0
NewsEmbed	97.5	73.9	89.1	77.0	84.2

Table 8: Accuracy scores for news classification datasets with logistic classifiers on top of frozen embeddings [supervised].

	Before	After
XNLI	76.4	72.8
PAWS-X	88.3	84.1
UDPOS	76.5	72.8
PANX	62.2	61.2
XQuAD	74.3	72.4
MLQA	70.6	67.5
TyDiQA	68.7	66.7
Tatoeba	80.1	77.6

Table 9: XTREME benchmarks on test dataset before and after weakly supervised learning.

4.2.4 XTREME Comparison. We compare XTREME benchmarks before and after weakly supervised learning and show the result in Table 9. We find that all tasks get moderate metric drop. Considering previous experiments about encoders capturing topical semantics given either short or long texts, it seems that NewsEmbed achieves outstanding performance over there via sacrificing its pre-trained reasoning and representation power at token level.

Based on this result, it becomes interesting to co-train masked language model objective with the InfoNCE and BCE objectives in the triplet network. We leave this as future work.

4.3 Ablation Studies

During the development of NewsEmbed, to better study alternative designs, we mainly resort to some in-house benchmarks. These include several semantic textual similarity corpora in the news domain, with the focus on evaluating sentence-to-document and document-to-document similarities. To measure model multilinguality more accurately, we record metrics in (English, Spanish, Hindi, Chinese) and compute the z-scores over specified step ranges and ablations for each benchmark. Then we aggregate all scores by taking the mean.

We mainly investigate the following ablations:

- *-Pre-training:* Randomly initializing the checkpoint.
- *-Alternate training:* Not using alternate training by mixing all <language, task> datasets together.
- *-Cross-accelerator negatives:* In-batch negatives are from the same accelerator in distributed learning.
- *-Translation augmentation:* Removing the translation in document triplets for non-English texts.
- *-Triplet classifier:* Removing the light-weight classifier to filter out noises in document triplets.
- *-Document-topic associations:* Removing the document-topic associations from the datasets.

Ablation	Z-score change	
	en	other languages
-Pre-training	-1.35	-1.51
-Alternate training	-0.73	-0.97
-Cross-accelerator negatives	-0.34	-0.67
-Translation augmentation	-0.42	-0.72
-Triplet classifier	-0.34	-0.70
-Document-topic associations	-0.45	-0.46

Table 10: Z-score changes of different ablation studies.

Table 10 shows the effect of removing each component from NewsEmbed model. First of all, we observe that the pre-training stage is the most critical although our weak supervision dataset is pretty large. Next, alternate training greatly improves optimization effectiveness, mainly because in-batch negatives are sharing the same languages. Similarly, increasing the effective batch size through cross-accelerator sharing also shows good improvement and noticeably brings more benefits for non-English. Our guess is that by making negative examples harder incorporates some false negatives, but it affects less for languages other than English due to smaller publisher density. Another study that gives better z-score for non-English languages is adding the translation augmentation. It verifies our conjecture about knowledge transfer by language alignment. The next ablation study shows the positive results about applying triplet classifier to clean data. It is interesting that non-English corpus improves significantly because of it, further confirming that the quality of the weak supervision is more important than the quantity. Lastly, the dataset of document-topic associations turns out to be helpful for both high- and low-resource languages, and we believe it is complementary to the document triplet dataset.

5 RELATED WORK

We have discussed in details in Section 1 about recent NLP developments in the area of text representation. The most related work is the recent pre-training/fine-tuning framework that significantly improve the encoder capability [19, 35]. Here we describe the rest of related work on contrastive learning and modeling of news.

Contrastive learning by augmentation. Contrastive learning can be considered as learning object representation by comparing. In computer vision, several works are proposed for self-supervised learning in a contrastive way [11, 25, 29, 42]. The high level idea is to encourage the model to learn representation of images so that different views or augmentations of the same image are close to the source in latent space. When it comes to textual data, Feng et al. [23] use a dual encoder with additive margin softmax on bilingual translation pairs. Fang and Xie [22] use back-translation as data augmentation approach that translate a sentence into a target language and back to the source language.

In our case, we mine large amount of document triplets as well as document-topic associations from the news corpus. This is based on a few observations that is unique to news domain. We consider our approach as weakly supervised because it derives noisy supervision signal based on these observations.

News modeling. Existing works about modeling news content mainly includes clustering, recommendation, and classification.

Works related to clustering can be categorized into online [41] and offline [52]. They usually compute textual representation by features such as bag-of-words [1], learned topics [52], or deep learning representations [27]. Our work differentiates from them by using a novel approach to mine web-scale knowledge and shows outstanding performance on clustering benchmarks.

For recommendation, a good representation of news articles and users are critical. Li et al. [37] represents news article using content, access patterns, named entities, popularity and recency. Okura et al. [44] uses denoising autoencoder to learn embeddings based on the similarities between news articles in the same and different categories. Next, de Souza Pereira Moreira et al. [17] predict next-item for users sessions by representing news based on text and meta-data and modeling session-based recommendation using Recurrent Neural Networks. Wu et al. [58] applies a convolutional neural network (CNN) to encode news articles followed by a word-level personalized attention network. Our triplet approach is similar to Okura et al. [44], but with a much higher granularity and scale of data. We focus more on the semantic representation by only using title and body to encode the document.

News classification tasks include fake news detection [34, 46] and new category classification. People use traditional approaches such as Naive Bayes and Support Vector Machines [15, 54], and neural network approaches such as LSTM [51] and CNN [18]. In our work, we propose to train a *universal* encoder for multilingual document representation, and support fine-tuning the model on downstream classification tasks.

6 CONCLUSION AND FUTURE WORK

In this work, we propose to study the problem of deriving a universal document embedding within the news domain. Based on observations unique to news domain, we collect weak supervision data in large scale and in different languages. An effective neural architecture, named NewsEmbed, is then proposed to co-train contrastive learning and multi-label classification from this data. The model demonstrates strong performance regardless of context length through a series of unsupervised and supervised evaluations, some of which are even out-of-domain.

Nevertheless, there are a few unsolved challenges. For instance, we use the full title with truncated body from the beginning up to 512 tokens because transformer encoder is expensive with longer input, while 80% articles have content longer than that. Some recent works have been proposed to accelerate the computation of long-text transformer with certain compromise on quality. We argue that by modeling the inductive bias where title and leading paragraphs can mostly summarize the whole document, one might achieve both efficiency and effectiveness. Another challenge would be encoding news with modern content types including image, audio and video in a multi-modal fashion. Aligning the space among these modalities and mutually enhancing through this process will become more critical as the streaming service is becoming increasingly diverse, including but not limited to podcast, radio, short video, live broadcast, etc. Lastly, our approach mostly relies on time-sensitivity to remove noise and therefore lacks training signal to model evolving topics in long running stories. We are actively exploring those directions as future work.

REFERENCES

- [1] C. C. Aggarwal and P. S. Yu. A framework for clustering massive text and categorical data streams. In *SDM*, pages 479–483, 2006.
- [2] M. Artetxe and H. Schwenk. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *TACL*, 7:597–610, 2019.
- [3] S. Ben-David, J. Blitzer, K. Crammer, and F. Pereira. Analysis of representations for domain adaptation. In *NeurIPS*, pages 137–144, 2007.
- [4] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *JMLR*, 3(Jan): 993–1022, 2003.
- [5] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov. Enriching word vectors with subword information. *TACL*, 5:135–146, 2017.
- [6] S. R. Bowman, G. Angeli, C. Potts, and C. D. Manning. A large annotated corpus for learning natural language inference. In *EMNLP*, 2015.
- [7] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- [8] D. Cer, M. Diab, E. Agirre, I. Lopez-Gazpio, and L. Specia. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation*, pages 1–14, 2017.
- [9] D. Cer, Y. Yang, S.-y. Kong, N. Hua, N. Limtiaco, R. S. John, N. Constant, M. Guajardo-Cespedes, S. Yuan, C. Tar, et al. Universal sentence encoder for english. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, 2018.
- [10] W.-C. Chang, F. X. Yu, Y.-W. Chang, Y. Yang, and S. Kumar. Pre-training tasks for embedding-based large-scale retrieval. In *ICLR*, 2020.
- [11] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020.
- [12] M. Chidambaram, Y. Yang, D. Cer, S. Yuan, Y. Sung, B. Strope, and R. Kurzweil. Learning cross-lingual sentence representations via a multi-task dual-encoder model. In *Proceedings of the 4th Workshop on Representation Learning for NLP*, pages 250–259, 2019.
- [13] A. Conneau and D. Kiela. Senteval: An evaluation toolkit for universal sentence representations. In *LREC*, 2018.
- [14] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, É. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov. Unsupervised cross-lingual representation learning at scale. In *ACL*, pages 8440–8451, 2020.
- [15] N. K. Conroy, V. L. Rubin, and Y. Chen. Automatic deception detection: Methods for finding fake news. *Proceedings of the Association for Information Science and Technology*, 52(1):1–4, 2015.
- [16] Y. Cui, W. Che, T. Liu, B. Qin, Z. Yang, S. Wang, and G. Hu. Pre-training with whole word masking for chinese bert. *arXiv preprint arXiv:1906.08101*, 2019.
- [17] G. de Souza Pereira Moreira, F. Ferreira, and A. M. da Cunha. News session-based recommendations using deep neural networks. In *Proceedings of the 3rd Workshop on Deep Learning for Recommender Systems*, pages 15–23, 2018.
- [18] N. Deligiannis, T. Huu, D. M. Nguyen, and X. Luo. Deep learning for geolocating social media users and detecting fake news. In *NATO Workshop*, 2018.
- [19] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186, 2019.
- [20] W. Dolan, C. Quirk, C. Brockett, and B. Dolan. Unsupervised construction of large paraphrase corpora: Exploiting massively parallel news sources. pages 350–356, 2004.
- [21] A. R. Fabbri, I. Li, T. She, S. Li, and D. Radev. Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model. In *ACL*, pages 1074–1084, 2019.
- [22] H. Fang and P. Xie. Cert: Contrastive self-supervised learning for language understanding. *arXiv preprint arXiv:2005.12766*, 2020.
- [23] F. Feng, Y. Yang, D. Cer, N. Arivazhagan, and W. Wang. Language-agnostic bert sentence embedding. *arXiv preprint arXiv:2007.01852*, 2020.
- [24] D. Greene and P. Cunningham. Practical solutions to the problem of diagonal dominance in kernel document clustering. In *ICML*, pages 377–384, 2006.
- [25] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doherty, B. A. Pires, Z. D. Guo, M. G. Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.
- [26] R. Guo, P. Sun, E. Lindgren, Q. Geng, D. Simcha, F. Chern, and S. Kumar. Accelerating large-scale inference with anisotropic vector quantization. In *ICML*, 2020.
- [27] X. Guo, L. Gao, X. Liu, and J. Yin. Improved deep embedded clustering with local structure preservation. In *IJCAI*, pages 1753–1759, 2017.
- [28] S. Gururangan, A. Marasović, S. Swayamdipta, K. Lo, I. Beltagy, D. Downey, and N. A. Smith. Don’t stop pretraining: Adapt language models to domains and tasks. In *ACL*, pages 8342–8360, 2020.
- [29] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020.
- [30] J. Hu, S. Ruder, A. Siddhant, G. Neubig, O. Firat, and M. Johnson. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *ICML*, pages 4411–4421. PMLR, 2020.
- [31] T. Joachims. Text categorization with support vector machines: Learning with many relevant features. In *ECML*, pages 137–142. Springer, 1998.
- [32] S.-M. Kim, P. Pantel, T. Chklovski, and M. Pennacchiotti. Automatically assessing review helpfulness. In *EMNLP*, pages 423–430, 2006.
- [33] K. Kowsari, K. Jafari Meimandi, M. Heidarysafa, S. Mendu, L. Barnes, and D. Brown. Text classification algorithms: A survey. *Information*, 10(4):150, 2019.
- [34] P. Ksieniewicz, M. Choraś, R. Kozik, and M. Woźniak. Machine learning methods for fake news classification. In *IDEAL*, pages 332–339, 2019.
- [35] G. Lample and A. Conneau. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*, 2019.
- [36] P. H. Le-Khac, G. Healy, and A. F. Smeaton. Contrastive representation learning: A framework and review. *IEEE Access*, 2020.
- [37] L. Li, D. Wang, T. Li, D. Knox, and B. Padmanabhan. Scene: a scalable two-stage personalized news recommendation system. In *SIGIR*, pages 125–134, 2011.
- [38] M. Marelli, S. Menini, M. Baroni, L. Bentivogli, R. Bernardi, R. Zamparelli, et al. A sick cure for the evaluation of compositional distributional semantic models. In *LREC*, pages 216–223, 2014.
- [39] O. Melamud, J. Goldberger, and I. Dagan. context2vec: Learning generic context embedding with bidirectional lstm. In *CoNLL*, pages 51–61, 2016.
- [40] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NeurIPS*, volume 26, pages 3111–3119, 2013.
- [41] S. Miranda, A. Znotins, S. B. Cohen, and G. Barzdins. Multilingual clustering of streaming news. In *EMNLP*, pages 4535–4544, 2018.
- [42] I. Misra and L. v. d. Maaten. Self-supervised learning of pretext-invariant representations. In *CVPR*, pages 6707–6717, 2020.
- [43] K. P. Murphy. *Machine learning: a probabilistic perspective*. 2012.
- [44] S. Okura, Y. Tagami, S. Ono, and A. Tajima. Embedding-based news recommendation for millions of users. In *SIGKDD*, pages 1933–1942, 2017.
- [45] A. v. d. Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [46] R. Oshikawa, J. Qian, and W. Y. Wang. A survey on natural language processing for fake news detection. In *LREC*, pages 6086–6093, 2020.
- [47] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543, 2014.
- [48] M. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer. Deep contextualized word representations. In *NAACL*, pages 2227–2237, 2018.
- [49] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 21:1–67, 2020.
- [50] W. M. Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical association*, 66(336):846–850, 1971.
- [51] H. Rashkin, E. Choi, J. Y. Jang, S. Volkova, and Y. Choi. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, pages 2931–2937, 2017.
- [52] R. Rehurek and P. Sojka. Software framework for topic modelling with large corpora. In *In Proceedings of the LREC 2010 workshop on new challenges for NLP frameworks*, 2010.
- [53] N. Reimers and I. Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *EMNLP*, pages 3973–3983, 2019.
- [54] K. Shu, D. Mahudeswaran, S. Wang, D. Lee, and H. Liu. Fakenewsnet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big Data*, 8(3):171–188, 2020.
- [55] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017.
- [56] R. Venkatesan, S.-M. Koon, M. H. Jakubowski, and P. Moulin. Robust image hashing. In *ICIP*, volume 3, pages 664–666, 2000.
- [57] A. Williams, N. Nangia, and S. Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *NAACL: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, 2018.
- [58] C. Wu, F. Wu, M. An, J. Huang, Y. Huang, and X. Xie. NPA: neural news recommendation with personalized attention. In *SIGKDD*, pages 2576–2584, 2019.
- [59] F. Wu, Y. Qiao, J.-H. Chen, C. Wu, T. Qi, J. Lian, D. Liu, X. Xie, J. Gao, W. Wu, et al. Mind: A large-scale dataset for news recommendation. In *ACL*, pages 3597–3606, 2020.
- [60] L. Xue, N. Constant, A. Roberts, M. Kale, R. Al-Rfou, A. Siddhant, A. Barua, and C. Raffel. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*, 2020.
- [61] F. Zhang, N. J. Yuan, D. Lian, X. Xie, and W.-Y. Ma. Collaborative knowledge base embedding for recommender systems. In *SIGKDD*, pages 353–362, 2016.
- [62] X. Zhang and Y. LeCun. Text understanding from scratch. *arXiv preprint arXiv:1502.01710*, 2015.
- [63] X. Zhang, J. Zhao, and Y. LeCun. Character-level convolutional networks for text classification. *NeurIPS*, 28:649–657, 2015.
- [64] C. Zhou, C. Sun, Z. Liu, and F. Lau. A c-lstm neural network for text classification. *arXiv preprint arXiv:1511.08630*, 2015.

A SHORT TEXT PACKING AT PRE-TRAINING

During data preparation, we observe that some documents and translation sentences are much shorter comparing to standard BERT input sequence length 512. This can cause lots of paddings in order to run on TPU, making the pre-training inefficient. To fully use sequence length, we apply a greedy algorithm to pack multiple short sequences into one training example. The packing algorithm achieves 8.5 compression ratio for the dataset of translation sentence pairs.

As shown in Algorithm 1, a priority instance packer has attributes as cache capacity C , max sequence length L , and min packing proportion ρ . The packer maintains a map M of partially packed instances with key as sequence length (with tie-breaker by instance id) ordered non-increasingly. In practice, we use Apache Beam to process the large scale data.

Algorithm 1: Training instance packing algorithm

```

Data:  $N$  Training Instances  $\{X_i\}_{i=1}^N$ 
Result: List of Packed Training Instances  $Y$ 
 $Y = []$ ;
for  $i$  in  $1:N$  do
  processed = False;
  for  $k, v$  in  $M$  do
    if  $\text{len}(X_i) + \text{len}(v) > L$  then
      | continue;
    end
    remove  $k$  from  $M$ ;
     $\tilde{v}$  = instance by concatenating  $v$  and  $X_i$ ;
    processed = True;
    if  $\text{len}(\tilde{v}) \geq \rho L$  then
      |  $Y.\text{append}(\tilde{v})$ ;
    else
      |  $M[(\text{len}(\tilde{v}), i)] = \tilde{v}$ ;
    end
    break;
  end
  if not processed then
     $M[(\text{len}(X_i), i)] = X_i$ ;
    if  $\text{len}(M) > C$  then
      | remove the first element  $(k_f, v_f)$  from  $M$ ;
      |  $Y.\text{append}(v_f)$ ;
    end
  end
end

```

B TECHNICAL DETAILS IN WEAKLY SUPERVISED LEARNING

B.1 Training

We initialize NewsEmbed from a BERT-style pre-trained checkpoint described in Section 3.1 and train triplet network with weak supervision using 256 TPU v3 chips. The batch size is 8,196 and we train the model for 400k steps, which is roughly three epochs. Similar to BERT, Adam optimizer is deployed with learning rate set as $5e-5$.

In the contrastive learning objective, i.e., Equation 1, temperature τ is set as 0.05 and the vectors in the augmented triplet are normalized by their L2-norm before computing the inner product. In-batch negatives Z include cross-accelerator negatives to allow us to fully realize the benefits of distributed training.

In the multi-label classification objective, i.e., Equation 2, T comprises of 600 hand-picked topics. Positive to negative ratio is empirically set as 25% through down-sampling mentioned in Section 2.2.2.

During the training, the total loss is computed as the summation of Equation 1 and Equation 2 without weights.

B.2 Languages and vocabulary

We select a total of 101 languages that have footprints in our proprietary system. To ensure a smoother distribution so that the model can be later trained better on low resource data, we use an exponential up-sampling approach [35] with smooth coefficient 0.7. More concretely, suppose the N languages have counts distributed as $\{n_i\}_{i=1\dots N}$ with $n_1 \geq n_2 \geq \dots \geq n_N$, we will up-sample the articles so that after sampling, language i has expected m_i articles, where

$$m_i = n_1 \left(\frac{p_i}{p_1} \right)^\alpha.$$

Here α is the smooth coefficient between 0 and 1 inclusive. If $\alpha = 0$, every language is re-sampled to the largest language. If $\alpha = 1$, the distribution will remain unchanged. If m_i is not an integer, we draw a random Bernoulli variable with probability of m_i 's fractional part to round it to an integer.

We use a large vocabulary size of 500K and apply wordpiece algorithm to learn the tokens from the above smoothed dataset. Note that the vocabulary is large because multilingual transformer models can benefit from allocating a higher proportion of the total number of parameters to the embedding layer [14] while the time complexity of embedding lookup operation is almost constant w.r.t. vocabulary size.

C GENERALIZATION CAPABILITY

It is challenging to model news partially because of the out-of-vocabulary concepts and evolving events. A good encoder not only needs to fit existing data well, but also is supposed to show outstanding generalization capability for future content.

Realizing this, we design the following experiment to measure performance degeneration over time. Firstly, we ensure our evaluation benchmarks only contain recently published documents after January 1, 2020. Next, two training datasets are re-collected by deleting documents published after January 1, 2020 and January 1, 2019 respectively. After that, we retrain NewsEmbed on them and compare their evaluation metrics with the model trained on the complete corpus (containing documents up to June 30, 2020). We consider the differences in benchmark metrics should reflect our system's generalization capability.

The reason we choose January 1, 2020 as the partition date is because of the overwhelming discussions about COVID-19 on the web regardless of languages. By training a model on datasets without any COVID-related topics, we expect this study can measure the model's staleness more accurately and guide us how often to re-train a model to retain reasonable performance.

It is worth noting that NewsEmbed is initialized on a pre-trained checkpoint. To prevent it from leaking information, we apply the same data filter at the pre-training stage.

Results are reported in Table 11. As one can tell, z-score of the model trained on fresher training dataset drops less, verifying that the model quality becomes worse over time. Moreover, due to the topic-drift brought by COVID-19, the loss becomes significantly large in the recent six months, regardless of the languages. This guides us to retrain the model every few months especially when there are major breaking events happening.

Ablation	Z-score change	
	en	other languages
Baseline (up to 06/30/2020)	0.00	0.00
- 6 months (up to 01/01/2020)	-1.74	-0.79
- 18 months (up to 01/01/2019)	-2.16	-1.27

Table 11: Z-score changes after removing recent documents from training datasets.