

Improving Sampling Methods for Fine-tuning SentenceBERT in Text Streams

Cristiano Mesquita Garcia^{1,3}[0000-0002-7475-146X], Alessandro Lameiras Koerich²[0000-0001-5879-7014], Alceu de Souza Britto Jr.^{3,4}[0000-0002-3064-3563], and Jean Paul Barddal³[0000-0001-9928-854X]

¹ Instituto Federal de Santa Catarina (IFSC), Câmpus Caçador, Brazil
`cristiano.garcia@ifsc.edu.br`

² École de Technologie Supérieure (ÉTS), Université du Québec, Montréal, Canada
`alessandro.koerich@etsmtl.ca`

³ Pontifícia Universidade Católica do Paraná (PUCPR), Curitiba, Brazil
`{alceu,jean.barddal}@ppgia.pucpr.br`

⁴ Universidade Estadual de Ponta Grossa (UEPG), Ponta Grossa, Brazil

Abstract. The proliferation of textual data on the Internet presents a unique opportunity for institutions and companies to monitor public opinion about their services and products. Given the rapid generation of such data, the text stream mining setting, which handles sequentially arriving, potentially infinite text streams, is often more suitable than traditional batch learning. While pre-trained language models are commonly employed for their high-quality text vectorization capabilities in streaming contexts, they face challenges adapting to concept drift—the phenomenon where the data distribution changes over time, adversely affecting model performance. Addressing the issue of concept drift, this study explores the efficacy of seven text sampling methods designed to fine-tune language models, thereby mitigating performance degradation selectively. We precisely assess the impact of these methods on fine-tuning the SBERT model using four different loss functions. Our evaluation, focused on Macro F1-score and elapsed time, employs two text stream datasets and an incremental SVM classifier to benchmark performance. Our findings indicate that Softmax loss and Batch All Triplets loss are particularly effective for text stream classification, demonstrating that larger sample sizes correlate with improved macro F1 scores. Notably, our proposed WordPieceToken ratio sampling method significantly enhances performance with the identified loss functions, surpassing baseline results.

Keywords: Text stream · Language Model · Concept drift · Sampling methods · Fine-tuning.

1 Introduction

The Internet has become part of daily life around the world. People and systems generate a vast amount of textual data through the Internet. Individuals can

chat with others, review products and services, and share comments and opinions through social media platforms, frequently working as social sensors [24]. Learning from social media posts can be relevant for institutions and governments, helping them quickly detect and respond to events [11, 19], for example.

Automatically learning from textual data leveraging machine learning mechanisms brings several challenges in batch processing, such as text standardization and vectorization. Text vectorization is essential since most machine learning methods expect numeric vectors as input. Traditional vector representations, such as Bag-of-Words (BOW) [12] and Term Frequency - Inverse of Document Frequency (TF-IDF) [21], can generate very-high-dimensional vectors, which can be disadvantageous to the machine learning model, increasing the computational cost.

In a textual stream scenario, the challenges are augmented. Due to the stream characteristics, e.g., data arriving on an instance-basis or in small batches and resource limitations [3, 9], generating vector representations through BOW and TF-IDF is complex. For instance, if the vector representations are generated in the first batch, new words in subsequent batches will not be represented. On the other hand, generating the representations as the batches arrive can lead to variable-dimension representations [10], a challenge since most machine learning algorithms require a fixed-dimension input.

Therefore, pre-trained language models have become popular in batch and stream scenarios due to their time-saving characteristics [10, 25]. SentenceBERT (SBERT) [20] is a popular pre-trained language model specific for sentence embedding generation. Although pre-trained models save time since training a language model from scratch is costly, adjustments may be necessary for domain adaptation. In addition, changes in data distribution over time (concept drift) are frequent phenomena in real-world data and can degrade a machine-learning model’s performance [9]. In textual data streams, those changes can emerge from sentiment changes, the appearance of particular words in different contexts, and so on [10]. Furthermore, computational linguistic studies using diachronic datasets attest that writing patterns change and word meanings evolve over time, e.g., semantic shift [4]. This work refers to changes as concept drift since semantic shifts are generally related to linguistic studies, which use long timespan, or diachronic, datasets and investigate those changes on a deeper, linguistic level. Therefore, to adapt to concept drift or a new domain, for example, it is important to adapt the language model. The fine-tuning process is a popular deep-learning-related process and can help adapt the language model [17, 22].

Due to the computational cost of the fine-tuning process, selecting representative instances to fine-tune the language model may provide valuable information while reducing the time spent [1, 23]. In this paper, we score the ability of different sampling methods in text selection for fine-tuning purposes. We also propose a sampling method, i.e., WordPieceToken ratio, whose results were promising in most scenarios evaluated. Considering the text stream setting, we assessed these methods intrinsically, i.e., in a downstream task. We also evaluated three ver-

sions of these sampling methods modified to account for the classes, totaling seven sampling methods.

The contributions of this paper are four-fold: (a) an extensive comparison among text sampling methods for fine-tuning purposes; (b) an analysis of the impact of the sampling methods considering the text stream setting; (c) an evaluation of loss functions for fine-tuning SBERT, and (d) a novel textual sampling method based on the ratio between Wordpieces and tokens of a text. The term Wordpieces represents a subword partition system present in BERT [8] that allows handling out-of-vocabulary tokens.

This paper is organized as follows: Section 2 presents important concepts for understanding this paper, including text streams and SentenceBERT. Section 3 presents the text-based sampling methods evaluated in this paper. Section 4 describes the experimental protocol, with datasets, settings, evaluation scenario, and results. Finally, Section 5 concludes this paper.

2 Background

This section presents core concepts for understanding this paper. We introduce text stream mining, SentenceBERT, and its fine-tuning process.

2.1 Text Stream Mining

According to Bifet et al. [3], data streams are “an abstraction that allows real-time analytics”. In a data stream, the items arrive individually or in small sequential batches, and the stream itself can be infinite [3, 9]. Different learning approaches have been developed to learn from data streams, including, for instance [9], the ability to learn incrementally, single-pass operations, elimination of input data as soon as possible after learning from it, and consumption of modest resources, i.e., processing power and time.

Text streams are a specialization of data streams in which texts arrive over time [10]. The challenges are extended in this scenario, mainly comprising natural language processing (NLP), such as text standardization, vocabulary, and representation maintenance. These NLP-related processing are challenging due to their complexity, which should meet the text stream constraints.

Frequently, text-related approaches leverage pre-trained language models [10, 25]. Using pre-trained models can help save time since training a language model from scratch is computationally costly in time and resources [23]. In addition, the language model can easily be reused in different scenarios. However, an important drawback is that texts generally suffer from concept drift. Concept drift is a phenomenon frequently observed in real-world datasets and corresponds to changes in data distribution over time [9, 10]. Leveraging pre-trained language models without accounting for concept drift can decrease performance in the downstream task since the texts would be represented using relatively old representations [10].

This paper leverages a pre-trained language model and evaluates the use of fine-tuning for language model updates in text stream settings, a less costly approach than training from scratch. In this paper, the selected language model is SentenceBERT [20].

2.2 SentenceBERT

SentenceBERT (SBERT) is an architecture that leverages pre-trained BERT models [20], such as BERT [8] and RoBERTa [18] models. SBERT leverages siamese networks to generate semantically meaningful representations that are compared using cosine similarity. A siamese architecture with a bi-encoder reduces the computational overhead while improving the quality of representations, compared to a cross-encoder to determine sentence similarity, as in BERT [20].

Additionally, SBERT provided significant improvements for semantic text similarity. Although the authors fine-tuned the SBERT model on natural language inference (NLI) data and also applied the model to semantic textual similarity (STS) task, SBERT demonstrated competitive results when being used as a text vectorization method for classification tasks [25].

Fine-tuning, data preparation, and loss functions SBERT allows several strategies for the fine-tuning process. Typically, SBERT requires texts and a label. Due to the siamese characteristic of SBERT, generally, it requires text pairs (or triplets) and a label. Depending on the strategy, this label can correspond to a class, relatedness degree between texts, or relatedness class between texts, e.g., contradiction, neutral, or entailment. The loss functions and respective strategies allowed by SBERT include:

- **Batch All Triplets loss (BATL)** [14], which requires single texts and their respective classes. Internally, same-class texts are treated as positive anchors, while distinct-class texts are considered negative anchors;
- **Cosine Similarity loss (CSL)**, which expects text pairs and a cosine similarity score as a label. A clear drawback is the demand for a cosine similarity, which requires an extra method/ground truth;
- **Contrastive Tension loss (CTL)** [5], which receives single texts without labels. In this case, exact texts are treated as positive anchors, and all the texts are randomly mixed to generate negative anchors. It uses a ratio to define the number of negative anchors for each positive anchor;
- **Multiple Negative Ranking loss (MNRL)** [13] uses only positive text pairs (or triplets with a negative anchor appended) without label. In this case, texts are mixed to generate negative anchors. A noticeable characteristic is the need for a ground truth to indicate positive texts;
- **Online Contrastive loss (OCL)** requires text pairs and a label indicating their relationship. In this case, the loss function is calculated per item;
- **Softmax loss (SL)** receives text pairs and a label indicating their relationship. Reimers and Gurevych [20] leveraged this loss function for natural language inference, and therefore, the possible labels were contradiction, neutral, or entailment.

The list above is non-exhaustive. This paper aims to evaluate text-based sampling methods (see Section 3) to gather more useful texts for fine-tuning. Considering the above loss functions, we selected BATL, CTL, OCL, and SL. We did not select CSL and MNRL because CSL depends on a similarity measure for the text pairs and would require extra information for this calculation. MNRL, on the other hand, requires positive pairs, or triplets, with a negative anchor. Although we could leverage the classes to generate positive pairs, choosing good-quality anchors can be challenging and require deeper analysis for an assertive selection.

The high cost of training language models is well-known [23]. Although fine-tuning is cheaper than training from scratch, using all new data can also lead to high costs [1, 23]. Thus, resorting to sampling methods can be beneficial in two aspects: (a) selecting more informative texts and (b) consuming fewer computational resources than using all new data.

3 Text-based Sampling Methods

This section presents the sampling methods used in this paper for selecting texts for fine-tuning purposes. We also propose the WordpieceToken ratio and later compare it to other text-based sampling methods.

In addition to each sampling method, except for the random sampling, we evaluate an extra scenario leveraging the text labels (classes). Therefore, the sampling methods correspond to their original version and the version that accounts for the class. Algorithm 1 provides the weighted sampling pseudocode. We highlight that, optionally, the observed classes’ frequencies can be used as an argument for the `WeightedSampling` function. These frequencies can also be calculated directly from the buffer, using the attribute `class`. The algorithm runs according to the following steps: (1) the buffer containing the stored items from the stream is iterated; (2) each item has its weight calculated, depending on the chosen sampling method; (3) if the classes’ frequencies are considered, then the items’ probabilities of less frequent classes are increased proportionally; (4) the weights are normalized; and (5) n_s instances are sampled from the buffer.

The text sampling methods employed in this paper are:

- **Length-based sampling** [1]: it ponders items by their length. This means we counted the number of tokens in each text and normalized them using the biggest and the lowest lengths. The main idea is that longer texts have more chance to encompass useful, novel information for the language model in the fine-tuning process;
- **Random sampling**: this sampling method randomly selects a given number of items from the buffer;
- **TF-IDF-based sampling**: Term Frequency - Inverse of Document Frequency (TF-IDF) [21] is a technique for measuring the importance of a given word in a document, considering a collection. The term frequency is the count of a token t in a text d . The inverse of document frequency is generally calculated as $\text{idf}(t) = \log \frac{n}{\text{document_frequency}(t)}$, where $\text{document_frequency}(t)$

Algorithm 1 Algorithm of weighted sampling

Require: n_s ▷ number of instances to be sampled
Require: `classes_frequencies` ▷ observed classes frequencies
Require: `buffer`
Ensure: `buffer` $\neq \emptyset$

```

function WEIGHTEDSAMPLING( $n_s$ , buffer, classes_frequencies: Optional)
  for each  $X_i \in$  buffer do
    Calculate  $X_i$ .weight according to the sampling method
    if classes_frequencies  $\neq$  None then
       $X_i$ .weight*  $\leftarrow X_i$ .weight  $\times \frac{\text{sum}(\text{classes\_frequencies})}{\text{classes\_frequencies}[X_i.\text{class}]}$ 
    end if
  end for
  Normalize weights
  Sample  $n_s$  instances using the calculated  $X_i$ .weight* as probability
end function

```

corresponds to the number of documents containing t , and n is the total number of documents. Then, the complete TF-IDF calculation is: $\text{TFIDF}(t, d) = \text{term_frequency}(t, d) \times \text{idf}(t)$. Given a text, the TF-IDF is calculated for each token in the text. Thus, the text’s weight is the sum of the TF-IDF values of the tokens present in the text. The rationale behind this approach is to select texts with more important words across the buffer.

In addition to these sampling methods, we propose a novel text sampling method named WordpieceToken ratio.

3.1 WordpieceToken ratio sampling

This paper proposes a novel sampling method based on the ratio between wordpieces and tokens. Wordpiece is a technique BERT-based models use to handle out-of-vocabulary (OOV) tokens [8]. For example, the word `institutionalization` could be partitioned into two wordpieces: [`institutional`, `##ization`], where `##` means that there is a previous partition.

The rationale behind this sampling method is that the bigger the ratio between wordpieces and tokens, the bigger the number of unknown words (by the language model) in the text. The ratio is calculated per instance. For example, let the text instance be: “Extramedullary toxicity was limited to hypothyroidism”, a 6-token sentence. When applied to a BERT tokenizer, it is split into [`extra`, `##med`, `##ulla`, `##ry`, `toxicity`, `was`, `limited`, `to`, `h`, `##yp`, `##oth`, `##yr`, `##oid`, `##ism`], resulting in 14 wordpieces. Thus, this instance’s weight is $\frac{14}{6} = 2.33$. After all the text instances from the buffer have their weights calculated, several instances are sampled considering the instances’ weights and class frequencies according to Algorithm 1.

Therefore, we hypothesize that sampling texts by weighting the Wordpiece-Token ratio may retrieve texts with more useful information for the fine-tuning process.

4 Experimental Results

This section provides the experimental results, including the experimental protocol, datasets, proposed scenario, evaluation metrics, and the results.

4.1 Experimental Protocol

Datasets Two datasets were used: Airbnb and Yelp.

- **Airbnb:** This dataset was obtained from the Inside Airbnb⁵, considering data related to the New York City. Since New York City is one of the most popular destinations in the United States⁶, the Airbnb dataset related to New York City has several reviews, enough for multiple sampling for text stream simulation (see Section 4.2). This dataset, by default, is not ready for classification tasks. Therefore, we leveraged: (a) a pre-trained model for language identification is used (lid.176.ftz [15, 16]), and (b) a pre-trained model for sentiment analysis to infer the reviews’ sentiment (Twitter RoBERTa Base Sentiment⁷ model [2]). Thus, the English reviews were filtered, and their sentiments were inferred. The sentiments, i.e., positive, negative, and neutral, were used as labels in the classification task. The processing steps are available on Github⁸.
- **Yelp:** This dataset is provided on Yelp Datasets⁹. Yelp consists of reviews collected regarding over 130 thousand businesses. These reviews are accompanied by a category from a scale of stars between 1 and 5. This category is used as a label in the text classification task.

An essential characteristic regarding the above datasets is the presence of a timestamp, which is crucial for text stream simulation. The data distributions are presented in Fig. 1. Noticeably, the datasets are imbalanced, reflecting the nature of real-world data.

Other relevant information on the experiments are:

- **Sample sizes:** we considered four sample sizes: 500, 1000, 2500, and 5000. Those values were chosen since they represent between 1% and 10% of the buffer size, which can be considered reasonable values;
- **Classifier:** we selected the Incremental Support Vector Machine (ISVM) as the classifier because it is updated incrementally according to the arrival of new data. As its batch counterpart [6], ISVM calculates an optimal hyperplane between instances of two distinct classes, and the surrounding instances are assumed as support vectors [7].

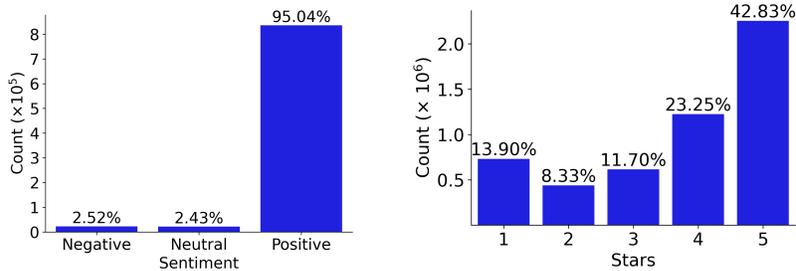
⁵ <http://insideairbnb.com/get-the-data>

⁶ Available at: <https://www.cntraveler.com/story/most-visited-american-cities>. Accessed on Jan 20th, 2024.

⁷ Available at: <https://huggingface.co/cardiffnlp/twitter-roberta-base-sentiment>.

⁸ <https://github.com/cristianomg10/methods-for-generating-drift-in-text-streams>

⁹ <https://www.yelp.com/dataset>



(a) Class distribution of Airbnb dataset after filtering. (b) Class distribution of Yelp dataset. Adapted from [25].

Fig. 1: Class distribution of the datasets used in this paper.

- **Hardware:** the hardware used in the experiments is a 13th Gen Intel(R) Core(TM) i9-13900K, 128 GB of RAM running Ubuntu 22.04 LTS, and 2 x GPU GeForce RTX 4090 (24GB).

4.2 Proposed Scenario

Considering the text stream mining setting, the proposed scenario is run as follows: (1) a text stream of length 200,000 sampled (stratified by class/label) from the original datasets; (2) the text stream classification is performed one-by-one; (3) a buffer accumulates the first 50,000 items of the text stream; (4) at the moment $t = 50,000$, a sampling method, among the methods presented in this paper, will sample a predefined number of items from those described in Section 4.1; (5) after sampling, the fine-tuning process for the language model, i.e., the SBERT model, is triggered using the sampled texts and the selected loss function; and (6) the evaluation metrics are calculated cumulatively in a test-then-train fashion, i.e., in a prequential manner.

This process was executed five times per sampling method and loss function. In each run, the entire stream was also sampled from the original dataset in a stratified manner. We used $t = 50,000$ as the effects of the fine-tuning and sampling methods would be easier to spot than at the end of the stream.

4.3 Loss Functions Settings

Considering the loss functions presented in Section 2.2, their inputs for fine-tuning were defined as follows:

- **BATL:** single texts and their respective classes;
- **CTL:** single texts;
- **OCL:** text pairs and a label, which we adapted for considering the distance between classes by calculating: $\text{label} = 1 - \frac{\text{abs}(X_1.\text{class} - X_2.\text{class})}{|\text{classes}|}$, where X_ϕ

is a text, $|\cdot|$ is the cardinality, and $\text{abs}(\cdot)$ is the absolute value. Thus, label is one if the classes are the same, indicating the similarity;

- **SL**: similarly to OCL, SL receives text pairs and a label, which we adapted to be the absolute distance, calculated as $\text{label} = \text{abs}(X_1.\text{class} - X_2.\text{class})$, where X_ϕ is a text, and $\text{abs}(\cdot)$ is the absolute value.

This paper leverages the pre-trained SBERT model `paraphrase-MiniLM-L6-v2`. When fine-tuning, we used 32-sized batches, 10 epochs, and 100 warmup steps, which are values frequent in the documentation.

4.4 Evaluation Metrics

Given the class imbalance of the datasets used in experimentation, results regarding Macro F1 Score and elapsed time were reported. In particular, the Macro F1 Score averages the harmonic mean of precision and recall obtained per class.

4.5 Results

Considering the presented scenario, loss functions, and sampling methods, the results obtained regarding Macro F1-Scores are demonstrated in Fig. 2. For readability issues, we kept different scales on the y-axis. The dashed lines correspond to the maximum and minimum Macro F1-Score values using SBERT without update and were used as a baseline. The x-axis regards the sampling methods, while the boxes correspond to the tested sample sizes. The Macro F1-Scores obtained for the baseline were: (a) 75.13 ± 0.28 (min: 74.86; max: 75.50) for the Airbnb dataset, and (b) 44.60 ± 0.12 (min: 44.46; max: 44.78) for Yelp dataset. Regarding the elapsed times, the values obtained were (in seconds): (a) 496.78 ± 0.70 (min: 495.89; max: 497.86) for the Airbnb dataset, and (b) 558.82 ± 3.19 (min: 553.36; max: 561.10) for Yelp dataset.

Considering the results obtained for the Airbnb dataset, we can see that the loss function is crucial for improving the results. CTL and OCL performed worse than when there were no updates, according to the dashed lines. Considering OCL and CTL, only the proposed WordPieceToken ratio sampling method (using class in the sampling, sample size 500, and CTL) performed equivalently to SBERT without update. Apart from this point, all the combinations for CTL and OCL performed worse than SBERT without update. Although sometimes equivalent to the competitors, the proposed WordPieceToken ratio (class) generally obtains higher averages than its peers with the same sample size, except for using the OCL. Furthermore, increasing the sample size for these loss functions degraded the results. We assume that these loss functions, together with the looseness of anchoring, ease the model to suffer from catastrophic forgetting.

Still analyzing the results for the Airbnb dataset, regarding BATL and SL, mostly in SL, all sampling methods were equivalent to SBERT without updates. However, using the Length or TF-IDF sampling method without accounting for the classes led to smaller Macro F1-Scores. In addition, it is possible to notice that, in the SL case, accounting for the classes in the sampling method

helped reach a better Macro F1-Score than the same sampling method without accounting for the class. In the BATL scenario, all methods perform similarly across the sample sizes, except for the WordPieceToken ratio, which obtained the highest Macro F1-Score using a sample size of 5000. Therefore, using a sample size of 5000 improved the performance in these cases.

Similar observations can be made by switching to the results concerning the Yelp dataset: CTL and OCL led to poorer results than BATL and SL. CTL provided the worst results for this dataset; the bigger the sample size, the worse the performance. Again, CTL seems to lead to catastrophic forgetting. It can be credited to its simple way of generating text pairs for fine-tuning. On the other hand, OCL obtained equivalent results across sampling methods and sample sizes, but all were worse than SBERT without updates. In addition, considering the sample sizes of 2500 and 5000, all methods showed increased performance.

BATL and SL functions led to increased performance for the Yelp dataset compared to the baseline (dashed lines). For BATL, all sampling methods using the sample size of 500 and 1000 are equivalent to the baseline. From the sample size of 1000, the At Random and WordPieceToken ratio sample methods reached Macro F1-Score values superior to the baseline. Regarding SL, most sampling methods are superior to the baseline from the sample size of 2500. Smaller samples led to decreased performance compared to the baseline.

Regarding the elapsed times, Fig. 3 shows measured values in each setting. Again, dashed lines correspond to the baseline, i.e., SBERT without update. For the Airbnb dataset using the BATL function, we noticed that the proposed WordPieceToken ratio took longer than the baseline to run, essentially from the sample size of 2500. For CTL, all methods, except for the length sampling method in both variations, took longer than the baseline. Specifically for Length (class), it was unstable, taking a reasonable time when using the sample size of 5000. It somehow can be expected since longer reviews would be selected, and the fine-tuning process would take longer. However, this should also be visualized for the Length sampling method, but it has not happened. We hypothesize that an anomaly in the GPU may have led to this increased variation. In OCL, all sampling methods have similar elapsed times: only sampling 5000 items led to higher run times than the baseline. At last, for SL, the Length- and TF-IDF-based methods took longer than the baseline from the sample size of 2500. In addition, Length (class) showed similar behavior to CTL regarding instability.

Regarding the Yelp dataset, considering the BATL function, random sampling led to shorter elapsed times. Length- and TF-IDF-based (with class) sampling methods led to longer elapsed times. Differently, in CTL, the proposed WordPieceToken ratio sampling reached the highest run times compared to all other methods. The same behavior happens in the SL function. For the OCL, the elapsed times for all methods are equivalent to the ones from SBERT without update, except for the WordPieceToken ratio sampling method.

Table 1 condenses the results obtained. The bold values are the best per dataset, sample size, and loss function (LF) combination, i.e., per row. The values in yellow and green are the best Macro F1 scores and elapsed times per

dataset/sample size. WordPieceToken ratio (class) obtained the best Macro F1-Scores in 5 out of 8 dataset/sample size pairs, and WordPieceToken ratio in 1 out of 8. Sampling at random was the fastest in 6 out of 8 dataset/sample size pairs. Furthermore, the best Macro F1-Scores increase with the sample size.

Although elapsed time is important and variations followed similar patterns, i.e., the higher the sample size, the longer the elapsed time, the differences between them show that they may not impede fine-tuning in text stream scenarios, depending on the hardware. The differences between maximum and minimum elapsed times are 187.07 seconds for Airbnb and 73.54 seconds for Yelp, which represent 37% and 13% of their respective average elapsed times.

5 Conclusion

Learning from text streams is challenging due to the constraints of text streams, such as time, resources, and one-pass processing [9, 10]. Furthermore, the existence of concept drifts in texts produced over time is well-known. A way to overcome textual concept drifts is by updating the language model. Updating (or fine-tuning) the language model is generally costly if all new data is considered.

This paper evaluates four different sampling methods, i.e., random sampling, length sampling, TF-IDF sampling, and WordPieceToken ratio sampling (proposed in this paper), where the latter three had a version that accounted for instances' classes, aiming at sampling important, informative texts from the buffer. Four loss functions were assessed in combination with the text sampling methods, i.e., Batch All Triplets loss (BATL), Contrastive Tension loss (CTL), Online Contrastive loss (OCL), and Softmax loss (SL).

We observed that the loss function plays a crucial role in improving the performance of the text classification task. Considering the scenarios assessed, CTL and OCL functions were insufficient to achieve satisfactory performance levels. On the other hand, BATL and SL functions aided in maintaining interesting performance levels, sometimes above the SBERT without update (our baseline), suggesting the SBERT can benefit from these loss functions. In addition, the elapsed times were comparable to the baseline. Therefore, BATL and SL were the most suitable functions for text stream classification among those evaluated in this paper. However, one must be careful when employing this approach because fine-tuning can create a bottleneck in real-time applications. Finally, text sampling methods are also crucial in fine-tuning. Our experiments suggest that the proposed WordPieceToken ratio method, especially leveraging the instances' classes, can retrieve more informative texts and favor the machine learning model's performance after fine-tuning.

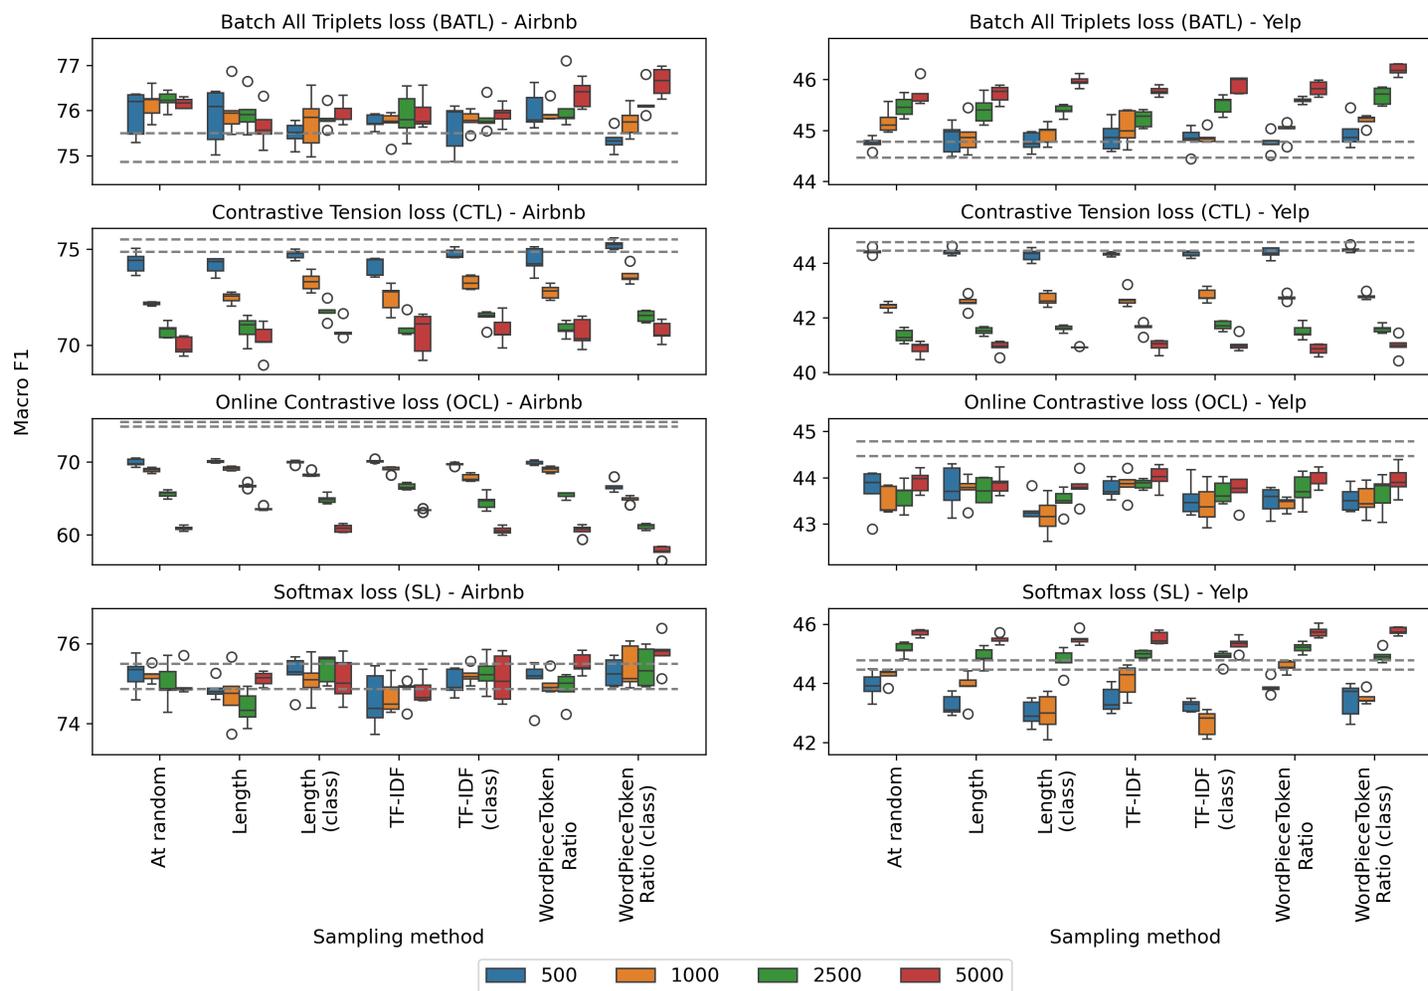


Fig. 2: Results for Airbnb (left) and Yelp (right). The dashed lines correspond to the maximum and minimum Macro F1-Score values using SBERT without update. The numbers in the legend correspond to the sample sizes.

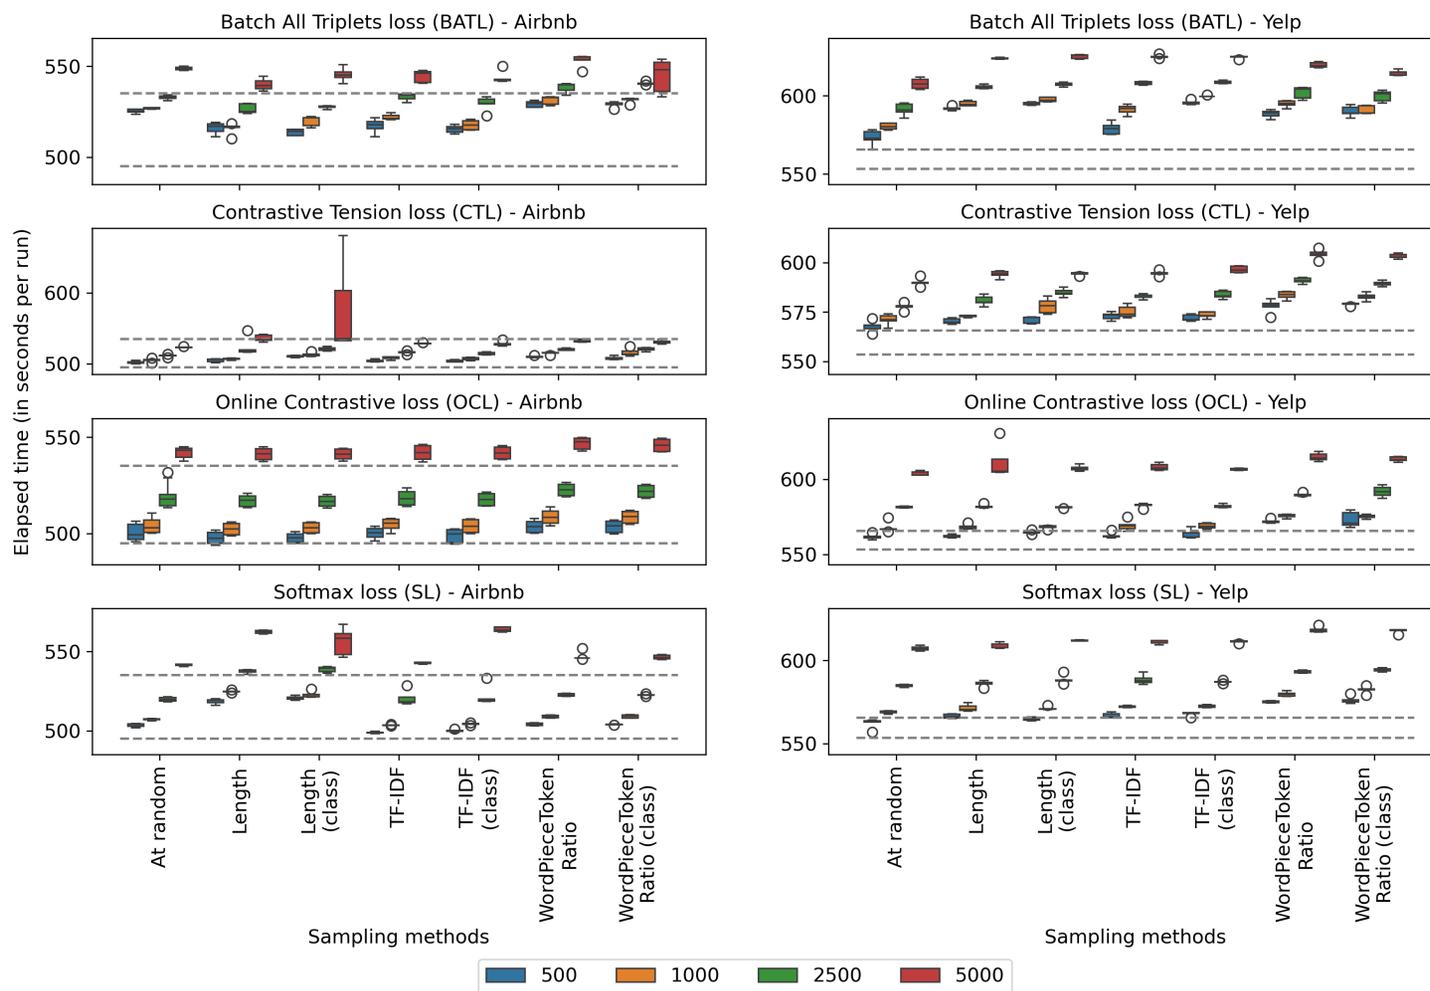


Fig. 3: Elapsed times for Airbnb (left) and Yelp (right). The dashed lines correspond to the maximum and minimum Macro F1-Score values using SBERT without update.

Table 1: Condensed results. Bold values are the best values obtained per row. Values in yellow (Macro F1) and green (elapsed time) are the best per dataset/sample size.

Dataset	LF	At random		Length		Length (class)		TF-IDF		TF-IDF (class)		WordPieceToken ratio		WordPieceT. (class)	
		Macro F1	Elaps. time	Macro F1	Elaps. time	Macro F1	Elaps. time	Macro F1	Elaps. time						
Airbnb (500)	BATL	75.94±0.51	525.29±1.09	75.86±0.63	516.03±3.20	75.49±0.27	513.74±1.74	75.75±0.16	517.22±3.97	75.63±0.55	515.42±2.17	76.01±0.43	529.26±1.71	75.34±0.25	529.00±1.55
	CTL	74.31±0.57	502.06±1.84	74.13±0.45	504.31±2.31	74.72±0.24	510.63±1.34	74.12±0.48	504.55±1.81	74.76±0.25	503.81±1.52	74.39±0.68	510.04±1.32	75.24±0.24	508.18±2.29
	OCL	69.92±0.49	500.51±4.21	70.13±0.22	497.78±3.27	69.93±0.26	497.78±2.52	70.12±0.21	500.30±2.81	69.70±0.21	498.81±3.43	69.95±0.28	503.60±3.05	66.71±0.74	503.60±3.23
	SL	75.24±0.44	503.34±1.12	74.86±0.25	518.56±1.62	75.25±0.47	520.60±1.19	74.58±0.72	498.92±0.48	75.03±0.33	499.92±0.75	75.06±0.57	504.14±0.86	75.29±0.36	503.97±0.27
Airbnb (1000)	BATL	76.15±0.35	526.95±0.48	76.00±0.53	515.73±3.28	75.74±0.62	519.87±2.88	75.70±0.32	522.05±1.64	75.79±0.22	517.72±2.71	75.95±0.21	531.35±2.33	75.75±0.33	531.44±1.53
	CTL	72.16±0.09	505.08±2.57	72.46±0.29	507.00±1.24	73.30±0.49	512.99±2.82	72.44±0.73	508.19±2.09	73.25±0.35	507.29±1.95	72.76±0.37	515.10±1.81	73.64±0.45	516.06±5.47
	OCL	68.91±0.32	503.90±3.84	69.12±0.25	502.31±3.06	68.31±0.37	502.92±2.82	68.97±0.43	504.76±3.00	67.86±0.49	503.85±3.45	68.96±0.39	508.66±3.74	64.88±0.46	508.48±3.26
	SL	75.23±0.20	507.20±0.53	74.71±0.70	524.81±0.69	75.09±0.51	522.96±2.10	74.67±0.44	503.55±0.41	75.21±0.23	504.27±0.56	75.00±0.27	508.90±0.87	75.42±0.55	509.07±1.04
Airbnb (2500)	BATL	76.23±0.21	532.92±1.24	75.96±0.44	526.52±2.68	75.84±0.24	527.48±0.85	75.90±0.51	533.20±2.16	75.85±0.33	529.35±4.03	76.10±0.57	538.23±2.76	76.20±0.35	540.58±0.78
	CTL	70.76±0.39	511.41±1.76	70.84±0.68	523.80±12.93	71.78±0.47	520.68±2.58	70.96±0.52	516.08±1.99	71.41±0.43	514.46±1.72	70.84±0.38	520.30±1.42	71.51±0.29	520.60±2.34
	OCL	65.61±0.47	519.29±6.45	66.71±0.31	516.95±2.99	64.91±0.59	516.74±2.84	66.63±0.41	518.33±3.89	64.56±1.05	517.66±3.32	65.46±0.42	522.70±3.29	61.16±0.37	521.83±3.31
	SL	75.00±0.53	519.83±1.29	74.40±0.42	537.47±1.17	75.39±0.36	538.80±1.83	74.83±0.33	520.59±4.74	75.25±0.44	521.87±6.34	74.89±0.40	522.76±0.90	75.41±0.50	522.59±0.59
Airbnb (5000)	BATL	76.16±0.12	548.77±0.93	75.66±0.44	539.97±3.28	75.94±0.26	545.38±3.80	75.95±0.38	544.53±3.41	75.91±0.23	543.80±3.49	76.37±0.31	552.95±3.40	76.64±0.32	544.69±9.47
	CTL	69.94±0.47	523.54±0.68	70.27±0.86	536.06±4.73	70.77±0.49	577.01±65.48	70.62±1.10	528.92±0.60	70.81±0.78	528.31±3.22	70.63±0.74	532.11±1.80	70.70±0.53	530.71±2.38
	OCL	60.90±0.29	541.87±2.84	63.61±0.24	541.25±3.14	60.90±0.52	541.18±2.90	63.37±0.18	541.97±3.90	60.62±0.50	541.88±3.27	60.61±0.73	546.76±2.96	57.73±0.75	545.78±3.48
	SL	75.02±0.38	541.39±0.54	75.12±0.17	562.32±1.08	75.10±0.57	556.26±8.92	74.83±0.34	542.58±0.51	75.13±0.61	563.87±1.56	75.50±0.27	546.98±2.86	75.77±0.45	546.37±1.39
Yelp (500)	BATL	44.75±0.12	573.14±5.07	44.85±0.31	591.92±1.27	44.78±0.19	594.89±0.87	44.89±0.30	579.11±3.89	44.83±0.25	595.63±1.35	44.77±0.19	588.47±2.56	44.96±0.31	590.06±3.50
	CTL	44.43±0.11	567.70±2.88	44.42±0.14	570.24±1.42	44.30±0.23	571.04±1.79	44.33±0.07	572.89±1.88	44.32±0.11	571.94±1.63	44.37±0.19	577.88±3.47	44.51±0.11	578.93±0.70
	OCL	43.73±0.50	561.79±1.79	43.77±0.49	562.14±1.05	43.33±0.28	564.83±1.13	43.76±0.21	562.65±2.03	43.55±0.39	564.05±2.98	43.50±0.31	572.15±1.28	43.54±0.27	573.24±5.18
	SL	43.93±0.46	562.21±2.95	43.27±0.35	566.43±1.35	42.98±0.44	564.91±1.02	43.45±0.46	567.24±1.46	43.25±0.20	567.83±1.32	43.89±0.25	575.07±0.55	43.43±0.60	576.29±2.29
Yelp (1000)	BATL	45.18±0.25	580.35±2.03	44.89±0.36	594.96±1.61	44.93±0.20	597.51±1.41	45.05±0.34	591.17±3.14	44.87±0.14	599.75±0.48	45.00±0.18	594.88±2.09	45.19±0.11	590.76±2.66
	CTL	42.42±0.16	571.09±2.78	42.56±0.26	572.81±0.45	42.69±0.25	578.06±3.92	42.70±0.31	575.08±3.13	42.83±0.25	573.45±1.60	42.74±0.11	583.48±2.05	42.79±0.12	582.68±1.84
	OCL	43.50±0.30	568.04±3.69	43.74±0.31	568.48±1.65	43.17±0.42	568.33±1.18	43.85±0.29	569.32±3.51	43.43±0.44	569.00±1.84	43.43±0.15	575.65±1.27	43.52±0.34	575.20±1.30
	SL	44.27±0.26	568.90±0.80	43.86±0.54	571.41±2.11	43.00±0.67	571.23±1.02	44.10±0.55	572.35±0.61	42.66±0.44	572.53±0.84	44.57±0.20	579.77±1.38	43.51±0.23	582.36±2.09
Yelp (2500)	BATL	45.47±0.21	591.20±3.90	45.40±0.28	605.69±1.19	45.39±0.11	607.26±1.19	45.23±0.17	608.03±1.13	45.50±0.19	608.62±0.88	45.59±0.06	602.05±3.93	45.68±0.17	599.80±3.57
	CTL	41.34±0.25	577.66±1.80	41.52±0.14	580.76±2.51	41.61±0.11	584.84±2.03	41.64±0.21	582.74±1.18	41.71±0.17	583.70±1.95	41.50±0.27	590.90±1.40	41.60±0.15	589.44±1.23
	OCL	43.60±0.31	581.54±0.52	43.73±0.27	581.98±1.18	43.50±0.26	581.37±0.40	43.86±0.11	582.63±1.47	43.68±0.26	582.26±1.07	43.74±0.35	589.63±1.17	43.65±0.41	591.88±3.68
	SL	45.16±0.23	584.80±0.75	44.90±0.33	585.97±1.66	44.75±0.42	588.60±2.74	44.98±0.14	588.49±2.93	44.88±0.23	587.16±0.77	45.20±0.18	593.36±0.75	44.93±0.22	594.54±1.03
Yelp (5000)	BATL	45.74±0.23	607.66±3.46	45.71±0.18	623.79±0.50	45.97±0.11	625.05±1.34	45.77±0.09	624.96±1.02	45.90±0.16	624.53±0.88	45.82±0.15	619.65±1.68	46.19±0.12	614.65±1.70
	CTL	40.88±0.27	590.04±2.02	40.93±0.23	594.08±1.76	40.92±0.02	594.33±0.79	40.98±0.23	594.59±1.29	41.04±0.27	596.56±1.70	40.84±0.20	604.23±2.38	40.98±0.37	603.21±1.27
	OCL	43.91±0.24	604.26±1.40	43.87±0.24	611.76±11.11	43.79±0.32	607.45±1.89	44.00±0.26	608.11±2.32	43.71±0.32	606.64±0.69	43.96±0.20	614.86±2.68	43.94±0.32	613.43±1.70
	SL	45.71±0.12	607.34±1.31	45.50±0.15	608.84±1.61	45.51±0.22	612.08±0.34	45.53±0.22	610.97±1.14	45.32±0.25	611.40±0.81	45.75±0.20	618.60±1.74	45.76±0.13	617.86±1.26

In future works, we intend to (a) extend to different stream mining tasks, (b) identify proper moments to automatically trigger the fine-tuning, (c) evaluate the balance between the sampling method and the resources consumption, and (d) evaluate other pre-trained BERT-based models with the combinations of loss functions and sampling method used in this paper.

Acknowledgements

Cristiano Mesquita Garcia is a grantee of a *Doutorado Sanduíche* scholarship provided by *Fundação Coordenação de Aperfeiçoamento de Pessoal de Nível Superior* (CAPES). This paper is also a byproduct of a CNPq project coordinated by Prof. Jean Paul Barddal under grant number 409371/2022-0 and is tied to CNPq International #441610/2023-4.

References

1. Amba Hombaiah, S., et al.: Dynamic Language Models for Continuously Evolving Content. In: Proceedings of the 27th ACM SIGKDD. pp. 2514–2524 (2021)
2. Barbieri, F., et al.: TweetEval: Unified Benchmark and Comparative Evaluation for Tweet Classification. In: Findings of the ACL: EMNLP (2020)
3. Bifet, A., et al.: Machine Learning for Data Streams: With Practical Examples in MOA. MIT Press (2023)
4. Bravo-Marquez, F., et al.: Incremental Word-Vectors for Time-evolving Sentiment Lexicon Induction. Cogn. Comput. pp. 1–17 (2022)
5. Carlsson, F., et al.: Semantic Re-tuning with Contrastive Tension. In: Int. Conf. on Learning Repr. (2020)
6. Cortes, C., Vapnik, V.: Support-Vector Networks. Mach. Learn. **20**, 273–297 (1995)
7. D’Andrea, E., et al.: Monitoring the public opinion about the vaccination topic from tweets analysis. Expert Systems with Applications **116**, 209–226 (2019)
8. Devlin, J., et al.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. . In: Proceedings of the 2019 Conf. of the North American Chapter of the ACL: Human Lang. Technol. (2019)
9. Gama, J., et al.: A Survey on Concept Drift Adaptation. ACM CSUR. **46**(4), 1–37 (2014)
10. Garcia, C.M., et al.: Concept Drift Adaptation in Text Stream Mining Settings: A Comprehensive Review. arXiv preprint arXiv:2312.02901 (2023)
11. Garcia, C.M., et al.: Event-driven Sentiment Drift Analysis in Text Streams: An Application in a Soccer Match. In: Proceedings of the 22nd Int. Conference on Machine Learning and Applications (ICMLA) (2023)
12. Harris, Z.S.: Distributional Structure. Word **10**(2-3), 146–162 (1954)
13. Henderson, M., et al.: Efficient Natural Language Response Suggestion for Smart Reply. arXiv preprint arXiv:1705.00652 (2017)
14. Hermans, A., et al.: In Defense of the Triplet Loss for Person Re-identification. arXiv preprint arXiv:1703.07737 (2017)
15. Joulin, A., et al.: Bag of Tricks for Efficient Text Classification. In: Proc. of the 15th Conf. of the Europ. Chapter of the ACL: Volume 2, Short Papers (2016)
16. Joulin, A., et al.: FastText.zip: Compressing Text Classification Models. arXiv:1612.03651 (2016)

17. Lee, J., et al.: BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* **36**(4), 1234–1240 (2020)
18. Liu, Y., et al.: RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint arXiv:1907.11692 (2019)
19. Pohl, D., et al.: Batch-based Active Learning: Application to Social Media Data for Crisis Management. *Exp. Syst. with Appl.* (2018)
20. Reimers, N., Gurevych, I.: Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In: *Proc. of the 2019 Conf. on Empir. Meth. in NLP. ACL* (2019)
21. Salton, G., Buckley, C.: Term-weighting Approaches in Automatic Text Retrieval. *Information Processing & Management* **24**(5), 513–523 (1988)
22. Schneider, E.T.R., et al.: CardioBERTpt: Transformer-based Models for Cardiology Language Representation in Portuguese. In: *2023 IEEE 36th Int. Symp. on Computer-Based Medical Syst. (CBMS)*. pp. 378–381. IEEE (2023)
23. Sharir, O., et al.: The cost of training NLP models: A concise overview. arXiv preprint arXiv:2004.08900 (2020)
24. Suprem, A., Pu, C.: ASSED: a Framework for Identifying Physical Events Through Adaptive Social Sensor Data Filtering. In: *Proceedings of the 13th ACM Int. Conf. on Distributed and Event-based Syst.* pp. 115–126 (2019)
25. Thuma, B.S., et al.: Benchmarking Feature Extraction Techniques for Textual Data Stream Classification. In: *2023 Int. Joint Conf. on Neural Netw. (IJCNN)*. pp. 1–8. IEEE (2023)