

Exploring Dimensionality Reduction Techniques in Multilingual Transformers

Álvaro Huertas-García^{a,*}, Alejandro Martín^a, Javier Huertas-Tato^a, David Camacho^{a,*}

^a*Departamento de Sistemas Informáticos, Universidad Politécnica de Madrid, Madrid, Spain*

Abstract

Background/introduction. Both in scientific literature and in industry, semantic and context-aware Natural Language Processing-based solutions have been gaining importance in recent years. The possibilities and performance shown by these models when dealing with complex Human Language Understanding tasks is unquestionable, from conversational agents to the fight against disinformation in social networks. In addition, considerable attention is also being paid to developing multilingual models to tackle the language bottleneck. The growing need to provide more complex models implementing all these features has been accompanied by an increase in their size, without being conservative in the number of dimensions required.

Methods. This paper aims to provide a comprehensive account of the impact of a wide variety of dimensional reduction techniques on the performance of different state-of-the-art multilingual siamese transformers, including unsupervised dimensional reduction techniques such as linear and nonlinear feature extraction, feature selection, and manifold techniques. In order to evaluate the effects of these techniques, we considered the multilingual extended version of Semantic Textual Similarity Benchmark (mSTSb) and two different baseline approaches, one using the embeddings from the pre-trained version of five models and another using their fine-tuned STS version.

Results and Conclusions. The results evidence that it is possible to achieve an average reduction of $91.58\% \pm 2.59\%$ in the number of dimensions of embeddings from pre-trained models requiring a fitting time $96.68\% \pm 0.68\%$ faster than the fine-tuning process. Besides, we achieve $54.65\% \pm 32.20\%$ dimensionality reduction in embeddings from fine-tuned models. The results of this study will significantly contribute to the understanding of how different tuning approaches affect performance on semantic-aware tasks and how dimensional reduction techniques deal with the high-dimensional embeddings computed for the STS task and their potential for other highly demanding NLP tasks.

Keywords: Dimensionality Reduction, Natural Language Processing, Semantic Textual Similarity, Multilingual Transformers, Language models.

1. Introduction

Natural Language Processing (NLP) includes various disciplines that provide a system with the ability to process and interpret natural language, just like humans use language as a communication and

*Corresponding authors.

Email addresses: `alvaro.huertas.garcia@upm.es` (Álvaro Huertas-García), `david.camacho@upm.es` (David Camacho)

reasoning tool [1]. Due to the recent increases in computational power, parallelization, the availability of large data sets, and recent advances in Artificial Intelligence, especially in the Machine Learning research field, NLP has been steadily proliferating and has garnered immense interest [1, 2]. In the very recent past, transformer-based architectures [3] have become an indispensable staple in the NLP field. Transformer models are able to capture latent syntactic-semantic information and encode text’s meaning as contextual vectors in a high-dimensional space referred as embeddings [2, 4]. In contrast to previous approaches, such as Statistical Natural Language Processing, the use of the attention mechanism provided by these architectures allows us to take into consideration a plethora of characteristics involved in human language.

The tremendous power of transformer-based models in Natural Language Understanding (NLU) and the new models that are continuously being proposed allow us to dramatically improve the state-of-the-art results in varied NLP tasks, including question answering, sentence classification, and sentence pair regression like Semantic Textual Similarity (STS) [2, 4, 5, 6]. The semantic evaluation performed in this STS task is one of the “levels of language” that determines the possible meanings of a sentence by focusing on the interactions among word-level [7]. It entails a high degree of complexity, given the large number of different linguistic features involved. In STS tasks, the systems need to compute how similar are two sentences considering all these features, returning a similarity score, usually ranging between 0 and 5 [8].

Regarding sentence-pair regression tasks, to overcome the massive computational overhead caused by the quadratic dependence on the input size of attention mechanism in transformers models [5, 9], the use of siamese architectures is a very effective method for deriving semantically meaningful sentence embeddings in an efficient way [5, 9]. This approach is also called dual-encoder, bi-encoder or siamese architectures. As explained by Humeau et al. [9], the training of a siamese architecture consists of two pre-trained transformer-based models with tied weights that can be fine-tuned for a specific task like computing separately semantic embeddings for a pair of sentences and measure their similarity using the extensively used cosine similarity function [10].

Despite the practical features of cosine similarity, such as symmetry and spatial interpretation, this similarity metric has complexity $O(N)$: time and memory grow linearly with the number of dimensions of the vectors compared [11]. Thus, dimensionality is a bottleneck for similarity computation and embedding storage [12]. Moreover, an increasing number of studies using ensemble approaches based on the concatenation of embeddings can be found in the literature [13, 14, 15], aiming to improve the results in state-of-the-art tasks but accentuating this issue. Given that the application of dimensionality reduction techniques can mitigate this bottleneck, it requires further exploration.

Although in the history of NLP, the focus has mainly been on proposing architectures for English tasks. Nevertheless, interest in developing multilingual NLP tools has grown recently to achieve diversity for transferring NLP capabilities to low-resource languages for which labelled (and unlabelled) data is scarce [16, 17]. The incorporation of tasks in various languages in the SemEval [18] and CLEF [19] competitions is clear evidence of this, surmounting the language bottleneck, but also the increasing number of newly introduced multilingual models [19], such as the recent BLOOM (BigScience Language Open-science Open-access Multilingual) [20] open-sourced model with 176B parameters generated by BigScience which has marked an inflexion point on the research and development of large language models (LLMs).

The present paper seeks to address how different dimensionality reduction techniques impact the performance of pre-computed embeddings from multilingual transformer-based models trained using a siamese architecture focusing on semantic similarity by employing four different approaches where the reduction techniques are compared with the pre-trained and fine-tuned versions of the models.

This research aims to expand the current knowledge further on the effect of dimensionality reduction techniques in Natural Language Processing [21, 22, 23, 24]. The following are the new contributions of our research concerning previous studies:

- A more comprehensive range of unsupervised dimension reduction techniques is included, from linear and nonlinear feature extraction to feature selection and manifold techniques.
- The effect of these techniques on pre-trained and fine-tuned precomputed embeddings in the Semantic Textual Similarity (STS) task is explored.
- In contrast to previous work that explored reduction techniques in classical static word embeddings, this paper investigates the effect of dimension reduction in state-of-the-art contextual-based transformer models.
- Unlike previous works focused on the English language, this research analyses multilingual models, overcoming the language bottleneck for the applicability of dimension reduction of the embeddings of these models.

The remainder of the article is organized as follows: Section 2 outlines previous related work on dimensionality reduction techniques, their application in Deep Learning, and the importance of multilingual semantics in NLP. Section 3 describes the approaches followed in this research to evaluate the impact of dimensionality reduction techniques, the multilingual models and the dimensionality reduction techniques applied, and Section 4 the data and processes used to fit and evaluate these techniques. The experimental results are discussed in Section 5. Finally, Section 6 summarizes the results of this work and draws conclusions.

2. Related Work

2.1. Dimensionality Reduction Techniques

Dimensionality reduction techniques aim to reduce the dimensionality of the data, removing irrelevant and redundant features while preserving critical information for subsequent applications, such as facilitating their visualization and understanding or leading to more compact models with better generalization ability [25, 26].

There are different non-mutual exclusive criteria to classify dimensionality reduction techniques. Firstly, according to the reduction approach, these techniques can be classified into two categories: feature selection and feature extraction techniques [27]. Feature selection involves selecting a subset of original features that are useful for building models, effectively removing some of the less relevant features from consideration [28, 26]. On the other hand, feature extraction transforms the original data into another feature space with certain criteria, creating new variables by combining information from the original variables that capture much of the information contained in those original variables [27, 25]. Additionally, feature extraction methods can be further subdivided into linear and nonlinear according to the variable combinations applied [29].

Secondly, according to the information available in the datasets, dimensionality reduction techniques can be classified as supervised and unsupervised [25, 26]. Supervised techniques require each data instance in the dataset to be labelled accordingly to the task, whereas unsupervised techniques are task-agnostic approaches and do not require labelled data.

2.2. Dimensional reduction of embeddings

In broad terms, embeddings can be categorised as pre-trained or downstream fine-tuned embeddings [30], and as pre-computed or on-the-fly embeddings [31, 32].

The first criteria to categorise embeddings is whether they come from pre-trained models for general tasks or are task-specific. Pre-trained embeddings are widely used as a starting point for downstream applications. A clear example is the current ‘Pre-train and Fine-tune’ paradigm of transformer models [30].

Training these models from scratch is prohibitively expensive in many cases. Alternatively, using self-supervised trained checkpoints of these models and their pre-trained embeddings as a starting point to be later fine-tuned for supervised downstream tasks is widely used. Unlike previous works in the literature that have only focused on reduced pre-trained embeddings [21, 22, 23], in this work we are interested of the evaluation of the impact of dimensionality reduction on both types of embeddings, pre-trained and downstream fine-tuned embeddings.

Pre-computed embeddings in NLP are widespread, i.e., embeddings that may or may not be adjusted to a task but generated beforehand and not at each use time. A straightforward application of pre-computed embeddings is the semantic search NLP task for locating relevant information from massive amounts of text data [33]. Semantic search task uses semantic textual similarity to compare an input text against a large set of texts from a database to extract relevant related information—typically, a distance metric such as cosine similarity ranks which content should be extracted to an input query. However, computing the database embeddings each time a query is introduced is infeasible. Alternatively, it is preferable to compute the embeddings once, store them, and use these pre-computed embeddings for subsequent requests [32]. With this in mind, it is important to note the usefulness of reducing embeddings dimensions, which can improve their utility in memory-constrained devices and benefit several real-world applications.

As Camastra & Vinciarelli mention [34], the use of more features than strictly necessary leads to several problems, pointing out that one of the main problems was the space needed to store the data. As the number of available information increases, the compression for storage purposes becomes even more critical [12]. Additionally, for the scope of this work, it cannot be ignored that the application of dimensional reduction techniques for reducing pre-computed embeddings dimensions neither improves the runtime nor the memory requirement for running the models. It only diminishes the needed space to store embeddings and increases the speed to make computations (i.e. to calculate the cosine similarity between two vectors), which also contributes to decrease the considerable impact on the energy and carbon footprints generated during the production use of the models when pre-computed and stored embeddings are required [35]. Research has tended to focus on implementing bigger and complex models rather than analysing methods to adjust the vector space to the desired task while being conservative on the number of dimensions required [35]. An additional problem is that the storage of high-dimensional embeddings is challenging when dealing with large volume datasets [32]; besides, researchers have already raised awareness of the storage limitation we are about to face if current technology is adopted and the storage utilization growth rate persists [12].

In terms of performance, dimensionality reduction techniques can also contribute. The performance of Machine Learning (ML) and, in particular, Deep Learning (DL) models is highly dependent on the choice of data representation (or features) to which they are applied [36]. For that reason, much effort during the deployment of ML and DL solutions is dedicated to obtaining a data representation that can support effective learning. As a result, different fields where dimension reduction techniques are combined with ML and DL complex models can be found in the literature. For example, from their application in time series forecasting [37] or health sciences for cancer classification [38] to Natural Language Processing (NLP) to improve the clustering of text documents [39] or sentiment classification of opinion texts [17].

Recently, the study of Raunak et al. [21, 22] has shed more light on the importance of reducing the size of embeddings produced by ML and DL models in the field of NLP. More specifically, these authors draw attention to reducing the size of classical GloVe [40] and FastText [41] pre-trained word embeddings using PCA-based post-processing algorithms, achieving similar or even better performance than the original embeddings.

Other works on the potential of reducing pre-computed embeddings dimensions have been carried out [23] exploring the effect of Principal Components Analysis (PCA) [42, 43] and Latent Semantic Analysis (LSA) [24] dimensionality reduction techniques as a post-processing step of pre-trained GloVe word embeddings for text classification tasks. These authors also corroborated the usefulness of PCA for

obtaining more accurate results with lower computational costs concluding that the PCA method is more suitable than LSA for dimensionality reduction. In the same way, Shimomoto et al. [44] propose solving topic classification and sentiment analysis by using PCA to transform pre-computed word embeddings of a text into a linear subspace. Their results showed the effectiveness of using the PCA subspace representation for text classification. This fact has already been proved by other authors [45] showing that the storage, memory, and computation required by these large embeddings typically results in low efficiency in NLP tasks, pointing out the importance of methods to compress pre-computed and pre-trained GloVe and FastText embeddings.

Additionally, researchers have explored dimensional reduction techniques for visualizing semantic embeddings of small text corpora [46]. The authors explored four dimension reduction strategies on pre-computed embeddings based on PCA and t-distributed stochastic neighbor embedding (t-SNE) [47], concluding that both methods preserved a significant amount of semantic information in the full embedding.

Similarly, the use of dimension reduction techniques is likewise interesting in Semantic Similarity. As discussed previously, in the Semantic Similarity task, the linear $O(N)$ complexity of cosine similarity is one of the reasons why this distance metric is widely used in the community and this study. However, the complexity of cosine similarity, although linear, is a limitation when dealing with a large number of dimensions and a large amount of data. Precisely, in tasks where cosine similarity is used as a distance metric, the large amount of data handled and the high number of dimensions of the embeddings generated by the models used represent a constraint in both computational time and storage that leaves the door open to the use of dimension reduction techniques.

To the best of our knowledge, in the literature, dimension reduction research on embeddings has focused on using classical pre-computed word embeddings, including the popular GloVe or FastText embeddings [21, 22, 23, 24, 46]. These classical word embeddings are static, word-level, and contextual-independent, and their main limitation is that they do not consider what context the word is being used. Moreover, the variety of embedding techniques explored is limited, focusing mainly on PCA. Likewise, these studies do not include multilingualism in their analyses, being limited to the English language.

Hence, the presented study follows the research line proposed by different authors [21, 22, 23, 24] but takes a step forward, including a broader range of techniques and evaluating the capability of **dimensionality reduction techniques** in both **pre-trained and fine-tuned pre-computed embeddings** from state-of-the-art contextual-based transformer models from the recently claimed **multilingual** point of view.

2.3. Siamese and Non-Siamese architectures

Training in non-siamese transformer architectures, known as cross encoders, requires that two sentences be passed to the transformer as input while a target value is predicted [9]. This approach requires feeding the model with both sentences, which causes massive computational overhead because the attention is quadratic to the sequence length. As previously described in the Introduction (Section 1), siamese architectures are the main alternative to cope with this massive computational overhead in transformers models. As Reimers & Gurevych [5] measured, in a V100 GPU, finding the pair with the highest similarity in a collection of $n = 10000$ sentences for a non-siamese model requires about 65 hours. On the other hand, the same task with a model trained with a siamese architecture is reduced to ~ 5 seconds.

Although it should also be noted that models trained in a non-siamese way generally obtain better results [48], since they generate the embeddings considering the interaction of both sentences, the type of architecture only affects the computational time to obtain the embeddings. Consequently, the architecture type does not affect the size of the computed embeddings, which makes dimension reduction techniques equally interesting and valuable for both architectures.

Therefore, it is noteworthy that the present research pays special attention to the dimension reduction of the embeddings of transformers models trained with siamese architecture because, in Semantic Textual

Similarity (STS) tasks, they are a widely used solution that obtains results in a reasonable time and at a feasible computational cost. However, these reduction techniques are equally helpful for other types of architecture.

2.4. Importance of Multilingual Semantics

Semantics has many applications in a wide range of domains and tasks. Recent developments regarding Information Retrieval tasks [33, 49, 48] have demonstrated the potential of combining semantic-aware models along with traditional baseline algorithms (e.g., BM25) [50]. Moreover, the use of semantic-aware models has proven to be an excellent approach to counteract informational disorders (i.e., misinformation, disinformation, malinformation, misleading information, or any other kind of information pollution) [51, 52, 53, 54] or to build automated fact-checking approaches [55]. Additionally, semantic similarity can be applied to organize data according to text properties, which is formally an unsupervised thematic analysis [56]. Following the same criterion, the measurement of semantic similarity between a sentence and each word in the sentence can be applied to extract the keywords with the highest semantic content from the sentence [57]. All these applications rely on measuring semantic textual similarity (STS), making STS a crucial task in NLP.

A significant limitation of these semantic-aware solutions is the language bottleneck [58]. Language constitutes one of the most significant barriers to be addressed since a model’s ability to handle multiple languages is essential for its widespread applications. In recent years, in the field of NLP, attention has been paid to multilingual models to achieve diversity in transferring NLP capabilities to low-resource languages [16, 18, 19]. Therefore, in this paper, special attention has been paid to multilingual transformers models, with the immediate goal of covering embedding dimensionality reduction in the semantic textual similarity task in the world’s most widely spoken languages, even though these techniques are relevant for monolingual and multilingual transformers.

Altogether, this work aims to broaden our knowledge of semantic-aware transformer-based models by analysing the impact of different dimensionality reduction techniques on the performance of multilingual siamese transformers on semantic textual similarity multilingual tasks. The results of this study will contribute significantly to understanding how different tuning approaches affect performance on semantic-aware tasks and how dimensional reduction techniques deal with the high-dimensional embeddings computed for the STS task from the recently claimed multilingual point of view.

3. Methodology

The main goal of this research lies in providing a deep analysis of the use of different dimensionality reduction techniques to reduce the size of the output embedding of different multilingual transformer models and their impact in the performance.

The following subsections describe in detail the techniques explored to reduce embeddings, present the multilingual transformer models included in this study and the different approaches applied to quantify the reduction margin and its effect on the performance.

3.1. Dimensionality Reduction Techniques

As previously mentioned, dimensionality reduction techniques can be grouped according to two non-mutual exclusive criteria. In this paper, we have included a whole range of types of dimensionality reduction techniques, including linear and nonlinear feature extraction and feature selection techniques. Nevertheless, since the transformers models included in this study are employed in a siamese architecture to determine the degree of similarity between a pair of sentences, they output a pair of non-concatenative embeddings between which the similarity is estimated using the cosine distance. Hence, for each labelled similarity score, there are two non-concatenative embeddings. For this reason, even though we have labelled the data, only unsupervised methods are explored.

The dimensionality reduction techniques explored in this project are:

- **Principal Component Analysis (PCA):** Principal Component Analysis [42, 43] is a powerful unsupervised linear feature extraction technique that computes a set of orthogonal directions from the covariance matrix that capture most of the variance in the data [59]. This is, it creates new uncorrelated variables that maximize variance, and at the same time, most existing structure in the data is retained. It is also important to note that this research uses a variant of PCA known as Incremental Principal Components Analysis (IPCA) [60]. This variant follows the same basic principles as PCA. However, it is much more memory efficient, as it applies PCA in batches, avoiding storing entire data in memory and allowing PCA to be applied on large datasets.
- **Independent Component Analysis (ICA)** [61]: Independent Component Analysis is an unsupervised feature extraction probabilistic method for learning a linear transformation to find components that are maximally independent between them and non-Gaussian (non-normal), but at the same time, they jointly maximize mutual information with the original feature space.
- **Kernel Principal Components Analysis (KPCA)** [62]: Kernel-based learning method for PCA. It makes use of kernel functions to construct a nonlinear version of PCA linear algorithm by first implicitly mapping the data into a nonlinear feature space and then performing linear PCA on the mapped patterns [59]. The kernels considered in this project are the Polynomial, Gaussian RBF, Hyperbolic Tangent (Sigmoid), and Cosine kernels.
- **Variance Threshold:** Unsupervised feature selection approach that removes all features with a variance below a threshold. Indeed, this technique selects a subset of features with large variances, considered more informative, without considering the desired outputs.
- **Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP):** The authors of UMAP [63] describe it as an algorithm that can be used for unsupervised dimension reduction based on manifold learning techniques and topological data analysis. In short, it first embeds data points in a new nonlinear fuzzy topological representation using neighbor graphs. Secondly, it learns a low-dimensional representation that preserves the maximum information of this space, minimizing Cross-Entropy. Compared to its counterparts, such as t-SNE, UMAP is fast, scalable, and allows better control of the desired balance between the local and global structure to be preserved. Two main parameters play a vital role in controlling this: (1) the number of sample points that defines a local neighborhood in the first step, and (2) the minimum distance between embedded points in low-dimensional space to be clustered in the second step. Larger values of the number of neighbors tend to preserve more global information in the manifold as UMAP has to consider larger neighborhoods to embed a point. Likewise, larger minimum distance values prevent UMAP from packing points together and preserve the overall topological structure.

According to the previous preprocessing steps required before dimensionality techniques, it should be noted that PCA and KPCA assume a Gaussian distribution, and the features must be normalized; otherwise, the variances will not be comparable. Therefore, the StandardScaler is applied beforehand. Regarding ICA, non-Gaussian distribution is assumed, and the data is already whitened by the algorithm, so no previous preprocessing step is necessary. For the Variance Threshold, the best standardization method is MinMaxScaler, as it transforms all features to the same scale but does not alter the initial variability. This allows the variance selection threshold set to affect all dimensions equally. Finally, since there is no Gaussian assumption under UMAP and the cosine distance calculation benefits from scaling the features to a given range, MinMaxScaler is applied before UMAP. A summary about the necessary considerations about the above scaling steps and the characteristics of the dimensionality reduction techniques applied in this project are listed in Table 1.

	PCA	KPCA	ICA	Variance Threshold	UMAP
Preprocessor	Standard	Standard		MinMax	MinMax
Scalation	✗	✗	✗	✗	✗
Normalization	✗	✗			
Unsupervised	✗	✗	✗	✗	✗
Feature Selection				✗	
Feature Extraction	✗	✗	✗		✗
Linear	✗		✗		
Non Linear		✗			✗

Table 1: Considerations about the previous scaling steps and the characteristics of the different dimensionality reduction techniques applied in this project.

Technique	Parameters
ICA	random_state = 0 max_iter = 320 whiten = True tol = 5e-4
	kernels = [sigmoid, polynomial, rbf, cosine] eigen_solver = arpack copy_X = False random_state = 0
Variance Threshold	threshold = [Min, Max, Decile of variance]
UMAP	precomputed_knn = True metric = cosine min_dist = 1 n_neighbors = [5, 10, 50, 100, 125] angular_rp_forest = True

Table 2: Parameters with non-default values used in the previous scaling steps and the dimensionality reduction techniques applied in this project.

Finally, it is worth mentioning that these dimensionality reduction techniques and preprocessing algorithms come from *scikit-learn* v1.0.2 [64], except for UMAP that belongs to *umap-learn* v0.5.2 [63]. For the sake of reproducibility, the different parameters and values used in the experiments are presented in Table 2. Finally, the variance threshold filters for the Variance Threshold technique tested are extracted by previously calculating the variance of each feature (i.e., 768 variances for 768 embedding dimensions), extracting the deciles, and including the maximum and minimum of these variances.

3.2. Multilingual Models

Regarding the models included, we have tried to include the open-source architectures widely used by the NLP community, such as the classic BERT and its robust version of multilingual nature, such as XLM-RoBERTa. The effects of dimensionality reduction and fine-tuning process were explored in the following pre-trained multilingual models extracted from Hugging Face [65]:

- **bert-base-multilingual-cased:** BERT [4] transformer model pre-trained on a large corpus of 104 languages Wikipedia articles using the self-supervised masked language modeling (MLM) objective with $\sim 177M$ parameters.
- **distilbert-base-multilingual-cased:** Distilled version of the previous model, being on average twice as fast as this model, totalizing $\sim 134M$ parameters [66].
- **xlm-roberta-base:** Base-sized XLM-RoBERTa [67] model totalizing $\sim 125M$ parameters. XLM-RoBERTa is RoBERTa model [68], a robusted version of BERT, pre-trained on CommonCrawl data containing 100 languages.
- **xlm-roberta-large:** Large-sized XLM-RoBERTa [67] model totalizing $\sim 355M$ parameters.
- **LaBSE:** Language-agnostic BERT Sentence Embedding [69] model trained for encoding and reducing the cosine distance between translation pairs with a siamese architecture based on BERT, a task related to semantic similarity. It was trained over 6 Billion translation pairs for 109 languages. The authors also reported that it has zero-shot capabilities as was able to produce decent results for other not seen languages.




Figure 1: Representation of the approaches followed to evaluate the impact of different dimensionality reduction techniques in multilingual transformers. Where T represents the pre-trained transformer model, T' the fine-tuned transformer model, R' the fitted dimensionality reduction technique, E_o the original output embedding, and E_r the reduced embedding.

Transformer-based models embed textual information into vectors of high dimensionality. The pre-trained multilingual models used in this study generate embeddings with 768 dimensions, 1024 in case of *xlm-roberta-large*. These default dimensions are the ones considered to be reduced since a different number of dimensions from the default would entail losing the knowledge acquired during the self-supervised pre-training phase (e.g., mask word prediction), in which we are interested in studying the effects of the dimension reduction techniques.

3.3. Evaluation Approaches

We have followed four approaches, aiming to evaluate and quantify the reduction margin and its effect in the performance while using different methodologies. These four approaches are shown in Figure 1.

We have included two different baseline approaches, one using the embeddings from pre-trained version of the five previous mentioned models (i.e., Approach 1) and another using their fine-tuned STS version (i.e., Approach 2). These baselines will allow us to have a standard to examine the dimensionality reduction capability and its effect on STS task performance.

For each baseline, an approach that applies dimensionality reduction on the generated embeddings is used. Consequently, Approach 3 applies dimension reduction using Approach 1 as the baseline. Similarly, Approach 4 uses Approach 2 as the baseline.

These four approaches are described as follows:

- **Approach 1 - Pre-trained models.** In the first approach we employ and directly evaluate the pre-trained models in the mSTSb test split without applying any dimensionality reduction. This approach is used as the baseline for Approach 3.
- **Approach 2 - Fine-tuned models.** In this second approach, the pre-trained models are fine-tuned downstream using the mSTSb train split and evaluated in the mSTSb test split without applying any dimensionality reduction technique. This approach is used as the baseline for Approach 4. The fine-tuning process will be discussed in more detail in Section 4.4.
- **Approach 3 - Reduced embeddings from pre-trained models.** In this approach, the embeddings generated by the pre-trained models from Approach 1 in the mSTSb train split are used to fit the different dimension reduction techniques and evaluate them in the mSTSb test split. Thus, an

analysis between the results achieved in Approach 1 and Approach 3 will help to understand the impact of dimensionality reduction techniques in the embeddings from pre-trained models

- **Approach 4 - Reduced embeddings from fine-tuned models.** This approach is equivalent to Approach 3 but uses the fine-tuned models in Approach 2, allowing us to assess the impact of dimensionality reduction techniques in fine-tuned embeddings.

The experimental design of these approaches will be discussed in more detail in Sections 4.3 and 4.5.

4. Experimental Setup

4.1. Data

The multilingual extended STS Benchmark (mSTSb) [48] train set is used for fine-tuning the multilingual transformers and fitting the variety of dimensional reduction techniques. This split consists of 16 languages¹ combined in 31 mono and cross-lingual tasks with 5,479 pair of sentences each one. Likewise, mSTSb test set is used to evaluate the performance of the models obtained from the different approaches. The mSTSb test set is also composed of 31 multilingual tasks with 1,379 pair of sentences per task.

To evaluate the performance in mSTSb, the sentence embeddings for each pair of sentences are computed and the semantic similarity is measured using the cosine similarity metric. Then, the Spearman correlation coefficient (ρ or r_s) is computed between the scores obtained and the gold standard scores, as it is recognised as an official metric used for semantic textual similarity tasks [70, 71].

It is important to note that the mSTSb data variety available for fitting (i.e., train split) totals +183k sentences (i.e., 16 languages with 5,749 pair of sentences each one). For linear PCA, this dataset is too large to fit in memory. To manage this situation, an Incremental PCA (IPCA) approach [60] is applied. As previously mentioned, IPCA simply fits the PCA in batches, being independent of the number of input data samples, but still dependent on the input data features.


Similarly, KPCA and UMAP are computationally more expensive than the linear counterparts [62, 59]. For this reason, these dimensionality reduction techniques were fitted using a subset of 10k pair of sentences (i.e., 20k sentences), ensuring always the number of data instances is larger than the number of dimensions. To perform this subsampling, the following requirements were taken into account: (1) all 16 languages must be equally represented, giving a total of 625 sentence pairs for each language; (2) all sentences present in the original train split will be present at least once in some language; (3) the representation of the different sentence pairs must be as random as possible. Following these criteria, we perform a sampling based on assigning sentences to a randomly selected language until we reach the maximum number of sampled data. To avoid bias in the order in which the sentences are assigned to the languages, the different sentence pairs are shuffled randomly at each iteration. As each language reaches the maximum data, that language is discarded. This ensures a random distribution of samples in each language but includes the full range of sentences present in the original train data.

4.2. Computational Resources

Regarding the computational resources used in this study, an Intel(R) Xeon(R) Bronze 3206R CPU at 1.90GHz is used to fit the reduction techniques in Approaches 3 and 4. On the other hand, the pre-computed embeddings used in the approaches and the fine-tuning experiments from Approach 2 are performed using a Quadro RTX 8000 48GB GPU.


¹ ar, cs, de, en, es, fr, hi, it, ja, nl, pl, pt, ru, tr, zh-CN, zh-TW

APPROACH 1




APPROACH 2

Step 1 - Fine-tuning




Step 2 - Test




APPROACH 3

Step 1 - Dimensionality Reduction Fitting




Step 2 - Test



APPROACH 4

Step 1 - Dimensionality Reduction Fitting



Step 2 - Test




Figure 2: Diagram of the methodology followed for testing the different Approaches on multilingual STS benchmark. Where T represents the pre-trained transformer model, T' the fine-tuned transformer model, R the dimensionality reduction technique, R' the fitted dimensionality reduction technique, V_i the cosine similarity score computed for i th sentence pair, G_i the gold standard similarity score for i th sentence pair, MSE for Mean Squared Error loss, and r_s the Spearman correlation coefficient.

4.3. Baseline Approaches: Approach 1 and Approach 2

As can be seen in Figure 2, baseline Approaches 1 and 2 share similar experimental designs. Both employ the mSTSb test data split to generate the embeddings of the sentence pairs. As mentioned previously, because they are siamese-trained models, one embedding is computed for each sentence in the sentence pair. Then their similarity value is calculated with the cosine distance and the Spearman coefficient is computed using the reference values to obtain the baseline performance for each multilingual model. The only difference between Approach 1 and Approach 2 is the version of the model used. In the case of Approach 1 it is the pre-trained version of the model, while in Approach 2 it is the fine-tuned version in the mSTSb task. This fine-tuning process is explained in more detail in the next subsection.

4.4. Transformers fine-tuning

The fine-tuning process of the models based on the siamese training strategy for Approach 2 was carried out following the methodology described by Reimers et al. [5, 58]. As depicted in the fine-tuning process of Approach 2 in Figure 2, for each sentence pair from the mSTSb training split, two networks with tied weights from the same transformer model compute the embeddings for each sentence separately. Then the cosine similarity between the two sentence embeddings is calculated. Finally, the mean squared error (MSE) loss between the predicted and the gold similarity scores is used as the objective function to update the tied weights. During training, the following hyperparameters were optimized: number of epochs, scheduler, weight decay, batch size, warmup ratio, and learning rate. The hyperparameter values explored, the required time for fine-tuning, and the results of the experiments can be consulted in Table 8 and Weight and Biases ².

4.5. Dimensionality reduced techniques fitting: Approach 3 and Approach 4

Figure 2 provides an overview of the process of applying and evaluating dimensional reduction techniques in Approach 3 and Approach 4. Both approaches follow the same methodology; the only difference is the pre-trained or fine-tuned version of the model used to compute the embeddings for fitting the dimensional reduction technique. The experimental design has two steps and it is repeated for each number of dimensions explored for each technique and each model. A wide range of the number of dimensions is explored for each dimensionality reduction technique, as shown in Section 5.

Firstly, a dimensionality reduction technique is fitted through the embeddings computed by a multilingual model using the mSTSb train split. As these are unsupervised techniques (i.e., the gold similarity scores are not included), this training step does not consider the relationship between pairs of sentences but uses the individual sentences to fit the technique. In other words, if the train split of mSTSb contains 10k pair of sentences from 16 languages, 20k separately sentences are used to fit a dimensional reduction technique.

Finally, the embeddings from paired sentences from the mSTSb test split are reduced with the fitted technique, then the cosine similarity distance function is applied to obtain a similarity value, and finally, the Spearman correlation coefficient to the gold similarity scores is computed. In contrast to the previous fitting step, in this second step the relationship between pairs of sentences is considered as we need to evaluate the techniques on mSTSb task.

4.6. Statistical comparison

Additionally, to test if the use of reduced embeddings has a significant impact on the performance in comparison to the baseline approaches, we compare the average Spearman correlation coefficient of the five multilingual siamese transformer models (see Section 3.2) between each pair of baseline and reduced

²https://wandb.ai/huertas_97/PaperDimRed

Model	Ap. 1 r_s	Best Technique	Dimensions	Ap. 3 r_s	Fitting Time
bert-base-multilingual-cased	0.4342	ICA	209	0.5019	4m 16s
distilbert-base-multilingual-cased	0.4531	ICA	169	0.523	2m 47s
xlm-roberta-base	0.3274	ICA	249	0.5269	7m 51s
xlm-roberta-large	0.2855	ICA	1024	0.5392	31m 22s
LaBSE	0.7096	ICA	129	0.7488	2m 27s

Table 3: Average Spearman r_s correlation coefficient comparison between Approach 1 (Ap. 1) and best dimensional reduction technique in Approach 3 (Ap. 3) for the multilingual Transformers.

Model	Ap. 2 r_s	Best Technique	Dimensions	Ap. 4 r_s	Fitting Time
bert-base-multilingual-cased-fine-tuned	0.7045	ICA	568	0.7117	12m 38s
distilbert-base-multilingual-cased-fine-tuned	0.6863	VarThres	692	0.6842	2s
xlm-roberta-base-fine-tuned	0.7470	VarThres	673	0.7495	3s
xlm-roberta-large-fine-tuned	0.8150	KPCA-sigmoid	1024	0.8176	20m 6s
LaBSE-fine-tuned	0.8242	KPCA-sigmoid	768	0.8243	19m 25s

Table 4: Average Spearman r_s correlation coefficient comparison between Approach 2 (Ap. 2) and best dimensional reduction technique in Approach 4 (Ap. 4) for the multilingual Transformers.

approaches (i.e., Approach 1 vs Approach 3, Approach 2 vs Approach 4). For this purpose, as we are comparing the same set of models in different approaches, the two-tailed paired T-test using a significance level of 0.05 is conducted to test the null hypothesis of identical average Spearman correlation coefficient scores.

5. Results

This section aims to summarise the effect of a wide variety of dimensionality reduction techniques on the performance of multilingual siamese transformers by comparing the baseline approaches (i.e., Approaches 1 and 2) with the reduced approaches (i.e., Approaches 3 and 4) for each model independently. It must be noted that this work does not pretend to provide a comparative analysis between the different models presented in Section 3.2 or to identify the best model for this task. In contrast, this work is focused on applying these dimensionality reduction techniques for reducing the dimensionality of the models’ embeddings. Thus, the application of different dimensionality reduction techniques does not affect the execution or the memory requirement for running the models. It only diminishes the needed space to store embeddings and increases the speed to compute the cosine similarity between them.

Due to space reasons, average results across the 31 monolingual and cross-lingual tasks are presented instead of a breakdown by language. The average of Spearman correlation coefficients is computed by transforming each correlation coefficient to a Fisher’s z value, averaging them, and then back-transforming to a correlation coefficient.

5.1. Approach 1 vs. Approach 3: Dimensionality reduction in pre-trained embeddings

As can be seen in Figure 3, for every model, the pre-trained performance on mSTSb (i.e., Approach 1) is improved using different dimensional reduction techniques. These results evidence that dimension reduction techniques are able to somehow adjust the knowledge present in the pre-trained embeddings to the semantic similarity task. This fact becomes even more significant in the case of LaBSE, a model with zero-shot capabilities trained on a task close to semantic similarity, which also greatly benefits from the use of dimension reduction techniques, increasing by 0.4 points the Spearman correlation coefficient (see Figure 3e and Table 3). For the rest of the models, it is equally remarkable that the dimension reduction techniques improve the pre-training performance, almost doubling score in the models with the XLM-RoBERTa architecture.

Clearly, the best technique in Approach 3 is ICA. Not only because it obtains the most remarkable improvement in pre-training performance for all models as shown in Table 3, but also because it is the




Figure 3: Approach 3 average Spearman r_s correlation coefficient in multilingual tasks from the mSTSb test as a function of the number of dimension for the different dimensionality reduction techniques grouped by model.

Model (Ap. 1 Avg r_s)	Technique	Threshold Performance Retained	Dimensions (% reduction)	Ap. 3 Avg r_s	Fitting Time
bert-base- multilingual -cased (0.4342)	IPCA	100%	209 (73%)	0.4251	27s
	ICA	100%	89 (88%)	0.4779	1m 10s
	poly	95%	249 (68%)	0.4130	49s
	rbf	95%	448 (42%)	0.4138	1m 38s
	sigmoid	100%	129 (83%)	0.4425	40s
	cosine	100%	209 (73%)	0.4350	36s
	UMAP	50%	129 (83%)	0.2176	35s
	VarThres	85%	161 (79%)	0.3727	2s
distilbert-base- multilingual -cased (0.4531)	IPCA	100%	209 (73%)	0.4553	38s
	ICA	100%	49 (94%)	0.4564	43s
	poly	95%	369 (52%)	0.4310	1m 6s
	rbf	95%	608 (21%)	0.4305	2m 13s
	sigmoid	100%	129 (83%)	0.4642	38s
	cosine	100%	209 (73%)	0.4537	33s
	UMAP	40%	49 (94%)	0.3942	18s
	VarThres	95%	238 (69%)	0.438	2s
xlm-roberta -base (0.3274)	IPCA	100%	89 (88%)	0.3711	38s
	ICA	100%	49 (94%)	0.4043	56s
	poly	100%	129 (83%)	0.3439	25s
	rbf	100%	129 (83%)	0.3439	50s
	sigmoid	100%	49 (94%)	0.3425	29s
	cosine	100%	89 (88%)	0.3709	15s
	UMAP	40%	10 (99%)	0.1320	15s
	VarThres	100%	52 (93%)	0.3310	2s
xlm-roberta -large (0.2885)	IPCA	100%	116 (89%)	0.3149	1m 37s
	ICA	100%	63 (94%)	0.3642	1m 52s
	poly	100%	223 (78%)	0.2927	45s
	rbf	100%	276 (73%)	0.2934	1m 8s
	sigmoid	100%	63 (94%)	0.2927	32s
	cosine	100%	116 (89%)	0.3191	25s
	UMAP	45%	10 (99%)	0.1365	15s
	VarThres	100%	598 (42%)	0.2917	3s
LaBSE (0.7096)	IPCA	100%	129 (83%)	0.7251	37s
	ICA	100%	89 (88%)	0.7431	1m 35s
	poly	100%	169 (78%)	0.7181	34s
	rbf	100%	408 (47%)	0.7106	1m 35s
	sigmoid	100%	89 (88%)	0.7127	34s
	cosine	100%	129 (83%)	0.7232	21s
	UMAP	70%	10 (99%)	0.5026	33s
	VarThres	85%	217 (72%)	0.6148	2s

Model (Ap. 2 Avg r_s)	Technique	Threshold Performance Retained	Dimensions (% reduction)	Ap. 4 Avg r_s	Fitting Time
bert-base- multilingual -cased -fine-tuned (0.7045)	IPCA	95%	49 (94%)	0.6710	34s
	ICA	100%	169 (78%)	0.7047	3m 37s
	poly	95%	129 (83%)	0.6716	31s
	rbf	95%	169 (78%)	0.6738	54s
	sigmoid	100%	329 (57%)	0.7048	1m 34s
	cosine	95%	49 (94%)	0.6707	11s
	UMAP	70%	10 (99%)	0.5398	32s
	VarThres	100%	393 (53%)	0.7046	2s
distilbert-base- multilingual -cased -fine-tuned (0.6863)	IPCA	95%	49 (94%)	0.6533	35s
	ICA	95%	49 (94%)	0.6556	56s
	poly	95%	129 (83%)	0.6542	30s
	rbf	95%	129 (83%)	0.6520	49s
	sigmoid	95%	49 (94%)	0.6601	29s
	cosine	95%	89 (88%)	0.6631	16s
	UMAP	75%	10 (99%)	0.5189	25s
	VarThres	95%	66 (91%)	0.6620	2s
xlm-roberta -base -fine-tuned (0.7470)	IPCA	95%	49 (94%)	0.7198	36s
	ICA	95%	49 (94%)	0.7208	59s
	poly	95%	89 (88%)	0.7112	25s
	rbf	95%	129 (83%)	0.7134	49s
	sigmoid	100%	289 (62%)	0.7472	1m 45s
	cosine	95%	49 (94%)	0.7195	11s
	UMAP	75%	10 (99%)	0.5724	31s
	VarThres	100%	411 (46%)	0.7491	3s
xlm-roberta -large -fine-tuned (0.8150)	IPCA	95%	63 (94%)	0.7910	51s
	ICA	95%	63 (94%)	0.7950	1m 16s
	poly	95%	63 (94%)	0.7774	23s
	rbf	95%	63 (94%)	0.7760	42s
	sigmoid	100%	223 (78%)	0.8151	50s
	cosine	95%	63 (94%)	0.7916	13s
	UMAP	80%	10 (99%)	0.6584	38s
	VarThres	95%	95 (91%)	0.7936	3s
LaBSE -fine-tuned (0.8242)	IPCA	95%	89 (88%)	0.8014	34s
	ICA	95%	89 (88%)	0.7986	2m 3s
	poly	95%	89 (88%)	0.7898	25s
	rbf	95%	129 (83%)	0.7932	47s
	sigmoid	100%	728 (5%)	0.8243	7m 11s
	cosine	95%	89 (88%)	0.8001	21s
	UMAP	80%	23 (97%)	0.6640	35s
	VarThres	95%	227 (70%)	0.7964	2s

Table 5: Analysis of the lowest reduced number of dimensions from Approach 3 that improves the result of the baseline Approach 1 for a specific performance threshold retained. For instance, 100% threshold represents that the technique achieves at least the 100% of the baseline approach score.

Table 6: Analysis of the lowest reduced number of dimensions from Approach 4 that improves the result of the baseline Approach 2 for a specific performance threshold retained. For instance, 100% threshold represents that the technique achieves at least the 100% of the baseline approach score.

technique that most quickly and with the fewest dimensions overcomes the pre-trained models of Approach 1 (see Table 5). From the results of this Table, the ICA technique improves the pre-trained Approach 1 performances reducing an average of $91.58\% \pm 2.59\%$ of the initial dimensions retaining 100% of the baseline Approach 1 performance. Remarkably, the two-tailed paired T-test comparing Approach 1 vs Approach 3 using the values for this technique from Table 3 resulted in $p = 0.041$, indicating that the improvement in performance is significant when using ICA as a dimensional reduction techniques for pre-trained embeddings. These findings corroborate the ideas of Raunak et al. [21], who maintained that reduced word embeddings can achieve similar or better performance than original pre-trained embeddings.

The most likely explanation for these results is the difference in the objective of ICA from the other feature extraction techniques. Even though all of them transform the initial space through combinations of dimensions into a new space, ICA technique is based on the optimization of mutual information [72, 61]. It tries to find a space where the new features (latent variables) are as independent as possible from each other but as dependent as possible from the initial space. Therefore, in the case of ICA, unlike other techniques such as PCA or KPCA, a higher number of components does not necessarily mean an increase in the information retained or an improvement in the result (as can be seen in Figures 3 and clearly in Figure 3e where there is a decrease from 169 dimensions onwards). This would explain why a low

number of dimensions would outperform Approach 1.

Likewise, the fact that it is the technique that achieves the best results in Approach 3 in all models could be due to the assumptions and characteristics of both the ICA and the pre-trained embeddings. First, the pre-trained embeddings probably include non-relevant and noisy variables as these embeddings are not adjusted to the STS task. Secondly, since ICA is a technique in which the original variables are related linearly to the latent variables, but for which the latent distribution is non-Gaussian, the noise present in pre-trained embeddings agnostic of the STS task could be managed appropriately.

Interestingly, these results also emphasize that the issue of non-Gaussianity is more relevant than the nonlinearity issue. Non-Gaussianity would be more important than how the initial variables are combined as the ICA technique outperforms both linear PCA and nonlinear KPCA. This is in good agreement with other studies comparing the performance of PCA and ICA as a method for feature extraction in visual object recognition tasks [73, 74].

Additionally, the presence of noisy variables in the pre-trained embeddings would be also corroborated by the low scores obtained from the Variance Threshold feature selection technique, which completely depends on the original variables and cannot manage these noisy distributions.

Consequently, ICA shows great properties to obtain compacted embeddings versions of pre-trained models with a large decrease of dimensions that improve the result in the task of semantic similarity at a multilingual level.

For all these reasons, we can understand unsupervised dimensionality reduction, specially ICA, as a method of fitting pre-trained models for downstream tasks. As it will be seen in the next section and as might be expected, this unsupervised dimensionality reduction downstream fitting is not comparable to a supervised fitting such as the fine-tuning of Approach 2. However, downstream fitting by unsupervised dimensionality reduction techniques may present interesting advantages such as the fact that being unsupervised is task agnostic resulting in models with higher generalizability and with a lower number of dimensions. Also, these dimensionality reduction techniques do not require GPUs and apply a more interpretable methodology than a Deep Learning model fine-tuning such as transformers.

5.2. Approach 2 vs. Approach 4: Dimensionality reduction in fine-tuned embeddings

Although it was stated at the beginning of this section that model comparison is not an objective of the paper, different versions of the same architecture have been included for a wider evaluation of the effects of dimensionality reduction techniques (i.e., *xlm-roberta-base* & *xlm-roberta-large*, *bert-base-multilingual-cased* & *distilbert-base-multilingual-cased*) have been included in the study. Based on the complexity and learning potential of the models, one would expect the *xlm-roberta-large* model to perform better than the *xlm-roberta-base* model. Similarly, the *bert-base-multilingual-cased* model would be expected to be superior to the *distilbert-base-multilingual-cased* model. In Approach 1, however, the opposite is true. Only when fine-tuning occurs in the task (Approaches 2 and 4) is it observed how the performance of the models is in line with the expected complexity (see Table 4). These results provide wider support for the importance of supervised fine-tuning.

Similarly, fine-tuning also alters the impact of dimensionality reduction techniques on the results of multilingual models. Compared to Approach 3, when fine-tuning, the feature selection techniques and nonlinearity becomes more important, ICA becomes less critical, and the minimum number of dimensions that outperform the baseline approach increases.

As can be seen in Figure 4 and Table 4, the promotion of Variance Threshold feature selection as one of the best techniques for some models in Approach 4 could be attributed to the fact that fine-tuning adjust the embeddings to the task, reducing the presence of noisy variable and taking advantage of the variable selection process. This would be in line with the results obtained in Approach 3, where feature extraction techniques more adequately handled the presence of unadjusted variables. This further supports the argument that they can reduce dimensions and generate new feature representations to help improve performance on learning problems.




Figure 4: Approach 4 average Spearman r_s correlation coefficient in multilingual tasks from the mSTSb test as a function of the number of dimensions for the different dimensionality reduction techniques grouped by model.

Model	Technique	Time Min Dimension	Time In-between Dimension	Time Max Dimension
bert-base-multilingual-cased	IPCA	27s (10)	29s (448)	30s (768)
	ICA	53s (10)	8m 43s (448)	52m 48s (768)
	poly	18s (10)	1m 26s (448)	2m 17s (768)
	rbf	36s (10)	1m 38s (448)	3m 29s (768)
	sigmoid	24s (10)	2m 15s (448)	5m 53s (768)
	cosine	8s (10)	55s (408)	1m 24s (768)
	UMAP	20s (10)	1m 30s (448)	3m 29s (768)
distilbert-base-multilingual-cased	VarThres	2s (65)	2s (516)	2s (767)
	IPCA	36s (10)	40s (448)	44s (768)
	ICA	29s (10)	8m 30s (448)	29m 31s (768)
	poly	16s (10)	1m 24s (448)	2m 35s (768)
	rbf	35s (10)	1m 44s (448)	2m 57s (768)
	sigmoid	23s (10)	3m 58s (448)	15m 40s (768)
	cosine	8s (10)	51s (448)	1m 51s (768)
xlm-roberta-base	UMAP	13s (10)	1m 38s (448)	4m 19s (768)
	VarThres	2s (66)	2s (507)	2s (767)
	IPCA	35s (10)	35s (448)	40s (768)
	ICA	56s (10)	9m 54s (448)	21m 9s (768)
	poly	16s (10)	1m 2s (448)	2m 12s (768)
	rbf	33s (10)	1m 38s (448)	2m 2s (768)
	sigmoid	23s (10)	1m 43s (448)	3m 26s (768)
xlm-roberta-large	cosine	7s (10)	56s (448)	1m 36s (768)
	UMAP	15s (10)	2m (448s)	4m 32s (768)
	VarThres	2s (1)	2s (448)	2s (768)
	IPCA	54s (10)	54s (490)	1m 46s (1024)
	ICA	57s (10)	23m 13s (490)	32m 15s (1024)
	poly	16s (10)	1m 24s (490)	3m 35s (1024)
	rbf	35s (10)	1m 30s (490)	4m 49s (1024)
LaBSE	sigmoid	23s (10)	2m 58s (490)	8m 28s (1024)
	cosine	9s (10)	1m 38s (490)	2m 41s (1024)
	UMAP	15s (10)	1m 57s (490)	7m 24s (1024)
	VarThres	3s (10)	3s (598)	3s (1024)
	IPCA	35s (10)	36s (448)	39s (768)
	ICA	36s (10)	12m 55s (448)	52m 39s (768)
	poly	17s (10)	1m 24s (448)	2m 35s (768)
LaBSE	rbf	35s (10)	1m 41s (448)	2m 50s (768)
	sigmoid	24s (10)	4m 25s (448)	19m 37s (768)
	cosine	8s (10)	55s (448)	1m 44s (768)
	UMAP	33s (10)	2m 11s (448)	4m 18s (768)
	VarThres	2s (94)	2s (615)	2s (767)

Table 7: Time required to fit the different dimension reduction techniques in the mSTSb task. Owing to the fact that a wide range of dimensions is explored, the fitting time for the minimum, in-between and maximum number of dimensions are reported for each model and technique. The measurements are performed on an Intel(R) Xeon(R) Bronze 3206R CPU at 1.90GHz.

Model	Time	Hyperparameters
bert-base-multilingual-cased	39m 18s	batch.size: 32
		lr: 2e-5
distilbert-base-multilingual-cased	15m 31s	epochs: 2
		scheduler: warmuplinear
xlm-roberta-base	27m 49s	warmup_ratio: 0.2
		wieght_decay: 0.2
xlm-roberta-large	1h 2m 32s	batch.size: 64
		lr: 2e-5
LaBSE	59m 39s	epochs: 2
		scheduler: warmuplinear
LaBSE	59m 39s	warmup_ratio: 0.3
		wieght_decay: 0.7
LaBSE	59m 39s	batch.size: 64
		lr: 5e-5
LaBSE	59m 39s	epochs: 2
		scheduler: warmuplinear_hard_restarts
LaBSE	59m 39s	warmup_ratio: 0.1
		wieght_decay: 0.5
LaBSE	59m 39s	batch.size: 64
		lr: 1e-5
LaBSE	59m 39s	epochs: 2
		scheduler: warmupcosine
LaBSE	59m 39s	warmup_ratio: 0.2
		wieght_decay: 0
LaBSE	59m 39s	batch.size: 32
		lr: 3e-6
LaBSE	59m 39s	epochs: 2
		scheduler: warmupcosine
LaBSE	59m 39s	warmup_ratio: 0.1
		wieght_decay: 0.5

Table 8: Time required to fine-tune the different models in the mSTSb task. The fine-tuning process is carried out in Quadro-RTX 8000 GPU. The best hyperparameters configuration considered for each model is also reported.

Furthermore, Table 6 shows the lack of ICA dominance and the emergence of the KPCA-sigmoid technique as the method that with fewer dimensions improves the baseline Approach 2 ($59.32\% \pm 29.92\%$ reduced dimensions retaining $99.00\% \pm 2.00\%$ baseline’s performance). It reveals that managing the non-Gaussianity issue is less relevant than the nonlinearity issue after fine-tuning. The fine-tuning process also impacts the reduction capabilities of the dimensionality reduction techniques since considering the technique that retains the maximum performance with the lowest dimensions for each model in Table 6 show that the initial dimensions from the baseline Approach 2 are reduced an average of $54.65\% \pm 32.20\%$. Although this average reduction is lower than the achieved earlier in the comparison of Approach 3 with the baseline Approach 1, it still remarkably that even after fine-tuning, the multilingual performance can be exceeded with half of the dimensions. Finally, the two-tailed paired T-test comparing Approach 2 vs Approach 4 using the values from Table 4 resulted in $p = 0.255$, revealing that there is no significant difference in performance using these dimensional reduction techniques.

5.3. Execution time comparison

Although it is out of the main scope of the paper, this subsection compares the different dimension reduction techniques regarding the execution time required to fit each technique. As mentioned in Section 4.1, it is fundamental to note that KPCA and UMAP were fitted using a subset of $10k$ pairs of

sentences due to their computational cost, which is more expensive than their linear counterparts. Therefore, for fair comparisons, IPCA, ICA, and variance threshold techniques are compared on the one hand and KPCA and UMAP on the other.

Finally, we analyse the reduction of computational time by comparing the fitting time of these techniques in Approach 3 and 4 with respect to the fine-tuning process performed in Approach 2, following the experimental design described in Section 4.

5.3.1. Dimensionality Reduction Techniques comparison: IPCA, ICA and Variance Threshold

Considering the wide range of dimensions explored for each model, to compare the different techniques in terms of their execution time, we analyze the fitting time for the minimum, in-between and maximum number of dimensions technique in the mSTSb train split for Approaches 3 and 4. The results of these tests and the dimensions associated with these times are presented in Table 7. It should be noted that the number of dimensions shown in this Table for the variance threshold technique differs from the other techniques since the number of reduced dimensions depends on the number of dimensions that exceed the variance threshold explored.

From Table 7, it can be noted that the technique that requires the longest fitting time of IPCA, ICA and variance threshold is clearly ICA, while variance threshold is the one that requires the shortest time.

Regarding computational complexity, IPCA has constant memory complexity $O(bd^2)$ [64], while the ICA algorithm is $O(2md(d+1)n)$ [75], where n is the number of samples of dimensions d , b is the batch size for the IPCA algorithm, and m is the number of iterations of the ICA algorithm. We can observe that when IPCA and ICA are compared, for the same number of dimensions and number of samples, the computational burden of the ICA technique is the number of iterations (m), since they critically increase execution time. As stated in [75], this effect is more noticeable as the dimension increases, making the use of ICA computationally unfeasible when the number of dimensions is moderate or high. Our results support these findings, as the ICA’s fitting time explodes when the number of dimensions is high (see Table 7). On the other hand, the IPCA execution time remains practically constant because of the fact that the batches avoid attempting to load all the data into memory at the same time.

Finally, as expected, the feature selection technique is the least time-consuming since it only requires computing the variance for each dimension and filtering it according to a selected variance threshold.

Consequently, there is a trade-off between the number of dimensions to be reduced, the fitting time, and the percentage of performance to be retained. Our recommendation would be that for the case of pre-trained embedding reduction, the ICA technique should be the priority if a low number of dimensions is explored, since it has been shown to be the best at retaining and improving the initial performance at a feasible cost. However, if the target number of dimensions to be reduced is medium or high, we would recommend the IPCA technique, since it can retain and improve performance with a constant computational cost. Finally, suppose that the execution time is clearly prioritized. In that case, we recommend the variance threshold feature selection technique, even if the number of dimensions is reduced and the performance retained is not the best, as this technique cannot handle the noise in pre-trained features not adjusted for a downstream task.

5.3.2. Dimensionality Reduction Techniques comparison: KPCA and UMAP

Regarding the KPCA technique, using a kernel function before performing linear PCA increases the execution time. These kernel functions map the data into a nonlinear feature space, so the execution time varies depending on the kernel type. From the results shown in Table 7, the sigmoid kernel requires the most execution time and the cosine kernel the least. Despite these previous results, the selection of the kernel type depends on each case, and different kernels should be explored, as previous studies suggest [59, 62].

Finally, it can be seen from the results shown in Table 7 that the execution time of UMAP is similar to KPCA. Nevertheless, several considerations must be taken into account in the case of the UMAP

technique. As discussed in Section 3.1, the number of neighbours parameter is proportional to the amount of global information retained. As shown in Table 2, different values of the number of neighbours are explored for each model. The number of neighbours that provided the best results were 10, 5, 5, 5 and 125 for each model in the same order as the models shown in Table 7. The value of this parameter determines the execution time of the technique, increasing the time the greater the number of neighbours to consider for embedding a sentence. This is the reason why the degree of similarity between KPCA and UMAP run times varies from one model to another. Similarly, although the Table only includes the execution times for the best UMAP parameter values found for each model, the choice of these parameters requires a prior exploration process that can be very long, depending on how intensive the search for the best parameter values is. Hence, this technique has an additional cost to consider.

5.3.3. Fitting Dimensionality Reduction Technique and Fine-tuning comparison

Another point to discuss is comparing the execution times required by dimension reduction techniques with the fine-tuning process. It should be noted that the fine-tuning times shown in Table 8 do not include the time required for hyperparameter optimization, which, depending on how extensive the search is, can extend the time required to obtain a fine-tuned model. Therefore, the time shown refers to the execution time needed to fine-tune each model once the optimal hyperparameters are found.

In practical application, the time required for fitting a dimensionality reduction technique on pre-trained embeddings instead of fine-tuning the model downstream is what matters. For this purpose, we use the fine-tuning process execution times shown in Table 8 compared to Table 5, which shows the fitting time of the dimension reduction techniques of Approach 3 that improves the result of the baseline Approach 1.

For space reasons, even though Table 5 shows a breakdown of the performance for each technique, the technique used to compare the fine-tuning of each model is the technique that, with the lowest number of dimensions, retains the highest possible performance (highlighted in bold in Table 5). As stated previously, the best technique for all models is ICA. This technique requires a fitting time of an average of $96.68\% \pm 0.68\%$ faster than the fine-tuning process on all models.

Taken together, we can see the advantage in execution time of using reduction techniques in Approach 3 with respect to the fine-tuning process carried out in Approach 2. Remarkably, it is exciting to point out the potential of ICA as an alternative to the fine-tuning process. Although it has previously been proven that fine-tuning achieves better performance than fitting dimension reduction techniques, this technique does not require GPU. This technique reduces the execution time by more than 96%. This fact, together with the ability to improve the performance of pre-trained embeddings and the generalisation capability of this unsupervised technique, reveals the potential of this technique as an alternative to fine-tuning.

5.4. UMAP: a case of study

In this section, we pay special attention to the recently proposed UMAP technique [63]. The case of UMAP shows that for the STS task, it is not a suitable technique to reduce the dimensionality of the embeddings since it is the one that retains the lowest percentage of the baseline performances in both pre-trained and fine-tuned embeddings (see Tables 5, 6). Considering this fact, it is interesting to note that the potential of this technique resides in the fact that it quickly saturates, i.e., the maximum retained performance is reached with a small number of dimensions in Approach 3 with an average of $94.65\% \pm 6.07\%$ of reduced initial dimensions retaining $49.00\% \pm 11.14\%$ performance with respect to the reference Approach 1, and more notably in Approach 4 with an average of $98.42\% \pm 0.72\%$ of reduced initial dimensions retaining $76.00\% \pm 3.74\%$ performance with respect to the reference Approach 2.

This saturation behaviour can be attributed to the demonstrated functionality of UMAP for generating visualizations by reducing high-dimensional data to 2- or 3-dimensions. As reported in other works [76, 77], the most significant potential of this technique lies precisely in visualization, where this saturation capability is exploited. UMAP is a nonlinear graph-based method for dimensionality reduction that is not

meant to extract features as methods like PCA do [63]. Instead, its primary goal is to represent the input data as a high-dimensional graph and then reconstruct it in a lower-dimensional space while retaining structure. These characteristics explain the lack of good results for the UMAP technique in the mSTSb task in Approach 3, where feature extraction has proven to be helpful for managing the noisy pre-trained features. Furthermore, this would explain why we observe UMAP saturation with fewer dimensions in Approach 4, as the embeddings are already fine-tuned and UMAP can exploit its potential to preserve information.

Previously, it has been proven that its fitting requires an execution time similar to other techniques such as KPCA. It is also important to note that its computational cost limits the total number of data instances to be used for fitting, since the first step is expensive, as it requires computing a graphical representation of the dataset and then learning an embedding for that graph.

Despite the previous conclusion, UMAP is time-consuming. This technique is highly dependent on the parameters used, and there are many parameters that can be explored. Additionally, the slow process of embedding new data in a fitted UMAP must also be taken into account. This is because UMAP does not build a function that directly maps high-dimensional points down to a reduced space. Instead, it first computes a graph representing the whole data and then learns an embedding for that graph. Thus, embedding a new point requires calculating its nearest neighbours from the fitting data and embedded in the learnt graph. Again, this process execution time increases with the number of neighbours value selected when. Therefore, UMAP can transform new data, albeit more slowly than other techniques that allow this.

5.5. Previous works comparison

Finally, we discuss our results concerning previous related work in dimensionality reduction of pre-computed embeddings. Before this discussion proceeds, it is essential to remark that our work differs from previous work in the literature as we explore a more comprehensive range of unsupervised dimension reduction techniques, evaluating them on pre-trained and fine-tuned precomputed embeddings from state-of-the-art multilingual contextual-based transformer models in the Semantic Textual Similarity (STS) task.

Our experimental results are in agreement with Raunak et al. [21, 22] that corroborates the hypothesis that reduction in pre-trained embeddings can maintain or improve the performance of the original embeddings. Similarly, our results are consistent with those obtained by Maria Mihaela et al. [23] and Shimomoto et al. [44]. The authors employ PCA as the only unsupervised technique. We have confirmed that this feature extraction technique has great potential to reduce the dimensions of pre-trained embeddings. Our results further reveal that the incremental PCA method (IPCA) is also suitable for embedding reduction. Besides, we found that for embeddings already fine-tuned to the downstream task, the use of the KPCA, the nonlinear version of PCA, is much more helpful for preserving performance and reducing dimensions.

Regarding execution time, our results are similar to those obtained by Saeed et al. [17]. These authors combined the unsupervised PCA technique and the supervised LDA technique with classical NLP techniques based on N-grams and classical ML models (Decision Trees, Logistic Regression, or Naive Bayes) for sentiment classification of monolingual Arabic texts. The authors reported an improvement in results over previous work by reducing the number of dimensions by up to 93% with a 97% shorter run time.

Even though the NLP task addressed in the present study is STS from a multilingual perspective (including Arabic) and not sentiment classification, our findings are in line with these previous results. The ICA technique reduces an average of $91.58\% \pm 2.59\%$ of the initial dimensions, retaining 100% of the baseline Approach 1 performance, requiring a fitting time of $96.68\% \pm 0.68\%$ faster than the fine-tuning process.

Given these points, we can safely conclude that our results are in line with previous work extending their findings to state-of-the-art contextual-based models from a multilingual approach. This paper provides new insights into dimensionality reduction techniques for a space- and time-efficient data representation.

6. Conclusion

In this investigation, the goal was to assess the impact of a variety of dimensionality reduction techniques on the performance of pre-computed multilingual siamese fashion transformers embeddings on semantic textual similarity tasks from mSTSb, to expand our knowledge of semantic-aware transformer-based models. To this end, two different baseline approaches are reduced (i.e., Approach 1 and Approach 2), one using the pre-trained version of the models and the second one further fine-tuning them on the downstream STS task. Particular attention is paid to analysing which techniques best and with the fewest dimensions improve the performance of the baseline approaches.

From the research carried out, it is possible to conclude that dimensionality reduction techniques can help to reduce the number of dimensions of the embeddings while improving the results if using pre-trained embeddings from Approach 1 and preserving the performance when using fine-tuned embeddings from Approach 2. Nevertheless, the dimensionality reduction is more considerable in the pre-trained version with an average of $91.58\% \pm 2.59\%$ compared to the average $54.65\% \pm 32.20\%$ of the fine-tuned version. Special attention is given to ICA in the pre-trained scenario, which proved to manage adequately the noisy variables present in not adjusted embeddings. This technique also proved to be a reasonable alternative to fit the models in the downstream task in an unsupervised way, leading to a generalised adjusted version of the models with downstream multitasking capabilities. Nevertheless, it has also been proved that this unsupervised fitting is not comparable to a supervised fine-tuning. On the other hand, the fine-tuned scenario revealed the relevance of feature selection techniques and the significance of nonlinear KPCA techniques for dimensionality reduction.

Additionally, although our execution time experiments reveal that ICA is the most time consuming technique among the dimension reduction techniques, reducing pre-trained embeddings (Approach 3) by gaining generalisation power with this unsupervised technique is still faster than performing downstream fine-tuning (Approach 2). Moreover, ICA does not require GPUs for fitting, reinforcing its potential as an alternative to the fine-tuning process. Overall, a good trade-off between performance against available computational resources, execution time, and the number of dimensions to be reduced must be considered when choosing the approach to follow.

The results of our experiments are consistent with previous results from the literature, corroborating the hypothesis that reduction in embeddings can maintain or improve the performance of the original embeddings by extending their evaluation to state-of-the-art contextual-based models from a multilingual approach. In this way, we can establish that dimensionality reduction techniques could be leveraged for contextualised embeddings as well.

To our knowledge, this is the first study to investigate the effect of dimensionality reduction techniques and transformers models in a multilingual semantic-awareness scenario. This study provides an analysis of alternatives to the storage limitation we are about to face if the current trends of using large datasets and the growth rate of storage utilisation persist. Based on the promising findings presented in this article, continued research into the impact of dimensionality reduction techniques in other highly demanded NLP tasks appears to be totally justified. Furthermore, in future research, we intend to focus on testing the reduced models presented in this work in real-world applications. Besides, we hope to carry out further experimental investigation including other dimensionality reduction approaches such as creating distilled version of pre-trained models.

As stated previously, the findings presented in this study of multilingual semantic similarity are of direct practical applicability. The combination of dimensionality reduction techniques with transformer

models could also help reduce the embeddings size and make possible ensemble approaches. Finally, further studies about the multitasking generalisation capabilities of ICA for pre-trained models are still required.

Data Availability Statement

The Multilingual Semantic Textual Similarity Benchmark (mSTSb) used for the purpose of this article can be found in: <https://github.com/Huertas97/Multilingual-STSB>

Acknowledgements

This research has been supported by the Spanish Ministry of Science and Education under FightDIS (PID2020-117263GB-I00), by MCIN/AEI/10.13039/501100011033/ and European Union NextGenerationEU/PRTR for XAI-Disinfodemics (PLEC2021-007681) grant, by Comunidad Autónoma de Madrid under S2018/ TCS-4566 (CYNAMON) grant, by BBVA Foundation grants for scientific research teams SARS-CoV-2 and COVID-19 under the grant: "CIVIC: *Intelligent characterisation of the veracity of the information related to COVID-19*", and by IBERIFIER (Iberian Digital Media Research and Fact-Checking Hub), funded by the European Commission under the call CEF-TC-2020-2, grant number 2020-EU-IA-0252. Finally, David Camacho has been supported by the Comunidad Autónoma de Madrid under: "Convenio Plurianual with the Universidad Politécnica de Madrid in the actuation line of *Programa de Excelencia para el Profesorado Universitario*".

Compliance with Ethical Standards

- **Disclosure of potential conflicts of interest:** The author declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.
- **Research involving human participants and/or animals:** This article does not contain any studies with human participants or animals performed by any of the authors.
- **Informed consent:** Informed consent was obtained from all individual participants included in the study.

References

- [1] D. W. Otter, J. R. Medina, J. K. Kalita, A Survey of the Usages of Deep Learning for Natural Language Processing, *IEEE Transactions on Neural Networks and Learning Systems* 32 (2) (2021) 604–624. doi:10.1109/TNNLS.2020.2979670.
- [2] Y. Tay, M. Dehghani, D. Bahri, D. Metzler, Efficient Transformers: A Survey, *ACM Computing Surveys* (2022). doi:10.1145/3530811.
- [3] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is All You Need, in: *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, Curran Associates Inc., Red Hook, NY, USA, 2017, p. 6000–6010. doi:10.5555/3295222.3295349.

- [4] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1, Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 4171–4186. [doi:10.18653/v1/N19-1423](https://doi.org/10.18653/v1/N19-1423).
- [5] N. Reimers, I. Gurevych, Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Association for Computational Linguistics, Hong Kong, China, 2019, pp. 3982–3992. [doi:10.18653/v1/D19-1410](https://doi.org/10.18653/v1/D19-1410).
- [6] J. Huertas-Tato, A. Martin, D. Camacho, BERTuit: Understanding Spanish language in Twitter through a native transformer (2022). [doi:10.48550/ARXIV.2204.03465](https://doi.org/10.48550/ARXIV.2204.03465).
- [7] K. R. Chowdhary, Natural Language Processing, Springer India, New Delhi, 2020, pp. 603–649. [doi:10.1007/978-81-322-3972-7_19](https://doi.org/10.1007/978-81-322-3972-7_19).
- [8] D. Cer, M. Diab, E. Agirre, I. Lopez-Gazpio, L. Specia, SemEval-2017 Task 1: Semantic Textual Similarity Multilingual and Crosslingual Focused Evaluation, in: Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017), Association for Computational Linguistics, Vancouver, Canada, 2017, pp. 1–14. [doi:10.18653/v1/S17-2001](https://doi.org/10.18653/v1/S17-2001).
- [9] S. Humeau, K. Shuster, M.-A. Lachaux, J. Weston, Poly-encoders: Architectures and Pre-training Strategies for Fast and Accurate Multi-sentence Scoring, in: International Conference on Learning Representations (ICLR), Online, 2020. [doi:10.48550/ARXIV.1905.01969](https://doi.org/10.48550/ARXIV.1905.01969).
- [10] V. Zhelezniak, A. Savkov, A. Shen, N. Hammerla, Correlation Coefficients and Semantic Textual Similarity, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 951–962. [doi:10.18653/v1/N19-1100](https://doi.org/10.18653/v1/N19-1100).
- [11] G. Sidorov, A. Gelbukh, H. Gómez-Adorno, D. Pinto, Soft similarity and soft cosine measure: similarity of features in vector space model, *Computación y Sistemas* 18 (3) (2014) 491 – 504. [doi:10.13053/cys-18-3-2043](https://doi.org/10.13053/cys-18-3-2043).
- [12] E. Cambria, H. Wang, B. White, Guest editorial: Big social data analysis, *Knowledge-Based Systems* 69 (2014) 1–2. [doi:https://doi.org/10.1016/j.knosys.2014.07.002](https://doi.org/10.1016/j.knosys.2014.07.002).
- [13] O. Araque, I. Corcuera-Platas, J. F. Sánchez-Rada, C. A. Iglesias, Enhancing deep learning sentiment analysis with ensemble techniques in social applications, *Expert Systems with Applications* 77 (2017) 236–246. [doi:https://doi.org/10.1016/j.eswa.2017.02.002](https://doi.org/10.1016/j.eswa.2017.02.002).
- [14] Y. Zhou, Y. Yang, H. Liu, X. Liu, N. Savage, Deep Learning Based Fusion Approach for Hate Speech Detection, *IEEE Access* 8 (2020) 128923–128929. [doi:10.1109/ACCESS.2020.3009244](https://doi.org/10.1109/ACCESS.2020.3009244).
- [15] A. Khan, A. Sohail, U. Zahoora, A. S. Qureshi, A survey of the recent architectures of deep convolutional neural networks, *Artificial Intelligence Review* 53 (8) (2020) 5455–5516. [doi:10.1007/s10462-020-09825-6](https://doi.org/10.1007/s10462-020-09825-6).

- [16] E. C. Chau, N. A. Smith, Specializing Multilingual Language Models: An Empirical Study, in: Proceedings of the 1st Workshop on Multilingual Representation Learning, Association for Computational Linguistics, Punta Cana, Dominican Republic, 2021, pp. 51–61. [doi:10.18653/v1/2021.mrl-1.5](https://doi.org/10.18653/v1/2021.mrl-1.5).
- [17] R. M. K. Saeed, S. Rady, T. F. Gharib, Optimizing Sentiment Classification for Arabic Opinion Texts, *Cognitive Computation* 13 (1) (2021) 164–178. [doi:10.1007/s12559-020-09771-z](https://doi.org/10.1007/s12559-020-09771-z).
- [18] A. Herbelot, X. Zhu, A. Palmer, N. Schneider, J. May, E. Shutova (Eds.), Proceedings of the Fourteenth Workshop on Semantic Evaluation, International Committee for Computational Linguistics, Barcelona (online), 2020.
- [19] N. Ferro, What Happened in CLEF... For a While?, in: Experimental IR Meets Multilinguality, Multimodality, and Interaction, Springer International Publishing, Cham, 2019, pp. 3–45. [doi:10.1007/978-3-030-28577-7_1](https://doi.org/10.1007/978-3-030-28577-7_1).
- [20] **Introducing The World’s Largest Open Multilingual Language Model: BLOOM** (Jul. 2022). URL <https://bigscience.huggingface.co/blog/bloom>
- [21] V. Raunak, V. Gupta, F. Metze, Effective Dimensionality Reduction for Word Embeddings, in: Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019), Association for Computational Linguistics, Florence, Italy, 2019, pp. 235–243. [doi:10.18653/v1/W19-4328](https://doi.org/10.18653/v1/W19-4328).
- [22] V. Raunak, V. Kumar, V. Gupta, F. Metze, On Dimensional Linguistic Properties of the Word Embedding Space, in: Proceedings of the 5th Workshop on Representation Learning for NLP, Association for Computational Linguistics, Online, 2020, pp. 156–165. [doi:10.18653/v1/2020.repl4nlp-1.19](https://doi.org/10.18653/v1/2020.repl4nlp-1.19).
- [23] T. Maria Mihaela, A. Anamaria, G. Simona Elena, A. Crisan, Post-Processing and Dimensionality Reduction for Extreme Learning Machine in Text Classification, *Economic Computation & Economic Cybernetics Studies & Research* 55 (4) (2021) 37–50. [doi:10.24818/18423264/55.4.21.03](https://doi.org/10.24818/18423264/55.4.21.03).
- [24] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, R. Harshman, Indexing by latent semantic analysis, *Journal of the American Society for Information Science* 41 (6) (1990) 391–407. [doi:10.1002/\(SICI\)1097-4571\(199009\)41:6<391::AID-ASII1>3.0.CO;2-9](https://doi.org/10.1002/(SICI)1097-4571(199009)41:6<391::AID-ASII1>3.0.CO;2-9).
- [25] W. Sun, Q. Du, Hyperspectral Band Selection A review, *IEEE Geoscience and Remote Sensing Magazine* 7 (2) (2019) 118–139. [doi:10.1109/MGRS.2019.2911100](https://doi.org/10.1109/MGRS.2019.2911100).
- [26] S. Solorio-Fernández, J. A. Carrasco-Ochoa, J. F. Martínez-Trinidad, A Review of Unsupervised Feature Selection Methods, *Artificial Intelligence Review* 53 (2) (2020) 907–948. [doi:10.1007/s10462-019-09682-y](https://doi.org/10.1007/s10462-019-09682-y).
- [27] A. E. Maxwell, T. A. Warner, F. Fang, Implementation of machine-learning classification in remote sensing: an applied review, *International Journal of Remote Sensing* 39 (9) (2018) 2784–2817. [doi:10.1080/01431161.2018.1433343](https://doi.org/10.1080/01431161.2018.1433343).
- [28] A. A. Patel, Hands-on unsupervised learning using Python: how to build applied machine learning solutions from unlabeled data, O’Reilly, Sebastopol, California, 2019.

- [29] Z. M. Hira, D. F. Gillies, A Review of Feature Selection and Feature Extraction Methods Applied on Microarray Data, *Advances in bioinformatics* 2015 (2015) 198363–13. [doi:10.1155/2015/198363](https://doi.org/10.1155/2015/198363).
- [30] D. Xu, I. E.-H. Yen, J. Zhao, Z. Xiao, Rethinking Network Pruning – under the Pre-train and Fine-tune Paradigm, in: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Association for Computational Linguistics, Online, 2021, pp. 2376–2382. [doi:10.18653/v1/2021.naacl-main.188](https://doi.org/10.18653/v1/2021.naacl-main.188).
- [31] D. Bahdanau, T. Bosc, S. Jastrzebski, E. Grefenstette, P. Vincent, Y. Bengio, Learning to Compute Word Embeddings On the Fly (2017). [doi:10.48550/ARXIV.1706.00286](https://doi.org/10.48550/ARXIV.1706.00286).
- [32] J. Johnson, M. Douze, H. Jégou, Billion-Scale Similarity Search with GPUs, *IEEE Transactions on Big Data* 7 (3) (2021) 535–547. [doi:10.1109/TBDATA.2019.2921572](https://doi.org/10.1109/TBDATA.2019.2921572).
- [33] B. Mitra, N. Craswell, An Introduction to Neural Information Retrieval, *Foundations and Trends® in Information Retrieval* 13 (1) (2018) 1–126. [doi:10.1561/15000000061](https://doi.org/10.1561/15000000061).
- [34] F. Camastra, A. Vinciarelli, Feature Extraction Methods and Manifold Learning Methods, in: *Machine Learning for Audio, Image and Video Analysis*, Springer London, London, 2008, pp. 305–341. [doi:10.1007/978-1-84800-007-0_11](https://doi.org/10.1007/978-1-84800-007-0_11).
- [35] E. Strubell, A. Ganesh, A. McCallum, Energy and Policy Considerations for Deep Learning in NLP, in: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Florence, Italy, 2019, pp. 3645–3650. [doi:10.18653/v1/P19-1355](https://doi.org/10.18653/v1/P19-1355).
- [36] Y. Bengio, A. Courville, P. Vincent, Representation Learning: A Review and New Perspectives, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35 (8) (2013) 1798–1828. [doi:10.1109/TPAMI.2013.50](https://doi.org/10.1109/TPAMI.2013.50).
- [37] S. W. Choi, B. H. S. Kim, Applying PCA to Deep Learning Forecasting Models for Predicting PM2.5, *Sustainability* 13 (7) (2021). [doi:10.3390/su13073726](https://doi.org/10.3390/su13073726).
- [38] D. Menaga, S. Revathi, Probabilistic Principal Component Analysis (PPCA) Based Dimensionality Reduction and Deep Learning for Cancer Classification, in: S. S. Dash, S. Das, B. K. Panigrahi (Eds.), *Intelligent Computing and Applications*, Springer Singapore, Singapore, 2021, pp. 353–368. [doi:10.1007/978-981-15-5566-4_31](https://doi.org/10.1007/978-981-15-5566-4_31).
- [39] N. Kushwaha, M. Pant, Textual Data Dimensionality Reduction - a Deep Learning Approach, *Multimedia Tools Applications* 79 (15–16) (2020) 11039–11050. [doi:10.1007/s11042-018-6900-x](https://doi.org/10.1007/s11042-018-6900-x).
- [40] J. Pennington, R. Socher, C. Manning, GloVe: Global Vectors for Word Representation, in: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, Doha, Qatar, 2014, pp. 1532–1543. [doi:10.3115/v1/D14-1162](https://doi.org/10.3115/v1/D14-1162).
- [41] P. Bojanowski, E. Grave, A. Joulin, T. Mikolov, Enriching Word Vectors with Subword Information, *Transactions of the Association for Computational Linguistics* 5 (2017) 135–146. [doi:10.1162/tacl_a_00051](https://doi.org/10.1162/tacl_a_00051).

- [42] K. Pearson, On lines and planes of closest fit to systems of points in space, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2 (11) (1901) 559–572. doi:10.1080/14786440109462720.
- [43] I. T. Jolliffe, J. Cadima, Principal component analysis: a review and recent developments, *Philosophical Transactions of the Royal Society a Mathematical Physical and Engineering Sciences* 374 (2065) (2016). doi:10.1098/rsta.2015.0202.
- [44] E. K. Shimomoto, F. Portet, K. Fukui, Text classification based on the word subspace representation, *Pattern analysis and applications : PAA* 24 (3) (2021) 1075–1093. doi:10.1007/s10044-021-00960-6.
- [45] H. Song, D. Zou, L. Hu, J. Yuan, Embedding Compression with Right Triangle Similarity Transformations, in: *Artificial Neural Networks and Machine Learning – ICANN 2020, Lecture Notes in Computer Science*, Springer International Publishing, Cham, 2020, pp. 773–785. doi:10.1007/978-3-030-61616-8_62.
- [46] R. Choudhary, S. Doboli, A. A. Minai, A Comparative Study of Methods for Visualizable Semantic Embedding of Small Text Corpora, in: *2021 International Joint Conference on Neural Networks (IJCNN)*, 2021, pp. 1–8. doi:10.1109/IJCNN52387.2021.9534250.
- [47] G. Hinton, S. Roweis, Stochastic Neighbor Embedding, in: *Proceedings of the 15th International Conference on Neural Information Processing Systems, NIPS’02*, MIT Press, Cambridge, MA, USA, 2002, p. 857–864.
- [48] Á. Huertas-García, J. Huertas-Tato, A. Martín, D. Camacho, Countering Misinformation Through Semantic-Aware Multilingual Models, in: *Intelligent Data Engineering and Automated Learning – IDEAL 2021*, Springer International Publishing, Cham, 2021, pp. 312–323. doi:10.1007/978-3-030-91608-4_31.
- [49] R. Nogueira, Z. Jiang, R. Pradeep, J. Lin, Document Ranking with a Pretrained Sequence-to-Sequence Model, in: *Findings of the Association for Computational Linguistics: EMNLP 2020*, Association for Computational Linguistics, Online, 2020, pp. 708–718. doi:10.18653/v1/2020.findings-emnlp.63.
- [50] S. Robertson, H. Zaragoza, M. Taylor, Simple BM25 Extension to Multiple Weighted Fields, in: *Proceedings of the Thirteenth ACM International Conference on Information and Knowledge Management, CIKM ’04*, Association for Computing Machinery, New York, NY, USA, 2004, p. 42–49. doi:10.1145/1031171.1031181.
- [51] C. Wardle, H. Derakhshan, *Information disorder: Toward an interdisciplinary framework for research and policy making*, Tech. rep., Council of Europe (Oct. 2017).
URL <https://rm.coe.int/information-disorder-toward-an-interdisciplinary-frameworkfor-researc/168076277c>
- [52] E. Carmi, S. J. Yates, E. Lockley, A. Pawluczuk, Data citizenship: Rethinking data literacy in the age of disinformation, misinformation, and malinformation, *Internet policy review* 9 (2) (2020). doi:10.14763/2020.2.1481.
- [53] J. Gaglani, Y. Gandhi, S. Gogate, A. Halbe, Unsupervised WhatsApp Fake News Detection using Semantic Search, in: *2020 4th International Conference on Intelligent Computing and Control Systems (ICICCS)*, 2020, pp. 285–289. doi:10.1109/ICICCS48265.2020.9120902.

- [54] Á. Huertas-García, J. Huertas-Tato, A. Martín, D. Camacho, CIVIC-UPM at CheckThat! 2021: Integration of Transformers in Misinformation Detection and Topic Classification, in: Proceedings of the Working Notes of CLEF 2021 - Conference and Labs of the Evaluation Forum, Vol. 2936 of CEUR Workshop Proceedings, CEUR-WS.org, Bucharest, Romania, 2021, pp. 520–530.
- [55] A. Martín, J. Huertas-Tato, Á. Huertas-García, G. Villar-Rodríguez, D. Camacho, FacTeR-Check: Semi-automated fact-checking through semantic similarity and natural language inference, *Knowledge-Based Systems* 251 (2022) 109265. [doi:10.1016/j.knosys.2022.109265](https://doi.org/10.1016/j.knosys.2022.109265).
- [56] M. Grootendorst, BERTopic: Neural topic modeling with a class-based TF-IDF procedure (2022). [doi:10.48550/ARXIV.2203.05794](https://doi.org/10.48550/ARXIV.2203.05794).
- [57] M. Grootendorst, KeyBERT: Minimal keyword extraction with BERT. (2020). [doi:10.5281/zenodo.4461265](https://doi.org/10.5281/zenodo.4461265).
- [58] N. Reimers, I. Gurevych, Making Monolingual Sentence Embeddings Multilingual using Knowledge Distillation, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), Association for Computational Linguistics, Online, 2020, pp. 4512–4525. [doi:10.18653/v1/2020.emnlp-main.365](https://doi.org/10.18653/v1/2020.emnlp-main.365).
- [59] K.-R. Muller, S. Mika, G. Ratsch, K. Tsuda, B. Scholkopf, An introduction to kernel-based learning algorithms, *IEEE Transactions on Neural Networks* 12 (2) (2001) 181–201. [doi:10.1109/72.914517](https://doi.org/10.1109/72.914517).
- [60] D. A. Ross, J. Lim, R.-S. Lin, M.-H. Yang, Incremental Learning for Robust Visual Tracking, *International Journal of Computer Vision* 77 (1-3) (2007) 125–141. [doi:10.1007/s11263-007-0075-7](https://doi.org/10.1007/s11263-007-0075-7).
- [61] A. Hyvärinen, Independent component analysis: recent advances, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371 (1984) (2013) 20110534. [doi:10.1098/rsta.2011.0534](https://doi.org/10.1098/rsta.2011.0534).
- [62] B. Schölkopf, A. Smola, K.-R. Müller, Nonlinear Component Analysis as a Kernel Eigenvalue Problem, *Neural Computation* 10 (5) (1998) 1299–1319. [doi:10.1162/089976698300017467](https://doi.org/10.1162/089976698300017467).
- [63] L. McInnes, J. Healy, N. Saul, L. Großberger, UMAP: Uniform Manifold Approximation and Projection, *Journal of Open Source Software* 3 (29) (2018) 861. [doi:10.21105/joss.00861](https://doi.org/10.21105/joss.00861).
- [64] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine Learning in Python, *Journal of Machine Learning Research* 12 (2011) 2825–2830. [doi:10.48550/ARXIV.1201.0490](https://doi.org/10.48550/ARXIV.1201.0490).
- [65] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, A. Rush, Transformers: State-of-the-Art Natural Language Processing, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Association for Computational Linguistics, Online, 2020, pp. 38–45. [doi:10.18653/v1/2020.emnlp-demos.6](https://doi.org/10.18653/v1/2020.emnlp-demos.6).
- [66] V. Sanh, L. Debut, J. Chaumond, T. Wolf, DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter (2019). [doi:10.48550/ARXIV.1910.01108](https://doi.org/10.48550/ARXIV.1910.01108).

- [67] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, V. Stoyanov, Unsupervised Cross-lingual Representation Learning at Scale (2019). [doi:10.48550/ARXIV.1911.02116](https://doi.org/10.48550/ARXIV.1911.02116).
- [68] Z. Liu, W. Lin, Y. Shi, J. Zhao, A Robustly Optimized BERT Pre-Training Approach with Post-Training, in: Chinese Computational Linguistics: 20th China National Conference, CCL 2021, Hohhot, China, August 13–15, 2021, Proceedings, Springer-Verlag, Berlin, Heidelberg, 2021, p. 471–484. [doi:10.1007/978-3-030-84186-7_31](https://doi.org/10.1007/978-3-030-84186-7_31).
- [69] F. Feng, Y. Yang, D. Cer, N. Arivazhagan, W. Wang, Language-agnostic BERT Sentence Embedding, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics, Vol. 1, Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 878–891. [doi:10.18653/v1/2022.acl-long.62](https://doi.org/10.18653/v1/2022.acl-long.62).
- [70] N. Reimers, P. Beyer, I. Gurevych, Task-Oriented Intrinsic Evaluation of Semantic Textual Similarity, in: Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers, The COLING 2016 Organizing Committee, Osaka, Japan, 2016, pp. 87–96.
- [71] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, S. Bowman, GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding, in: Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 353–355. [doi:10.18653/v1/W18-5446](https://doi.org/10.18653/v1/W18-5446).
- [72] C. M. Bishop, Pattern Recognition and Machine Learning (Information Science and Statistics), Springer-Verlag, Berlin, Heidelberg, 2006.
- [73] C. Liu, Enhanced independent component analysis and its application to content based face image retrieval, IEEE Transactions on Systems, Man, and Cybernetics, Part B-Cybernetics 34 (2) (2004) 1117–1127. [doi:10.1109/TSMCB.2003.821449](https://doi.org/10.1109/TSMCB.2003.821449).
- [74] H. Kemal Ekenel, B. Sankur, Multiresolution face recognition, Image and Vision Computing 23 (5) (2005) 469–477. [doi:10.1016/j.imavis.2004.09.002](https://doi.org/10.1016/j.imavis.2004.09.002).
- [75] V. Laparra, G. Camps-Valls, J. Malo, Iterative Gaussianization: From ICA to Random Rotations, IEEE Transactions on Neural Networks 22 (4) (2011) 537–549. [doi:10.1109/TNN.2011.2106511](https://doi.org/10.1109/TNN.2011.2106511).
- [76] J. Cao, M. Spielmann, X. Qiu, X. Huang, D. M. Ibrahim, A. J. Hill, F. Zhang, S. Mundlos, L. Christiansen, F. J. Steemers, C. Trapnell, J. Shendure, The single-cell transcriptional landscape of mammalian organogenesis, Nature 566 (7745) (2019) 496+. [doi:10.1038/s41586-019-0969-x](https://doi.org/10.1038/s41586-019-0969-x).
- [77] S. Carter, Z. Armstrong, L. Schubert, I. Johnson, C. Olah, Activation Atlas, Distill<https://distill.pub/2019/activation-atlas> (2019). [doi:10.23915/distill.00015](https://doi.org/10.23915/distill.00015).