

# Fast, Effective, and Self-Supervised: Transforming Masked Language Models into Universal Lexical and Sentence Encoders

Fangyu Liu, Ivan Vulić, Anna Korhonen, Nigel Collier

Language Technology Lab, TAL, University of Cambridge

{f1399, iv250, alk23, nhc30}@cam.ac.uk

## Abstract

Previous work has indicated that pretrained Masked Language Models (MLMs) are not effective as universal lexical and sentence encoders off-the-shelf, i.e., without further task-specific fine-tuning on NLI, sentence similarity, or paraphrasing tasks using annotated task data. In this work, we demonstrate that it is possible to turn MLMs into effective lexical and sentence encoders even without any additional data, relying simply on self-supervision. We propose an extremely simple, fast, and effective contrastive learning technique, termed Mirror-BERT, which converts MLMs (e.g., BERT and RoBERTa) into such encoders in 20–30 seconds with no access to additional external knowledge. Mirror-BERT relies on identical *and* slightly modified string pairs as positive (i.e., synonymous) fine-tuning examples, and aims to maximise their similarity during “identity fine-tuning”. We report huge gains over off-the-shelf MLMs with Mirror-BERT both in lexical-level and in sentence-level tasks, across different domains and different languages. Notably, in sentence similarity (STS) and question-answer entailment (QNLI) tasks, our self-supervised Mirror-BERT model even matches the performance of the Sentence-BERT models from prior work which rely on annotated task data. Finally, we delve deeper into the inner workings of MLMs, and suggest some evidence on why this simple Mirror-BERT fine-tuning approach can yield effective universal lexical and sentence encoders.

## 1 Introduction

Transfer learning with pretrained Masked Language Models (MLMs) such as BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) has been widely successful in NLP, offering unmatched performance in a large number of tasks (Wang et al., 2019a). Despite the wealth of semantic knowledge stored in the MLMs (Rogers et al., 2020), they do not produce high-quality lexical and sentence embeddings when used off-the-shelf, without further




Figure 1: Illustration of the main concepts behind the proposed self-supervised Mirror-BERT method. The same text sequence can be observed from two additional “views”: 1) by performing random span masking in the input space, and/or 2) by applying dropout (inside the BERT/RoBERTa MLM) in the feature space, yielding identity-based (i.e., “mirrored”) positive examples for fine-tuning. A contrastive learning objective is then applied to encourage such “mirrored” positive pairs to obtain more similar representations in the embedding space relatively to negative pairs.

task-specific fine-tuning (Feng et al., 2020; Li et al., 2020). In fact, previous work has shown that their performance is sometimes even below static word embeddings and specialised sentence encoders (Cer et al., 2018) in lexical and sentence-level semantic similarity tasks (Reimers and Gurevych, 2019; Vulić et al., 2020b; Litschko et al., 2021).

In order to address this gap, recent work has trained dual-encoder networks on labelled external resources to convert MLMs into universal language encoders. Most notably, Sentence-BERT (SBERT, Reimers and Gurevych 2019) further trains BERT and RoBERTa on Natural Language Inference (NLI, Bowman et al. 2015; Williams et al. 2018) and sentence similarity data (Cer et al., 2017) to obtain high-quality universal sentence embeddings. Recently, SapBERT (Liu et al., 2021) self-

aligns phrasal representations of the same meaning using synonyms extracted from the UMLS (Bodenreider, 2004), a large biomedical knowledge base, obtaining lexical embeddings in the biomedical domain that reach state-of-the-art (SotA) performance in biomedical entity linking tasks. However, both SBERT and SapBERT require annotated (i.e., human-labelled) data as external knowledge: it is used to instruct the model to produce similar representations for text sequences (e.g., words, phrases, sentences) of similar/identical meanings.

In this paper, we fully dispose of any external supervision, demonstrating that the transformation of MLMs into universal language encoders can be achieved without task-labelled data. We propose a fine-tuning framework termed *Mirror-BERT*, which simply relies on duplicating and slightly augmenting the existing text input (or their representations) to achieve the transformation, and show that it is possible to learn universal lexical and sentence encoders with such “mirrored” input data through self-supervision (see Fig. 1). The proposed Mirror-BERT framework is also extremely efficient: the whole MLM transformation can be completed in less than one minute on two 2080Ti GPUs.

Our findings further confirm a general hypothesis from prior work (Liu et al., 2021; Ben-Zaken et al., 2020; Glavaš and Vulić, 2021) that fine-tuning exposes the wealth of (semantic) knowledge stored in the MLMs. In this case in particular, we demonstrate that the Mirror-BERT procedure can rewire the MLMs to serve as universal language encoders even without any external supervision. We further show that data augmentation in both input space and feature space are key to the success of Mirror-BERT, and they provide a synergistic effect.

**Contributions.** **1)** We propose a completely self-supervised approach that can quickly transform pretrained MLMs into capable universal lexical and sentence encoders, greatly outperforming off-the-shelf MLMs in similarity tasks across different languages and domains. **2)** We investigate the rationales behind why Mirror-BERT works at all, aiming to understand the impact of data augmentation in the input space as well as in the feature space. We release our code and models at <https://github.com/cambridgeltl/mirror-bert>.

## 2 Mirror-BERT: Methodology

Mirror-BERT consists of three main parts, described in what follows. First, we create positive

pairs by duplicating the input text (§2.1). We then further process the positive pairs by simple data augmentation operating either on the input text or on the feature map inside the model (§2.2). Finally, we apply standard contrastive learning, ‘attracting’ the texts belonging to the same class (i.e., positives) while pushing away the negatives (§2.3).

### 2.1 Training Data through Self-Duplication

The key to success of dual-network representation learning (Henderson et al., 2019; Reimers and Gurevych, 2019; Humeau et al., 2020; Liu et al., 2021, *inter alia*) is the construction of positive and negative pairs. While negative pairs can be easily obtained from randomly sampled texts, positive pairs usually need to be manually annotated. In practice, they are extracted from labelled task data (e.g., NLI) or knowledge bases that store relations such as synonymy or hypernymy (e.g., PPDB, Pavlick et al. 2015; BabelNet, Ehrmann et al. 2014; WordNet, Fellbaum 1998; UMLS).

Mirror-BERT, however, does not rely on any external data to construct the positive examples. In a nutshell, given a set of non-duplicated strings  $\mathcal{X}$ , we assign individual labels ( $y_i$ ) to each string and build a dataset  $\mathcal{D} = \{(x_i, y_i) | x_i \in \mathcal{X}, y_i \in \{1, \dots, |\mathcal{X}|\}\}$ . We then create self-duplicated training data  $\mathcal{D}'$  simply by repeating every element in  $\mathcal{D}$ . In other words, let  $\mathcal{X} = \{x_1, x_2, \dots\}$ . We then have  $\mathcal{D} = \{(x_1, y_1), (x_2, y_2), \dots\}$  and  $\mathcal{D}' = \{(x_1, y_1), (\bar{x}_1, \bar{y}_1), (x_2, y_2), (\bar{x}_2, \bar{y}_2), \dots\}$  where  $x_1 = \bar{x}_1, y_1 = \bar{y}_1, x_2 = \bar{x}_2, y_2 = \bar{y}_2, \dots$ . In §2.2, we introduce data augmentation techniques (in both input space and feature space) applied on  $\mathcal{D}'$ . Each positive pair  $(x_i, \bar{x}_i)$  yields two different points/vectors in the encoder’s representation space (see again Fig. 1), and the distance between these points should be minimised.

### 2.2 Data Augmentation

We hypothesise that applying certain ‘corruption’ techniques to (i) parts of input text sequences or (ii) to their representations, or even (iii) doing both in combination, does not change their (captured) meaning. We present two ‘corruption’ techniques as illustrated in Fig. 1. First, we can directly mask parts of the input text. Second, we can erase (i.e., dropout) parts of their feature maps. Both techniques are rather simple and intuitive: (i) even when masking parts of an input sentence, humans can usually reconstruct its semantics; (ii) dropping a small subset of neurons or representation dimen-

x1: Economist Paul Krugman mainly works on trade models.  
x̄1: Econ [MASK] Paul Krugman mainly works on trade models.

Figure 2: An example of input data augmentation via random span masking.

sions, the representations of a neural network will not drift too much.

**Input Augmentation: Random Span Masking.** The idea is inspired by random cropping in visual representation learning (Hendrycks et al., 2020). In particular, starting from the mirrored pairs  $(x_i, y_i)$  and  $(\bar{x}_i, \bar{y}_i)$ , we randomly replace a consecutive string of length  $k$  with [MASK] in either  $x_i$  or  $\bar{x}_i$ . The example (Fig. 2) illustrates the random span masking procedure with  $k = 5$ .

**Feature Augmentation: Dropout.** The random span masking technique, operating directly on text input, can be applied only with sentence/phrase-level input; word-level tasks involve only short strings, usually represented as a single token under the sentence-piece tokeniser. However, data augmentation in the feature space based on dropout, as introduced below, can be applied to any input text.

Dropout (Srivastava et al., 2014) randomly drops neurons from a neural net during training with a probability  $p$ . In practice, it results in the erasure of each element with a probability of  $p$ . It has mostly been interpreted as implicitly bagging a large number of neural networks which share parameters at test time (Bouthillier et al., 2015). Here, we take advantage of the dropout layers in BERT/roBERTa to create augmented views of the input text. Given a pair of identical strings  $x_i$  and  $\bar{x}_i$ , their representations in the embedding space slightly differ due to the existence of multiple dropout layers in the BERT/roBERTa architecture (Fig. 6). The two data points in the embedding space can be seen as two augmented views of the same text sequence, which can be leveraged for fine-tuning.<sup>1</sup>

It is possible to combine data augmentation via random span masking and feature augmentation via dropout; this variant is also evaluated later.

### 2.3 Contrastive Learning

Let  $f(\cdot)$  denote the encoder model. The encoder is then fine-tuned on the data constructed in §2.2.

<sup>1</sup>The dropout augmentations are naturally a part of the BERT/roBERTa network. That is, no further actions need to be taken to implement them. Note that random span masking is applied on only one side of the positive pair while dropout is applied on all data points.




Figure 3: As the same vector goes through the same dropout layer separately, the outcomes are independent. Consequently, two fully identical strings fed to the single BERT/roBERTa model yield different representations in the MLM embedding space.

Given a batch of data  $\mathcal{D}'_b$ , we leverage the standard InfoNCE loss (Oord et al., 2018) to cluster/attract the positive pairs together and push away the negative pairs in the embedding space:

$$\mathcal{L}_b = - \sum_{i=1}^{|\mathcal{D}'_b|} \log \frac{\exp(\cos(f(x_i), f(\bar{x}_i))/\tau)}{\sum_{x_j \in \mathcal{N}_i} \exp(\cos(f(x_i), f(x_j))/\tau)}. \quad (1)$$

$\tau$  denotes a temperature parameter;  $\mathcal{N}_i$  denotes all negatives of  $x_i$ , which includes all  $x_j, \bar{x}_j$  where  $i \neq j$  in the current data batch (i.e.,  $|\mathcal{N}_i| = |\mathcal{D}'_b| - 2$ ). Intuitively, the numerator is the similarity of the self-duplicated pair (the positive example) and the denominator is the sum of the similarities between  $x_i$  and all other strings besides  $\bar{x}_i$  (the negatives).<sup>2</sup>

## 3 Experimental Setup

**Evaluation Tasks: Lexical.** We evaluate on domain-general and domain-specific tasks: word similarity and biomedical entity linking (BEL). For the former, we rely on the Multi-SimLex evaluation set (Vulić et al., 2020a): it contains human-elicited word similarity scores for multiple languages. For the latter, we use NCBI-disease (NCBI, Doğan et al. 2014), BC5CDR-disease, BC5CDR-chemical (BC5-d, BC5-c, Li et al. 2016), AskAPatient (Limsopatham and Collier, 2016) and COMETA (stratified-general split, Basaldella et al. 2020) as our evaluation datasets. The first three datasets are in the scientific domain (i.e., the data have been extracted from scientific papers), while the latter two

<sup>2</sup>We also experimented with another state-of-the-art contrastive learning scheme proposed by Liu et al. (2021). There, hard triplet mining combined with multi-similarity loss (MS loss) is used as the learning objective. InfoNCE and triplet mining + MS loss work mostly on par, with slight gains of one variant in some tasks, and vice versa. For simplicity and brevity, we report the results only with InfoNCE.

are in the social media domain (i.e., extracted from online forums discussing health-related topics). We report Spearman’s rank correlation coefficients ( $\rho$ ) for word similarity; accuracy @1/@5 is the standard evaluation measure in the BEL task.

**Evaluation Tasks: Sentence-Level.** Evaluation on the intrinsic sentence textual similarity (STS) task is conducted on the standard SemEval 2012-2016 datasets (Agirre et al., 2012, 2013, 2014, 2015, 2016), STS Benchmark (STS-b, Cer et al. 2017), SICK-Relatedness (SICK-R, Marelli et al. 2014) for English; STS SemEval-17 data is used for Spanish and Arabic (Cer et al., 2017), and we also evaluate on Russian STS.<sup>3</sup> We report Spearman’s  $\rho$  rank correlation. Evaluation in the question-answer entailment task is conducted on QNLI (Rajpurkar et al., 2016; Wang et al., 2019b). It contains 110k English QA pairs with binary entailment labels.<sup>4</sup>

**Evaluation Tasks: Cross-Lingual.** We also assess the benefits of Mirror-BERT on cross-lingual representation learning, evaluating on cross-lingual word similarity (CLWS, Multi-SimLex is used) and bilingual lexicon induction (BLI). We rely on the standard mapping-based BLI setup (Artetxe et al., 2018), and training and test sets from Glavaš et al. (2019), reporting accuracy @1 scores (with CSLS as the word retrieval method, Lample et al. 2018).

**Mirror-BERT: Training Resources.** For fine-tuning (general-domain) lexical representations, we use the top 10k most frequent words in each language. For biomedical name representations, we randomly sample 10k names from the UMLS. In sentence-level tasks, for STS, we sample 10k sentences (without labels) from the training set of the STS Benchmark; for Spanish, Arabic and Russian, we sample 10k sentences from the Wiki-Matrix dataset (Schwenk et al., 2021). For QNLI, we sample 10k sentences from its training set.

**Training Setup and Details.** The hyperparameters of word-level models are tuned on SimLex-999 (Hill et al., 2015); biomedical models are tuned on COMETA (zero-shot-general split). Sentence-level models are tuned on the dev set of STS-b.  $\tau$  in Eq. (1) is 0.04 (biomedical and sentence-level models); 0.2 (word-level). Dropout rate  $p$  is 0.1. Sentence-level models use a random span masking

<sup>3</sup>[github.com/deepmipt/deepPavlovEval](https://github.com/deepmipt/deepPavlovEval)

<sup>4</sup>We follow the setup of Li et al. (2020) and adapt QNLI to an unsupervised task by computing the AUC scores (on the development set,  $\approx 5.4k$  pairs) using 0/1 labels and cosine similarity scores of QA embeddings.

lang. →	EN	FR	ET	AR	ZH	RU	ES	PL	avg.
fastText	<u>.528</u>	<u>.560</u>	<b>.447</b>	<u>.409</u>	.428	<b>.435</b>	<b>.488</b>	<u>.396</u>	<u>.461</u>
BERT	.267	.020	.106	.220	.398	.202	.177	.217	.201
+ Mirror	<b>.556</b>	<b>.621</b>	<u>.308</u>	<b>.538</b>	<b>.639</b>	<u>.365</u>	.296	<b>.444</b>	<b>.471</b>
mBERT	.105	.130	.094	.101	.261	.109	.095	.087	.123
+ Mirror	<u>.358</u>	<u>.341</u>	<u>.134</u>	<u>.097</u>	<u>.501</u>	.210	<u>.332</u>	<u>.141</u>	<u>.264</u>

Table 1: Word similarity evaluation on Multi-SimLex. “BERT” denotes monolingual BERT models in each language (see the Appendix). “mBERT” denotes multilingual BERT. **Bold** and underline denote highest and second-highest scores per column, respectively.

dataset → model ↓	scientific language				social media language					
	NCBI		BC5-d		BC5-c		AskAPatient		COMETA	
	@1	@5	@1	@5	@1	@5	@1	@5	@1	@5
SapBERT	<b>.920</b>	<b>.956</b>	<b>.935</b>	<b>.960</b>	<b>.965</b>	<b>.982</b>	<b>.705</b>	<b>.889</b>	<b>.659</b>	<b>.779</b>
BERT	.676	.770	.815	.891	.798	.912	.382	.433	.404	.477
+ Mirror	<u>.872</u>	<u>.921</u>	<u>.921</u>	<u>.949</u>	<u>.957</u>	<u>.971</u>	<u>.555</u>	<u>.695</u>	<u>.547</u>	<u>.647</u>
PubMedBERT	.778	.869	.890	.938	.930	.946	.425	.496	.468	.532
+ Mirror	<u>.909</u>	<u>.948</u>	<u>.930</u>	<u>.962</u>	<u>.958</u>	<u>.979</u>	<u>.590</u>	<u>.750</u>	<u>.603</u>	<u>.713</u>

Table 2: Biomedical entity linking (BEL) evaluation.

rate of  $k = 5$ , while  $k = 2$  for biomedical phrase-level models; we do not employ span masking for word-level models (an analysis is in the Appendix). All lexical models are trained for 2 epochs, max token length is 25. Sentence-level models are trained for 1 epoch with a max sequence length of 50.

All models use AdamW (Loshchilov and Hutter, 2019) as the optimiser, with a learning rate of  $2e-5$ , batch size of 200 (400 after duplication). In all tasks, for all ‘Mirror-tuned’ models, unless noted otherwise, we create final representations using [CLS], instead of another common option: *mean-pooling* ( $mp$ ) over all token representations in the last layer (Reimers and Gurevych, 2019).<sup>5 6</sup>

## 4 Results and Discussion

### 4.1 Lexical-Level Tasks

**Word Similarity (Tab. 1).** SotA static word embeddings such as fastText (Mikolov et al., 2018) typically outperform off-the-shelf MLMs on word similarity datasets (Vulić et al., 2020a). However, our results demonstrate that the Mirror-BERT procedure indeed converts the MLMs into much stronger word encoders. The Multi-SimLex results on 8 languages from Tab. 1 suggest that the fine-

<sup>5</sup>For ‘non-Mirrored’ original MLMs, the results with  $mp$  are reported instead; they produce much better results than using [CLS]; see the Appendix.

<sup>6</sup>All reported results are averages of three runs. In general, the training is very stable, with negligible fluctuations.

model↓, dataset→	STS12	STS13	STS14	STS15	STS16	STS-b	SICK-R	avg.
SBERT*	<b>.719</b>	.774	<b>.742</b>	.799	.747	.774	<b>.721</b>	<u>.754</u>
BERT-CLS	.215	.321	.213	.379	.442	.203	.427	.314
BERT-mp	.314	.536	.433	.582	.596	.464	.528	.493
+ Mirror	.670	.801	.713	.812	.743	.764	.699	.743
+ Mirror (drophead)	<u>.691</u>	.811	.730	<u>.819</u>	.757	.780	.691	<u>.754</u>
RoBERTa-CLS	.090	.327	.210	.338	.388	.317	.355	.289
RoBERTa-mp	.134	.126	.124	.203	.224	.129	.320	.180
+ Mirror	.646	<u>.818</u>	.734	.802	<u>.782</u>	<u>.787</u>	<u>.703</u>	<u>.753</u>
+ Mirror (drophead)	.666	<b>.827</b>	<u>.740</u>	<b>.824</b>	<b>.797</b>	<b>.796</b>	.697	<b>.764</b>

Table 3: English STS. \*We were able to reproduce the scores reported in the original Sentence-BERT (SBERT, Reimers and Gurevych 2019) paper. However, we found mean-pooling over all tokens (including padded ones) yield better performance (.754 vs .749). We thus report the stronger baseline.

model↓, lang.→	ES	AR	RU	avg.
BERT	.599	.455	.552	.533
+ Mirror	<u>.709</u>	<b>.669</b>	<u>.673</u>	<b>.684</b>
mBERT	.610	.447	.616	.558
+ Mirror	<b>.755</b>	<u>.594</u>	<b>.692</b>	<u>.680</u>

Table 4: STS evaluation in other languages.

tuned +*Mirror* variant substantially improves the performance of base MLMs (both monolingual and multilingual ones), even beating fastText in 5 out of the 8 evaluation languages.<sup>7</sup>

We also observe that it is essential to have a strong base MLM. While Mirror-BERT does offer substantial performance gains with all base MLMs, the improvement is more pronounced when the base model is strong (e.g., EN, ZH).

**Biomedical Entity Linking (Tab. 2).** The goal of BEL is to map a biomedical name mention to a controlled vocabulary (usually a node in a knowledge graph). Considered a downstream application in BioNLP, the BEL task also helps evaluate and compare the quality of biomedical name representations: it requires pairwise comparisons between the biomedical mention and all surface strings stored in the biomedical knowledge graph.

The results from Tab. 2 suggest that our +*Mirror* transformation achieves very strong gains on top of the base MLMs, both BERT and PubMedBERT (Gu et al., 2020). We note that PubMedBERT is a current SotA MLM in the biomedical domain, and performs significantly better than BERT, both before and after +*Mirror* fine-tuning. This highlights the necessity of starting from a domain-specific model when possible. On scientific datasets, the self-supervised *PubMedBERT+Mirror* model is

<sup>7</sup>Language codes: see the Appendix for a full listing.




Figure 4: Unsupervised QNLI: ROC curves and AOC scores.

very close to SapBERT, which fine-tunes PubMedBERT with more than 10 million synonyms extracted from the external UMLS knowledge base.

However, in the social media domain, *PubMedBERT+Mirror* still cannot match the performance of knowledge-guided SapBERT. This in fact reflects the nature and complexity of the task domain. For the three datasets in the scientific domain (NCBI, BC5-d, BC5-c), strings with similar surface forms tend to be associated with the same concept. On the other hand, in the social media domain, semantics of very different surface strings might be the same.<sup>8</sup> This also suggests that Mirror-BERT adapts PubMedBERT to a very good surface-form encoder for biomedical names, but dealing with more difficult synonymy relations (e.g. as found in the social media) does need external knowledge.<sup>9</sup>

## 4.2 Sentence-Level Tasks

**English STS (Tab. 3).** Regardless of the base model (BERT/RoBERTa), applying +*Mirror* fine-

<sup>8</sup>For instance, *HCQ* and *Plaquenil* refer to exactly the same concept on online health forums: *Hydroxychloroquine*.

<sup>9</sup>Motivated by these insights, in future work we will also investigate a combined approach that blends self-supervision and external knowledge (Vulić et al., 2021), which could also be automatically mined (Su, 2020; Thakur et al., 2021).

lang.→	EN-FR	EN-ZH	EN-HE	FR-ZH	FR-HE	ZH-HE	avg.
mBERT	.163	.118	.071	.142	.104	.010	.101
+ Mirror	<b>.454</b>	<b>.385</b>	<b>.133</b>	<b>.465</b>	<b>.163</b>	<b>.179</b>	<b>.297</b>

Table 5: Cross-lingual word similarity results.

lang.→	EN-FR	EN-IT	EN-RU	EN-TR	IT-FR	RU-FR	avg.
BERT	.014	.112	.154	.150	.025	.018	.079
+ Mirror	<b>.458</b>	<b>.378</b>	<b>.336</b>	<b>.289</b>	<b>.417</b>	<b>.345</b>	<b>.371</b>

Table 6: BLI results.

tuning greatly boosts performance across all English STS datasets. Surprisingly, on average, *RoBERTa+Mirror*, fine-tuned with only 10k sentences without any external supervision, is on-par with the SBERT model, which is trained on the merged SNLI (Bowman et al., 2015) and MultiNLI (Williams et al., 2018) datasets, containing 570k and 430k sentence pairs, respectively.

**Spanish, Arabic and Russian STS (Tab. 4).** The results in the STS tasks on other languages, which all have different scripts, again indicate very large gains, using both monolingual language-specific BERTs and mBERT as base MLMs. This confirms that Mirror-BERT is a language-agnostic method.

**Question-Answer Entailment (Fig. 4).** The results indicate that our *+Mirror* fine-tuning consistently improves the underlying MLMs. The *RoBERTa+Mirror* variant even shows a slight edge over the supervised SBERT model (.709 vs. .706).

### 4.3 Cross-Lingual Tasks

We observe huge gains across all language pairs in CLWS (Tab. 5) and BLI (Tab. 6) after running the Mirror-BERT procedure. For language pairs that involve Hebrew, the improvement is usually smaller. We suspect that this is due to mBERT itself containing poor semantic knowledge for Hebrew. This finding aligns with our prior argument that a strong base MLM is still required to obtain prominent gains from running Mirror-BERT.

### 4.4 Further Discussion and Analyses

**Running Time.** The Mirror-BERT procedure is extremely time-efficient. While fine-tuning on NLI (SBERT) or UMLS (SapBERT) data can take hours, Mirror-BERT with 10k positive pairs completes the conversion from MLMs to universal language encoders within a minute on two NVIDIA RTX 2080Ti GPUs. On average, 10-20 seconds is needed for 1 epoch of the Mirror-BERT procedure.




Figure 5: The impact of the number of fine-tuning “mirrored” examples ( $x$ -axis) on the task performance ( $y$ -axis). The scores across tasks are not directly comparable, and are based on different evaluation metrics (§3).

**Input Data Size (Fig. 5).** In our main experiments in §4.1-§4.3, we always use 10k examples for Mirror-BERT tuning. In order to assess the importance of the fine-tuning data size, we run a relevant analysis for a subset of base MLMs, and on a subset of English tasks. In particular, we evaluate the following: (i) BERT, Multi-SimLex (EN) (word-level); (ii) PubMedBERT, COMETA (biomedical phrase-level); (iii) RoBERTa, STS12 (sentence-level). The results indicate that the performance in all tasks reaches its peak in the region of 10k-20k examples and then gradually decreases, with a steeper drop on the the word-level task.<sup>10 11</sup>

**Random Span Masking + Dropout? (Tab. 7).** We conduct our ablation studies on the English STS tasks. First, we experiment with turning off dropout, random span masking, or both. With both techniques turned off, we observe large performance drops of *RoBERTa+Mirror* and *BERT+Mirror* (see also the Appendix). Span masking appears to be the more important factor: its absence causes a larger decrease. However, the best performance is achieved when both dropout and random span masking are leveraged, suggesting a synergistic effect when the two augmentation techniques are used together.

**Other Data Augmentation Types? Dropout vs. Drophead (Tab. 7).** Encouraged by the effectiveness of random span masking and dropout for Mirror-BERT, a natural question to pose is: can

<sup>10</sup>We suspect that this is due to the inclusion of lower-frequency words into the fine-tuning data: embeddings of such words typically obtain less reliable embeddings (Pilehvar et al., 2018).

<sup>11</sup>For word-level experiments, we used the top 100k words in English according to Wikipedia statistics. For phrase-level experiments, we randomly sampled 100k names from UMLS. For sentence-level experiments we sampled 100k sentences from SNLI and MultiNLI datasets (as the STS training set has fewer than 100k sentences).

model configuration	avg. $\rho$
RoBERTa + Mirror	.753
- dropout + drophead	.764 $\uparrow$ .011
- dropout	.732 $\downarrow$ .021
- span mask	.717 $\downarrow$ .036
- dropout & span mask	.682 $\downarrow$ .071

Table 7: Ablation study: (i) replacing dropout with drophead; (ii) the synergistic effect of dropout and random span masking in the English STS tasks.




Figure 6: Under controlled dropout, if two strings are identical, they will have an identical set of dropout masks throughout the encoding process.

other augmentation types work as well? Recent work points out that pretrained MLM are heavily overparameterised and most Transformer heads can be pruned without hurting task performance (Voita et al., 2019; Kovaleva et al., 2019; Michel et al., 2019). Zhou et al. (2020) propose a *drop-head* method: it randomly prunes attention heads at MLM training as a regularisation step. We thus evaluate a variant of Mirror-BERT where the dropout layers are replaced with such dropheads:<sup>12</sup> this results in even stronger STS performance, cf. Tab. 7. In short, this hints that the Mirror-BERT framework might benefit from other data and feature augmentation techniques in future work.<sup>13</sup>

**Regularisation or Augmentation? (Tab. 8).** When using dropout, is it possible that we are simply observing the effect of adding/removing regularisation instead of the augmentation benefit? To answer this question, we design a simple probe that attempts to disentangle the effect of regular-

<sup>12</sup>Drophead rates for BERT and RoBERTa are set to the default values of 0.2 and 0.05, respectively.

<sup>13</sup>Besides the drophead-based feature space augmentation, in our side experiments, we also tested input space augmentation techniques such as whole-word masking, random token masking, and word reordering; they typically yield performance similar or worse to random span masking. We also point to very recent work that explores text augmentation in a different context (Wu et al., 2020; Meng et al., 2021). We leave a thorough search of optimal augmentation techniques for future work.

model configuration (MLM=RoBERTa)	$\rho$ on STS12
random span masking $\times$ ; dropout $\times$	.562
random span masking $\times$ ; dropout $\checkmark$	.648 $\uparrow$ .086
random span masking $\times$ ; controlled dropout $\checkmark$	.452 $\downarrow$ .110

Table 8: Probing the impact of dropout.

isation versus augmentation; we turn off random span masking but leave the dropout on (so that the regularisation effect remains). However, instead of assigning independent dropouts to every individual string (rendering each string slightly different), we control the dropouts applied to a positive pair to be identical. As a result, it holds  $f(x_i) = f(\bar{x}_i)$ , when  $x_i \equiv \bar{x}_i, \forall i \in \{1, \dots, |\mathcal{D}|\}$ . We denote this as “controlled dropout”. In Tab. 8, we observe that, during the +*Mirror* fine-tuning, controlled dropout largely underperforms standard dropout and is even worse than not using dropout at all. As the only difference between controlled and standard dropout is the augmented features for positive pairs in the latter case, this suggests that the gain from +*Mirror* indeed stems from the data augmentation effect rather than from regularisation.

**Mirror-BERT Improves Isotropy? (Fig. 7).** We argue that the gains with Mirror-BERT largely stem from its reshaping of the embedding space geometry. *Isotropy* (i.e., uniformity in all orientations) of the embedding space has been a favourable property for semantic similarity tasks (Arora et al., 2016; Mu and Viswanath, 2018). However, Ethayarajh (2019) shows that (off-the-shelf) MLMs’ representations are anisotropic: they reside in a narrow cone in the vector space and the average cosine similarity of (random) data points is extremely high. Sentence embeddings induced from MLMs without fine-tuning thus suffer from spatial anisotropy (Li et al., 2020; Su et al., 2021). Is Mirror-BERT then improving isotropy of the embedding space?<sup>14</sup> To investigate this claim, we inspect (1) the distributions of cosine similarities and (2) an isotropy score, as defined by Mu and Viswanath (2018).

First, we randomly sample 1,000 sentence pairs from the Quora Question Pairs (QQP) dataset. In

<sup>14</sup>Some preliminary evidence from Tab. 7 already leads in this direction: we observe large gains over the base MLMs even without any positive examples, that is, when both span masking and dropout are not used (i.e., it always holds  $x_i = \bar{x}_i$  and  $f(x_i) = f(\bar{x}_i)$ ). During training, this leads to a constant numerator in Eq. (1). In this case, learning collapses to the scenario where all gradients solely come from the negatives: the model is simply pushing all data points away from each other, resulting in a more isotropic space.




Figure 7: Cosine similarity distribution over 1k sentence pairs sampled from QQP. Blue and orange mean positive and negative similarities, respectively.

Fig. 7, we plot the distributions of pairwise cosine similarities of BERT representations before (Figs. 7a and 7b) and after the +Mirror tuning (Fig. 7c). The overall cosine similarities (regardless of positive/negative) are greatly reduced and the positives/negatives become easily separable.

We also leverage a quantitative isotropy score (IS), proposed in prior work (Arora et al., 2016; Mu and Viswanath, 2018), and defined as follows:

$$IS(\mathcal{V}) = \frac{\min_{\mathbf{c} \in \mathcal{C}} \sum_{\mathbf{v} \in \mathcal{V}} \exp(\mathbf{c}^\top \mathbf{v})}{\max_{\mathbf{c} \in \mathcal{C}} \sum_{\mathbf{v} \in \mathcal{V}} \exp(\mathbf{c}^\top \mathbf{v})} \quad (2)$$

where  $\mathcal{V}$  is the set of vectors,<sup>15</sup>  $\mathcal{C}$  is the set of all possible unit vectors (i.e., any  $\mathbf{c}$  so that  $\|\mathbf{c}\| = 1$ ) in the embedding space. In practice,  $\mathcal{C}$  is approximated by the eigenvector set of  $\mathbf{V}^\top \mathbf{V}$  ( $\mathbf{V}$  is the stacked embeddings of  $\mathcal{V}$ ). The larger the IS value, more isotropic an embedding space is (i.e., a perfectly isotropic space obtains the IS score of 1).

IS scores in Tab. 9 confirm that the +Mirror fine-tuning indeed makes the embedding space more isotropic. Interestingly, with both data augmentation techniques switched off, a naive expectation is that IS will increase as the gradients now solely come from negative examples, pushing apart points in the space. However, we observe the increase of IS only for word-level representations. This hints at more complex dynamics between isotropy and gradients from positive and negative examples, where positives might also contribute to isotropy in some settings. We will examine these dynamics more in future work.<sup>16</sup>

**Learning New Knowledge or Exposing Available Knowledge?** Running Mirror-BERT for more epochs, or with more data (see Fig. 5) does not re-

<sup>15</sup> $\mathcal{V}$  comprises the corresponding text data used for Mirror-BERT fine-tuning (10k items for each task type).

<sup>16</sup>Introducing positive examples also naturally yields stronger task performance, as the original semantic space is better preserved. Gao et al. (2021) provide an insightful analysis on the balance of learning uniformity and alignment preservation, based on the method of Wang and Isola (2020).

level→	word	phrase	sentence
BERT	.169	.205	.222
+ Mirror	.599	<b>.252</b>	<b>.265</b>
+ Mirror (w/o aug.)	<b>.825</b>	.170	.255

Table 9: IS of word, phrase, and sentence-level models.

model	$\rho$
fastText	.528
BERT-CLS	.105
BERT-mp	.267
+ Mirror	<b>.556</b>
+ Mirror (random string)	.393
+ Mirror (random string, lr $5e-5$ )	.481

Table 10: Running Mirror-BERT with a set of ‘zero-semantics’ random strings. Evaluation is conducted on Multi-SimLex (EN).




Figure 8: Blue: words in Multi-SimLex (EN) follow a long-tail distribution. Yellow: BERT+Mirror trained with frequent words tend to perform better.

sult in performance gains. This hints that, instead of gaining new knowledge from the fine-tuning data, Mirror-BERT in fact ‘rewires’ existing knowledge in MLMs (Ben-Zaken et al., 2020). To further verify this, we run Mirror-BERT with random ‘zero-semantics’ words, generated by uniformly sampling English letters and digits, and evaluate on (EN) Multi-SimLex. Surprisingly, even these data can transform off-the-shelf MLMs into effective word encoders: we observe a large improvement over the base MLM in this extreme scenario, from  $\rho = 0.267$  to 0.481 (Tab. 10). We did a similar experiment on the sentence-level and observed similar trends. However, we note that using the actual English texts for fine-tuning still performs better as they are more ‘in-domain’ (with further evidence and discussions in the following paragraph).

**Selecting Examples for Fine-Tuning.** Using raw text sequences from the end task should be the default option for Mirror-BERT fine-tuning since they are in-distribution by default, as semantic similarity models tend to underperform when faced with a domain shift (Zhang et al., 2020). In the *general-domain* STS tasks, we find that using sentences



extracted from the STS training set, Wikipedia articles, or NLI datasets all yield similar STS performance after running Mirror-BERT (though optimal hyperparameters differ). However, porting *BERT+Mirror* trained on STS data to QNLI results in AUC drops from .674 to .665. This suggests that slight or large domain shifts do affect task performance, further corroborated by our findings from fine-tuning with fully random strings (see before).

Further, Fig. 8 shows a clear tendency that more frequent strings are more likely to yield good task performance. There, we split the 100k most frequent words from English Wikipedia into 10 equally sized fine-tuning buckets of 10k examples each, and run *+Mirror* fine-tuning on BERT with each bucket. In sum, using frequent in-domain examples seems to be the optimal choice.

## 5 Related Work

Self-supervised text representations have a large body of literature. Here, due to space constraints, we provide a highly condensed summary of the most related work. Even prior to the emergence of large pretrained LMs (PLMs), most representation models followed the distributional hypothesis (Harris, 1954) and exploited the co-occurrence statistics of words/phrases/sentences in large corpora (Mikolov et al., 2013a,b; Pennington et al., 2014; Kiros et al., 2015; Hill et al., 2016; Logeswaran and Lee, 2018). Recently, DeCLUTR (Giorgi et al., 2021) follows the distributional hypothesis and formulates sentence embedding training as a contrastive learning task where span pairs sampled from the same document are treated as positive pairs. Very recently, there has been a growing interest in using individual raw sentences for self-supervised contrastive learning on top of PLMs.

Wu et al. (2020) explore input augmentation techniques for sentence representation learning with contrastive objectives. However, they use it as an auxiliary loss during full-fledged MLM pre-training from scratch (Rethmeier and Augenstein, 2021). In contrast, our post-hoc approach offers a lightweight and fast self-supervised transformation from any pretrained MLM to a universal language encoder at lexical or sentence level.

Carlsson et al. (2021) use two distinct models to produce two views of the same text, where we rely on a single model, that is, we propose to use dropout and random span masking within the same model to produce the two views, and demonstrate

their synergistic effect. Our study also explores word-level and phrase-level representations and tasks, and to domain-specialised representations (e.g., for the BEL task).

SimCSE (Gao et al., 2021), a work concurrent to ours, adopts the same contrastive loss as Mirror-BERT, and also indicates the importance of data augmentation through dropout. However, they do not investigate random span masking as data augmentation in the input space, and limit their model to general-domain English sentence representations only, effectively rendering SimCSE a special case of the Mirror-BERT framework. Other concurrent papers explore a similar idea, such as Self-Guided Contrastive Learning (Kim et al., 2021), ConSERT (Yan et al., 2021), and BSL (Zhang et al., 2021), *inter alia*. They all create two views of the same sentence for contrastive learning, with different strategies in feature extraction, data augmentation, model updating or choice of loss function. However, they offer less complete empirical findings compared to our work: we additionally evaluate on (1) lexical-level tasks, (2) tasks in a specialised biomedical domain and (3) cross-lingual tasks.

## 6 Conclusion

We proposed Mirror-BERT, a simple, fast, self-supervised, and highly effective approach that transforms large pretrained masked language models (MLMs) into universal lexical and sentence encoders within a minute, and without any external supervision. Mirror-BERT, based on simple unsupervised data augmentation techniques, demonstrates surprisingly strong performance in (word-level and sentence-level) semantic similarity tasks, as well as on biomedical entity linking. The large gains over base MLMs are observed for different languages with different scripts, and across diverse domains. Moreover, we dissected and analysed the main causes behind Mirror-BERT’s efficacy.

## Acknowledgements

We thank the reviewers and the AC for their considerate comments. We also thank the LTL members and Xun Wang for insightful feedback. FL is supported by Grace & Thomas C.H. Chan Cambridge Scholarship. AK and IV are supported by the ERC Grant LEXICAL (no. 648909) and the ERC PoC Grant MultiConvAI (no. 957356). NC kindly acknowledges grant-in-aid funding from ESRC (grant number ES/T012277/1).

## References

- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Iñigo Lopez-Gazpio, Montse Maritxalar, Rada Mihalcea, German Rigau, Larraitz Uribe, and Janyce Wiebe. 2015. [SemEval-2015 task 2: Semantic textual similarity, English, Spanish and pilot on interpretability](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 252–263, Denver, Colorado. Association for Computational Linguistics.
- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2014. [SemEval-2014 task 10: Multilingual semantic textual similarity](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 81–91, Dublin, Ireland. Association for Computational Linguistics.
- Eneko Agirre, Carmen Banea, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2016. [SemEval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 497–511, San Diego, California. Association for Computational Linguistics.
- Eneko Agirre, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. 2012. [SemEval-2012 task 6: A pilot on semantic textual similarity](#). In *\*SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 385–393, Montréal, Canada. Association for Computational Linguistics.
- Eneko Agirre, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and Weiwei Guo. 2013. [\\*SEM 2013 shared task: Semantic textual similarity](#). In *Second Joint Conference on Lexical and Computational Semantics (\*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, pages 32–43, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. 2016. [A latent variable model approach to PMI-based word embeddings](#). *Transactions of the Association for Computational Linguistics*, 4:385–399.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018. [A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 789–798, Melbourne, Australia. Association for Computational Linguistics.
- Marco Basaldella, Fangyu Liu, Ehsan Shareghi, and Nigel Collier. 2020. [COMETA: A corpus for medical entity linking in the social media](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3122–3137, Online. Association for Computational Linguistics.
- Elad Ben-Zaken, Shauli Ravfogel, and Yoav Goldberg. 2020. [BitFit: Simple parameter-efficient fine-tuning for transformer-based masked language-model](#). Technical Report.
- Olivier Bodenreider. 2004. [The unified medical language system \(umls\): integrating biomedical terminology](#). *Nucleic acids research*, 32(suppl\_1):D267–D270.
- Xavier Bouthillier, Kishore Konda, Pascal Vincent, and Roland Memisevic. 2015. [Dropout as data augmentation](#). *ArXiv preprint*, abs/1506.08700.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Xingyu Cai, Jiaji Huang, Yuchen Bian, and Kenneth Church. 2021. [Isotropy in the contextual embedding space: Clusters and manifolds](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Fredrik Carlsson, Amaru Cuba Gyllensten, Evangelia Gogoulou, Erik Ylipää Hellqvist, and Magnus Sahlgren. 2021. [Semantic re-tuning with contrastive tension](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. [Universal sentence encoder for English](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, Brussels, Belgium. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference*

- of the North American Chapter of the Association for Computational Linguistics: *Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Rezarta Islamaj Doğan, Robert Leaman, and Zhiyong Lu. 2014. Ncbi disease corpus: a resource for disease name recognition and concept normalization. *Journal of biomedical informatics*, 47:1–10.
- Maud Ehrmann, Francesco Cecconi, Daniele Vannella, John Philip McCrae, Philipp Cimiano, and Roberto Navigli. 2014. Representing multilingual data as linked data: the case of BabelNet 2.0. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 401–408, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Kawin Ethayarajh. 2019. How contextual are contextualized word representations? comparing the geometry of BERT, ELMo, and GPT-2 embeddings. 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)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics.
- Christiane Fellbaum. 1998. *WordNet*. MIT Press.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. Language-agnostic bert sentence embedding. *ArXiv preprint*, abs/2007.01852.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple contrastive learning of sentence embeddings. *ArXiv preprint*, abs/2104.08821.
- John Giorgi, Osvald Nitski, Bo Wang, and Gary Bader. 2021. DeCLUTR: Deep contrastive learning for unsupervised textual representations. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 879–895, Online. Association for Computational Linguistics.
- Goran Glavaš, Robert Litschko, Sebastian Ruder, and Ivan Vulić. 2019. How to (properly) evaluate cross-lingual word embeddings: On strong baselines, comparative analyses, and some misconceptions. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 710–721, Florence, Italy. Association for Computational Linguistics.
- Goran Glavaš and Ivan Vulić. 2021. Is supervised syntactic parsing beneficial for language understanding tasks? an empirical investigation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3090–3104, Online. Association for Computational Linguistics.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2020. Domain-specific language model pretraining for biomedical natural language processing. *ArXiv preprint*, abs/2007.15779.
- Zellig S Harris. 1954. Distributional structure. *Word*, 10(2-3):146–162.
- Matthew Henderson, Ivan Vulić, Daniela Gerz, Iñigo Casanueva, Paweł Budzianowski, Sam Coope, Georgios Spithourakis, Tsung-Hsien Wen, Nikola Mrkšić, and Pei-Hao Su. 2019. Training neural response selection for task-oriented dialogue systems. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5392–5404, Florence, Italy. Association for Computational Linguistics.
- Dan Hendrycks, Norman Mu, Ekin Dogus Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. 2020. Augmix: A simple data processing method to improve robustness and uncertainty. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Felix Hill, Kyunghyun Cho, and Anna Korhonen. 2016. Learning distributed representations of sentences from unlabelled data. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1367–1377, San Diego, California. Association for Computational Linguistics.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. SimLex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.
- Samuel Humeau, Kurt Shuster, Marie-Anne Lachaux, and Jason Weston. 2020. Poly-encoders: Architectures and pre-training strategies for fast and accurate multi-sentence scoring. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Taeuk Kim, Kang Min Yoo, and Sang-goo Lee. 2021. Self-guided contrastive learning for BERT sentence representations. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2528–2540, Online. Association for Computational Linguistics.
- Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 3294–3302.

- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. [Revealing the dark secrets of BERT](#). 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)*, pages 4365–4374, Hong Kong, China. Association for Computational Linguistics.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. [Word translation without parallel data](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020. [On the sentence embeddings from pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9119–9130, Online. Association for Computational Linguistics.
- Jiao Li, Yueping Sun, Robin J Johnson, Daniela Sciaky, Chih-Hsuan Wei, Robert Leaman, Allan Peter Davis, Carolyn J Mattingly, Thomas C Wieggers, and Zhiyong Lu. 2016. [Biocreative v cdr task corpus: a resource for chemical disease relation extraction](#). *Database*, 2016.
- Nut Limsopatham and Nigel Collier. 2016. [Normalising medical concepts in social media texts by learning semantic representation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1014–1023, Berlin, Germany. Association for Computational Linguistics.
- Robert Litschko, Ivan Vulić, Simone Paolo Ponzetto, and Goran Glavaš. 2021. [Evaluating multilingual text encoders for unsupervised cross-lingual retrieval](#). In *Proceedings of 43rd European Conference on Information Retrieval (ECIR 2021)*, pages 342–358.
- Fangyu Liu, Ehsan Shareghi, Zaiqiao Meng, Marco Basaldella, and Nigel Collier. 2021. [Self-alignment pretraining for biomedical entity representations](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4228–4238, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *ArXiv preprint*, abs/1907.11692.
- Lajanugen Logeswaran and Honglak Lee. 2018. [An efficient framework for learning sentence representations](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. [A SICK cure for the evaluation of compositional distributional semantic models](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 216–223, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Yu Meng, Chenyan Xiong, Payal Bajaj, Saurabh Tiwary, Paul Bennett, Jiawei Han, and Xia Song. 2021. [Coco-lm: Correcting and contrasting text sequences for language model pretraining](#). *ArXiv preprint*, abs/2102.08473.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. [Are sixteen heads really better than one?](#) In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 14014–14024.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. [Efficient estimation of word representations in vector space](#). *ArXiv preprint*, abs/1301.3781.
- Tomas Mikolov, Edouard Grave, Piotr Bojanowski, Christian Puhresch, and Armand Joulin. 2018. [Advances in pre-training distributed word representations](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Tomás Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013b. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3111–3119.
- Jiaqi Mu and Pramod Viswanath. 2018. [All-but-the-top: Simple and effective postprocessing for word representations](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. [Representation learning with contrastive predictive coding](#). *ArXiv preprint*, abs/1807.03748.

- Ellie Pavlick, Pushpendre Rastogi, Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2015. [PPDB 2.0: Better paraphrase ranking, fine-grained entailment relations, word embeddings, and style classification](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 425–430, Beijing, China. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Mohammad Taher Pilehvar, Dimitri Kartsaklis, Victor Prokhorov, and Nigel Collier. 2018. [Card-660: Cambridge rare word dataset - a reliable benchmark for infrequent word representation models](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1391–1401, Brussels, Belgium. Association for Computational Linguistics.
- Sara Rajaei and Mohammad Taher Pilehvar. 2021. [A cluster-based approach for improving isotropy in contextual embedding space](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 575–584, Online. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Nils Rethmeier and Isabelle Augenstein. 2021. [A primer on contrastive pretraining in language processing: Methods, lessons learned and perspectives](#). *ArXiv preprint*, abs/2102.12982.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in BERTology: What we know about how BERT works](#). *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2021. [Wiki-Matrix: Mining 135M parallel sentences in 1620 language pairs from Wikipedia](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1351–1361, Online. Association for Computational Linguistics.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: a simple way to prevent neural networks from overfitting](#). *The journal of machine learning research*, 15(1):1929–1958.
- Jianlin Su. 2020. [Simbert: Integrating retrieval and generation into bert](#). Technical report.
- Jianlin Su, Jiarun Cao, Weijie Liu, and Yangyiwen Ou. 2021. [Whitening sentence representations for better semantics and faster retrieval](#). *ArXiv preprint*, abs/2103.15316.
- Nandan Thakur, Nils Reimers, Johannes Daxenberger, and Iryna Gurevych. 2021. [Augmented SBERT: Data augmentation method for improving bi-encoders for pairwise sentence scoring tasks](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 296–310, Online. Association for Computational Linguistics.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Ivan Vulić, Simon Baker, Edoardo Maria Ponti, Ulla Petti, Ira Leviant, Kelly Wing, Olga Majewska, Eden Bar, Matt Malone, Thierry Poibeau, Roi Reichart, and Anna Korhonen. 2020a. [Multi-SimLex: A large-scale evaluation of multilingual and crosslingual lexical semantic similarity](#). *Computational Linguistics*, 46(4):847–897.
- Ivan Vulić, Edoardo Maria Ponti, Anna Korhonen, and Goran Glavaš. 2021. [LexFit: Lexical fine-tuning of pretrained language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5269–5283, Online. Association for Computational Linguistics.
- Ivan Vulić, Edoardo Maria Ponti, Robert Litschko, Goran Glavaš, and Anna Korhonen. 2020b. [Probing pretrained language models for lexical semantics](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7222–7240, Online. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer

- Levy, and Samuel R. Bowman. 2019a. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 3261–3275.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019b. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Tongzhou Wang and Phillip Isola. 2020. [Understanding contrastive representation learning through alignment and uniformity on the hypersphere](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 9929–9939. PMLR.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Zhuofeng Wu, Sinong Wang, Jiatao Gu, Madian Khabsa, Fei Sun, and Hao Ma. 2020. [Clear: Contrastive learning for sentence representation](#). *ArXiv preprint*, abs/2012.15466.
- Yuanmeng Yan, Rumei Li, Sirui Wang, Fuzheng Zhang, Wei Wu, and Weiran Xu. 2021. [ConSERT: A contrastive framework for self-supervised sentence representation transfer](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5065–5075, Online. Association for Computational Linguistics.
- Yan Zhang, Ruidan He, Zuozhu Liu, Lidong Bing, and Haizhou Li. 2021. [Bootstrapped unsupervised sentence representation learning](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5168–5180, Online. Association for Computational Linguistics.
- Yan Zhang, Ruidan He, Zuozhu Liu, Kwan Hui Lim, and Lidong Bing. 2020. [An unsupervised sentence embedding method by mutual information maximization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1601–1610, Online. Association for Computational Linguistics.
- Wangchunshu Zhou, Tao Ge, Furu Wei, Ming Zhou, and Ke Xu. 2020. [Scheduled DropHead: A regularization method for transformer models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1971–1980, Online. Association for Computational Linguistics.

## A Language Codes

EN	English
ES	Spanish
FR	French
PL	Polish
ET	Estonian
FI	Finnish
RU	Russian
TR	Turkish
IT	Italian
ZH	Chinese
AR	Arabic
HE	Hebrew

Table 11: Language abbreviations used in the paper.

## B Additional Training Details

### Most Frequent 10k/100k Words by Language.

The most frequent 10k words in each language were selected based on the following list:

<https://github.com/oprogramador/most-common-words-by-language>.

The most frequent 100k English words in Wikipedia can be found here:

<https://gist.github.com/h3xx/1976236>.

**[CLS] or Mean-Pooling?** For MLMs, the consensus in the community, also validated by our own experiments, is that mean-pooling performs better than using [CLS] as the final output representation. However, for Mirror-BERT models, we found [CLS] (before pooling) generally performs better than mean-pooling. The exception is BERT on sentence-level tasks, where we found mean-pooling performs better than [CLS]. In sum, sentence-level *BERT+Mirror* models are fine-tuned and tested with mean-pooling while all other Mirror-BERT models are fine-tuned and tested with [CLS]. We also tried representations after the pooling layer, but found no improvement.

**Training Stability.** All task results are reported as averages over three runs with different random seeds (if applicable). In general, fine-tuning is very stable and the fluctuations with different random seeds are very small. For instance, on the sentence-level task STS, the standard deviation is  $< 0.002$ . On word-level, standard deviation is a bit higher, but is generally  $< 0.005$ . Note that the randomly

dropout rate→	0.05	0.1*	0.2	0.3	0.4
BERT + Mirror	.740	.743	<b>.748</b>	<b>.748</b>	.731
RoBERTa + Mirror	<b>.755</b>	.753	.737	.694	.677

Table 12: Average  $\rho$  across STS tasks with different dropout rates. \* default dropout rate for all models in other experiments.

random span mask rate→	2	5*	10	15	20
BERT + Mirror	.741	<b>.743</b>	.720	.690	.616
RoBERTa + Mirror	.750	.753	<b>.757</b>	.743	.706

Table 13: Avg.  $\rho$  across STS tasks with different random span masking rates. \* default mask rates for all models in other experiments.

sampled training sets are fixed across all experiments, and changing the training corpus for each run might lead to larger fluctuations.

## C Details of *Mirror-BERT* Trained on Random Strings

We pointed out in the main text that *BERT+Mirror* trained on random strings can outperform MLMs by large margins. With standard training configurations, BERT improves from .267 (BERT-mp) to .393 with *+Mirror*. When learning rate is increased to  $5e-5$ , the MLM fine-tuned with random strings performs only around 0.07 lower than the standard *BERT+Mirror* model fine-tuned with the 10k most frequent English words.

## D Dropout and Random Span Masking Rates

**Dropout Rate (Tab. 12).** The performance trends conditioned on dropout rates are generally the same across word-level, phrase-level and sentence-level fine-tuning. Here, we use the STS task as a reference point. BERT prefers larger dropouts (0.2 & 0.3) and is generally more robust. RoBERTa prefers a smaller dropout rate (0.05) and its performance decreases more steeply with the increase of the dropout rate. For simplicity, as mentioned in the main paper, we use the default value of 0.1 as the dropout rate for all models.

**Random Span Masking Rate (Tab. 13).** Interestingly, the opposite holds for random span masking: RoBERTa is more robust to larger masking rates  $k$ , and is much more robust than BERT to this hyperparameter.

level→ model↓	word		phrase		sentence	
	MVN	IS	MVN	IS	MVN	IS
BERT-CLS	13.79	.043	12.8	.028	12.73	.062
BERT-mp	7.89	.169	6.82	.205	6.93	.222
+ Mirror	2.11	.599	<b>5.91</b>	<b>.252</b>	<b>5.57</b>	<b>.265</b>
+ Mirror (w/o aug.)	<b>0.71</b>	<b>.825</b>	8.16	.170	5.75	.255

Table 14: Full table for MVN and IS of word-, phrase-, and sentence-level models. Higher is better, that is, more isotropic with IS, while the opposite holds for MVN (lower scores mean more isotropic representation spaces).

## E Mean-Vector $l_2$ -Norm (MVN)

To supplement the quantitative evidence already suggested by the Isotropy Score (IS) in the main paper, we additionally compute the mean-vector  $l_2$ -norm (MVN) of embeddings. In the word embedding literature, mean-centering has been a widely studied post-processing technique for inducing better semantic representations. [Mu and Viswanath \(2018\)](#) point out that mean-centering is essentially increasing spatial isotropy by shifting the centre of the space to the region where actual data points reside in. Given a set of representation vectors  $\mathcal{V}$ , we define MVN as follows:

$$\text{MVN}(\mathcal{V}) = \left\| \sum_{\mathbf{v} \in \mathcal{V}} \frac{\mathbf{v}}{|\mathcal{V}|} \right\|_2. \quad (3)$$

The lower MVN is, the more mean-centered an embedding is. As shown in [Tab. 14](#), MVN aligns with the trends observed with IS. This further confirms our intuition that *+Mirror* tuning makes the space more isotropic and shifts the centre of space close to the centre of data points.

Very recently, [Cai et al. \(2021\)](#) defined more metrics to measure spatial isotropy. [Rajae and Pilehvar \(2021\)](#) also used [Eq. \(2\)](#) for analysing sentence embedding’s isotropiness.

## F Evaluation Dataset Details

All datasets used and links to download them can be found in the code repository provided. The Russian STS dataset is provided by <https://github.com/deepmipt/deepPavlovEval>. The Quora Question Pair (QQP) dataset is downloaded at <https://www.kaggle.com/c/quora-question-pairs>.

## G Pretrained Encoders

A complete listing of URLs for all used pretrained encoders is provided in [Tab. 15](#). For monolingual MLMs of each language, we made the best effort to select the most popular one (based on download counts). For computational tractability of the large number of experiments conducted, all models are BASE models (instead of LARGE).

## H Full Tables

Here, we provide the complete sets of results. In these tables we include both MLMs w/ features extracted using both mean-pooling (“mp”) and [CLS] (“CLS”).

For full multilingual word similarity results, view [Tab. 16](#). For full Spanish, Arabic and Russian STS results, view [Tab. 3](#). For full cross-lingual word similarity results, view [Tab. 18](#). For full BLI results, view [Tab. 19](#). For full ablation study results, view [Tab. 20](#). For full MVN and IS scores, view [Tab. 14](#).

## I Number of Model Parameters

All BERT/RobERTa models in this paper have  $\approx 110\text{M}$  parameters.

## J Hyperparameter Optimisation

[Tab. 21](#) lists the hyperparameter search space. Note that the chosen hyperparameters yield the overall best performance, but might be suboptimal on any single setting (e.g. different base model).

## K Software and Hardware Dependencies

All our experiments are implemented using PyTorch 1.7.0 and [huggingface.co](https://huggingface.co) transformers 4.4.2, with Automatic Mixed Precision (AMP)<sup>17</sup> turned on during training. Please refer to the GitHub repo for details. The hardware we use is listed in [Tab. 22](#).

<sup>17</sup><https://pytorch.org/docs/stable/amp.html>



model	URL
fastText	<a href="https://fasttext.cc/docs/en/crawl-vectors.html">https://fasttext.cc/docs/en/crawl-vectors.html</a>
SBERT	<a href="https://huggingface.co/sentence-transformers/bert-base-nli-mean-tokens">https://huggingface.co/sentence-transformers/bert-base-nli-mean-tokens</a>
SapBERT	<a href="https://huggingface.co/cambridgeltl/SapBERT-from-PubMedBERT-fulltext">https://huggingface.co/cambridgeltl/SapBERT-from-PubMedBERT-fulltext</a>
BERT (English)	<a href="https://huggingface.co/bert-base-uncased">https://huggingface.co/bert-base-uncased</a>
RoBERTa (English)	<a href="https://huggingface.co/roberta-base">https://huggingface.co/roberta-base</a>
mBERT	<a href="https://huggingface.co/bert-base-multilingual-uncased">https://huggingface.co/bert-base-multilingual-uncased</a>
Turkish BERT	<a href="https://huggingface.co/dbmdz/bert-base-turkish-uncased">dbmdz/bert-base-turkish-uncased</a>
Italian BERT	<a href="https://huggingface.co/dbmdz/bert-base-italian-uncased">dbmdz/bert-base-italian-uncased</a>
French BERT	<a href="https://huggingface.co/camembert-base">https://huggingface.co/camembert-base</a>
Spanish BERT	<a href="https://huggingface.co/dccuchile/bert-base-spanish-wwm-uncased">https://huggingface.co/dccuchile/bert-base-spanish-wwm-uncased</a>
Russian BERT	<a href="https://huggingface.co/DeepPavlov/rubert-base-cased">https://huggingface.co/DeepPavlov/rubert-base-cased</a>
Chinese BERT	<a href="https://huggingface.co/bert-base-chinese">https://huggingface.co/bert-base-chinese</a>
Arabic BERT	<a href="https://huggingface.co/aubmindlab/bert-base-arabertv02">https://huggingface.co/aubmindlab/bert-base-arabertv02</a>
Polish BERT	<a href="https://huggingface.co/dkleczyk/bert-base-polish-uncased-v1">https://huggingface.co/dkleczyk/bert-base-polish-uncased-v1</a>
Estonian BERT	<a href="https://huggingface.co/tartuNLP/EstBERT">https://huggingface.co/tartuNLP/EstBERT</a>

Table 15: A listing of HuggingFace & fastText URLs of all pretrained models used in this work.

language→	EN	FR	ET	AR	ZH	RU	ES	PL	avg.
fastText	<u>.528</u>	<u>.560</u>	<b>.447</b>	<u>.409</u>	.428	<b>.435</b>	<b>.488</b>	<u>.396</u>	<u>.461</u>
BERT-CLS	.105	.050	.160	.210	.277	.177	.152	.257	.174
BERT-mp	.267	.020	.106	.220	.398	.202	.177	.217	.201
+ Mirror	<b>.556</b>	<b>.621</b>	.308	<b>.538</b>	<b>.639</b>	<u>.365</u>	.296	<b>.444</b>	<b>.471</b>
mBERT-CLS	.062	.046	.074	.047	.204	.063	.039	.051	.073
mBERT-mp	.105	.130	.094	.101	.261	.109	.095	.087	.123
+ Mirror	.358	.341	.134	.097	<u>.501</u>	.210	<u>.332</u>	.141	.264

Table 16: Word similarity evaluation on Multi-SimLex (Spearman’s  $\rho$ ).

model↓, lang.→	ES	AR	RU	avg.
BERT-CLS	.526	.308	.470	.435
BERT-mp	.599	.455	.552	.535
+ Mirror	<u>.709</u>	<b>.669</b>	<u>.673</u>	<b>.684</b>
mBERT-CLS	.421	.326	.430	.392
mBERT-mp	.610	.447	.616	.558
+ Mirror	<b>.755</b>	<u>.594</u>	<b>.692</b>	<u>.680</u>

Table 17: Full Spanish, Arabic and Russian STS evaluation. Spearman’s  $\rho$  correlation reported.

lang.→	EN-FR	EN-ZH	EN-HE	FR-ZH	FR-HE	ZH-HE	avg.
mBERT-CLS	.059	.053	.032	.042	.024	.050	.043
mBERT-mp	.163	.118	.071	.142	.104	.010	.101
+ Mirror	<b>.454</b>	<b>.385</b>	<b>.133</b>	<b>.465</b>	<b>.163</b>	<b>.179</b>	<b>.297</b>

Table 18: Full cross-lingual word similarity evaluation on Multi-SimLex (Spearman’s  $\rho$ ).

lang.→	EN-FR	EN-IT	EN-RU	EN-TR	IT-FR	RU-FR	avg.
BERT-CLS	.045	.049	.108	.109	.046	.068	.071
BERT-mp	.014	.112	.154	.150	.025	.018	.079
+ Mirror	<b>.458</b>	<b>.378</b>	<b>.336</b>	<b>.289</b>	<b>.417</b>	<b>.345</b>	<b>.371</b>

Table 19: Full Bilingual Lexicon Induction results (accuracy reported). “EN-FR” means en mapped to FR.

model configuration↓, dataset→	STS12	STS13	STS14	STS15	STS16	STS-b	SICK-R	avg.
BERT + Mirror	.674	.796	.713	.814	.743	.764	.703	.744
- dropout	.646	.770	.691	.800	.726	.745	.701	.726↓.018
- random span masking	.641	.775	.684	.777	.737	.749	.658	.717↓.027
- dropout & random span masking	.587	.695	.617	.688	.683	.674	.614	.651↓.093
RoBERTa + Mirror	.648	.819	.732	.798	.780	.787	.706	.753
- dropout	.619	.795	.706	.802	.777	.727	.698	.732↓.021
- random span masking	.616	.786	.689	.766	.743	.756	.663	.717↓.036
- dropout & random span masking	.562	.730	.643	.744	.752	.708	.638	.682↓.071

Table 20: Full table for the synergistic effect of dropout and random span masking in sentence similarity tasks.

hyperparameters	search space
learning rate	{ $5e-5$ , $2e-5^*$ , $1e-5$ }
batch size	{100, 200*, 300}
training epochs	{1*, 2*, 3, 5}
$\tau$ in Eq. (1)	{0.03, 0.04*, 0.05, 0.07, 0.1, 0.2*, 0.3}

Table 21: Hyperparameters along with their search grid. \* marks the values used to obtain the reported results. The hparams are not always optimal in every setting but generally performs (close to) the best.

hardware	specification
RAM	128 GB
CPU	AMD Ryzen 9 3900x 12-core processor $\times$ 24
GPU	NVIDIA GeForce RTX 2080 Ti (11 GB) $\times$ 2

Table 22: Hardware specifications of the used machine. When encountering out-of-memory error, we also used a second server with two NVIDIA GeForce RTX 3090 (24 GB).