

Generating Datasets with Pretrained Language Models

Timo Schick and Hinrich Schütze
Center for Information and Language Processing
LMU Munich, Germany
schickt@cis.lmu.de

Abstract

To obtain high-quality sentence embeddings from pretrained language models (PLMs), they must either be augmented with additional pre-training objectives or finetuned on a large set of labeled text pairs. While the latter approach typically outperforms the former, it requires great human effort to generate suitable datasets of sufficient size. In this paper, we show how PLMs can be leveraged to obtain high-quality sentence embeddings without the need for labeled data, finetuning or modifications to the pretraining objective: We utilize the generative abilities of large and high-performing PLMs to generate entire datasets of labeled text pairs from scratch, which we then use for finetuning much smaller and more efficient models. Our fully unsupervised approach outperforms strong baselines on several semantic textual similarity datasets.¹

1 Introduction

While pretrained language models (PLMs) achieve strong results for many NLP tasks (Peters et al., 2018; Radford et al., 2018; Devlin et al., 2019), they do not produce good sentence embeddings out of the box (Reimers and Gurevych, 2019). Recent approaches address this by augmenting or replacing the language modeling objective with likewise unsupervised sentence-level objectives (e.g., Zhang et al., 2020; Li et al., 2020), but they typically lag behind their supervised counterparts trained on human-annotated sentence pairs. Unfortunately, obtaining large amounts of high-quality training data can be both difficult and prohibitively expensive (Bowman et al., 2015; Agirre et al., 2016). Furthermore, with larger and larger model sizes (Radford et al., 2019; Raffel et al., 2020; Brown et al., 2020; Fedus et al., 2021), it becomes increasingly challenging to finetune PLMs.

¹Our code and datasets are publicly available at <https://github.com/timoschick/dino>.

Task: Write two sentences that mean the same thing.

Sentence 1: “A man is playing a flute.”

Sentence 2: “He’s playing a flute.”

Task: Write two sentences that are somewhat similar.

Sentence 1: “A man is playing a flute.”

Sentence 2: “A woman has been playing the violin.”

Task: Write two sentences that are on completely different topics.

Sentence 1: “A man is playing a flute.”



Sentence 2: “A woman is walking down the street.”

Figure 1: **Continuations** generated by GPT2-XL with DINO for three different **task descriptions**. We investigate two different unsupervised approaches to generating sentence-similarity datasets: (i) The **input sentence** is given and only the **continuation** is generated. This requires that an (unlabeled) set of sentences is available. (ii) Both **input sentence** and **continuation** are generated. This does not rely on the availability of any resources.

To alleviate both problems, we explore a novel approach to obtaining high-quality sentence embeddings: We mimic the creation of NLI datasets by human crowdworkers (Bowman et al., 2015; Williams et al., 2018), but replace human annotators with large PLMs. This allows us to automatically create entire datasets from scratch that can be used for supervised training of much smaller models. Not only does this solve the problem of limited training data, it also provides a viable path to leverage big models like GPT-3 (Brown et al., 2020) without requiring any updates to their parameters. As illustrated in Figure 1, our approach is based on recent methods for providing instructions to PLMs (e.g., Radford et al., 2019; Brown et al., 2020; Schick and Schütze, 2020, 2021a). We use the *self-debiasing* approach of Schick et al. (2021) to ensure that each generated text pair is not only a

good fit for a given similarity label, but also *not* a good fit for other labels. We refer to our method as **Datasets from Instructions (DINO)**.

In summary, our contributions are as follows:

- We introduce DINO, a method for automatically generating labeled datasets of arbitrary size by providing PLMs with instructions.
- We release STS- (read as “STS-Dino”), the first textual similarity dataset generated completely automatically, without any human annotation effort.
- We show that Sentence-RoBERTa (Reimers and Gurevych, 2019) trained on STS- outperforms strong baselines on several semantic textual similarity datasets.

2 Related Work

There are many unsupervised approaches to obtaining sentence embeddings, for example by averaging word embeddings (Mikolov et al., 2013; Pennington et al., 2014; Bojanowski et al., 2017) or with carefully designed sentence-level objectives (Le and Mikolov, 2014; Kiros et al., 2015). Ensembling several methods improves results (Pörner and Schütze, 2019; Pörner et al., 2020). Recent work obtains sentence representations by supplementing BERT (Devlin et al., 2019) or other PLMs with additional unsupervised objectives (Zhang et al., 2020; Li et al., 2020; Wu et al., 2020; Giorgi et al., 2020). Often, labeled datasets such as paraphrase databases (Wieting and Gimpel, 2018) or natural language inference datasets (Conneau et al., 2017; Cer et al., 2018; Reimers and Gurevych, 2019) are used for supervised learning.

Some approaches *augment* existing datasets with automatically generated examples (Anaby-Tavor et al., 2020; Papanikolaou and Pierleoni, 2020; Yang et al., 2020; Mohapatra et al., 2020; Kumar et al., 2021), but in contrast to our work, all of these approaches require that there already exists a labeled dataset for finetuning the generator. Providing PLMs with task descriptions for zero- or few-shot learning has been studied extensively (e.g., Radford et al., 2019; Puri and Catanzaro, 2019; Brown et al., 2020; Schick and Schütze, 2020, 2021b,a; Weller et al., 2020; Gao et al., 2021; Tam et al., 2021). However, none of these approaches is suitable for generating sentence embeddings.

Closely related to our work, Efrat and Levy (2020) examine the ability of PLMs to follow natu-

Task: Write two sentences that i_y .

Sentence 1: “ \mathbf{x}_1 ”

Sentence 2: “

Figure 2: Instruction template $I_y(\mathbf{x}_1)$ for similarity label y and input sentence \mathbf{x}_1 ; i_y is described in Section 3. See Figure 1 for three instantiations of the template.

ral language instructions for generating examples in place of human crowdworkers, but find that their approach performs poorly.

3 Datasets from Instructions

Let M be a PLM with vocabulary V , $X = V^*$ the set of all token sequences and Y a finite set of semantic similarity labels. Our aim is to generate a dataset $Z \subset X \times X \times Y$ of text *pairs* $(\mathbf{x}_1, \mathbf{x}_2)$ with corresponding similarity labels y . For $x \in V$ and $\mathbf{x} \in X$, we denote with $p_M(x | \mathbf{x})$ the probability that M assigns to x as a continuation of \mathbf{x} .

We first assume that **we already have access to a set $X_1 \subset X$ of texts** (e.g., a set of sentences that are typical of the domain of interest). This is a realistic setting for many real-world applications, where large amounts of unlabeled text are abundant, but it is difficult to obtain *interesting* and (for our task) *useful* text pairs and labels. DINO requires a set of *instructions* $\mathcal{I} = \{I_y | y \in Y\}$ where each $I_y \in \mathcal{I}$ is a function that, given an input $\mathbf{x}_1 \in X_1$, prompts its recipient to generate an appropriate second text \mathbf{x}_2 . We use the instruction template in Figure 2 and consider three levels of similarity ($Y = \{0, 0.5, 1\}$), where

$$i_y = \begin{cases} \text{mean the same thing} & \text{if } y = 1 \\ \text{are somewhat similar} & \text{if } y = 0.5 \\ \text{are on completely different topics} & \text{if } y = 0 \end{cases}$$

is loosely based on Cer et al. (2017)’s five-level similarity scheme. Note that for all y , I_y ends with an opening quotation mark, which allows us to treat the first quotation mark generated by the PLM as a sign that it is done.

For a given $\mathbf{x}_1 \in X_1$ and $y \in Y$, we could directly use the instructions I_y to obtain \mathbf{x}_2 by continuously sampling tokens

$$x_k \sim p_M(x_k | I_y(\mathbf{x}_1), x_1, \dots, x_{k-1})$$

starting from $k = 1$ until x_k is a quotation mark and setting $\mathbf{x}_2 = x_1, \dots, x_{k-1}$. However, we may

want the PLM to generate a text \mathbf{x}_2 that is not only a good fit for instruction $I_y(\mathbf{x}_1)$, but also *not* a good fit for some other instruction $I_{y'}(\mathbf{x}_1)$. We refer to y' as a *counterlabel* for y and denote the set of y' 's counterlabels as $\text{CL}(y)$. For example, $1 \in \text{CL}(0.5)$ means that for $y = 0.5$, we want M to generate a sentence \mathbf{x}_2 that is similar to ($y = 0.5$), but at the same time does *not* have the same meaning as ($y = 1$) sentence \mathbf{x}_1 . We achieve this using Schick et al. (2021)'s self-debiasing algorithm: When sampling the token x_k , we consider not just $p_y = p_M(x_k | I_y(\mathbf{x}_1), x_1, \dots, x_{k-1})$ [x_k 's probability given $I_y(\mathbf{x}_1)$], but also $p_{y'}$ [x_k 's probability given $I_{y'}(\mathbf{x}_1)$], for all $y' \in \text{CL}(y)$. We penalize each token x_k for which p_y is lower than *any* $p_{y'}$ by multiplying its probability with a factor $\alpha = \exp(\lambda \cdot \delta_y)$ where

$$\delta_y = p_y - \max_{y' \in \text{CL}(y)} p_{y'}$$

is the difference between x_k 's probability given $I_y(\mathbf{x}_1)$ and its maximum probability given $I_{y'}(\mathbf{x}_1)$ for any $y' \in \text{CL}(y)$, and the *decay constant* λ is a hyperparameter.

For settings where **no set of unlabeled texts** X_1 is available, a straightforward approach would be to use the phrase shown in Figure 2 up to and including the first quotation mark as an instruction to let the PLM generate both \mathbf{x}_1 and \mathbf{x}_2 . However, this approach has at least two issues: First, generated texts may not match the required schema (e.g., the model may never produce the string ‘‘Sentence 2:’’). Second, the set of texts \mathbf{x}_1 should ideally be highly diverse, whereas we want to give the model less leeway when generating \mathbf{x}_2 , so we may want to use different sampling strategies for \mathbf{x}_1 and \mathbf{x}_2 .

We solve both problems as follows: We first use I_y (Figure 2) up to and including the first quotation mark (the one right after ‘‘Sentence 1:’’ to generate \mathbf{x}_1 ; we stop as soon as the model produces a quotation mark. We run this procedure repeatedly until we have a sufficient number of sentences. These are gathered into a set X_1 and then we proceed exactly as in the case where X_1 is already given.

4 Experiments

We evaluate DINO on several English semantic textual similarity datasets: the STS tasks 2012–2016 (Agirre et al., 2012, 2013, 2014, 2015, 2016), the STS benchmark (STSb) (Cer et al., 2017), and the SICK-Relatedness dataset (SICK) (Marelli et al.,

2014). For all tasks, we adopt the unsupervised setting without task-specific training examples.

We use DINO to generate $\text{STS-}\mathcal{L} \subset X \times X \times Y$, a dataset of text pairs with semantic similarity labels. We generate two variants:

- $\text{STS-}\mathcal{L}\text{-}\mathbf{x}_2$, for which we make use of STSb to obtain a set of texts X_1 ;
- $\text{STS-}\mathcal{L}\text{-}\mathbf{x}_1\mathbf{x}_2$, where the set of sentences X_1 is generated from scratch.

We use GPT2-XL as PLM with a decay constant of $\lambda = 100$ and the set of counterlabels $\text{CL}(y) = \{y' \in Y \mid y' > y\}$. That is, we do not restrict the PLM when generating texts for $y = 1$, but for $y = 0.5$ ($y = 0$) we encourage it not to generate texts \mathbf{x}_2 that mean the same thing as (are somewhat similar to) \mathbf{x}_1 . We apply top- p (Holtzman et al., 2020) and top- k (Fan et al., 2018; Holtzman et al., 2018) sampling with $p = 0.9$, $k = 5$ and generate up to 40 output tokens. For each $\mathbf{x}_1 \in X_1$ and $y \in Y$, we generate up to two corresponding \mathbf{x}_2 's.² For $\text{STS-}\mathcal{L}\text{-}\mathbf{x}_1\mathbf{x}_2$, we obtain X_1 by generating 15,000 sentences using only top- p sampling (again with $p = 0.9$) and no top- k sampling to ensure more diversity in the generated output. We remove all examples where $\mathbf{x}_1 = \mathbf{x}_2$ (as those provide no training signal to the model) and split the datasets 90/10 into training and validation.

To assess the quality of the generated datasets, we use them to train Sentence-RoBERTa (Reimers and Gurevych, 2019), a biencoder architecture based on RoBERTa (base) (Liu et al., 2019) that measures the similarity of two texts by computing the cosine similarity of their embeddings. As our datasets contain many noisy examples, we use a technique similar to label smoothing (Szegedy et al., 2016) and replace similarity scores of 0 and 1 with 0.1 and 0.9, respectively. Additionally, for each \mathbf{x}_1 , we sample two \mathbf{x}_2 's from *other* dataset entries and augment the dataset with $(\mathbf{x}_1, \mathbf{x}_2, 0)$. We use the default parameters of Reimers and Gurevych (2019) with a batch size of 32 and train for at most one epoch; the exact number of training steps is determined based on Spearman's rank correlation on the $\text{STS-}\mathcal{L}$ validation set.

Results We compare S-RoBERTa (base) trained on datasets generated with DINO to S-BERT and S-RoBERTa finetuned on NLI data as well as Universal Sentence Encoder (USE) (Cer et al., 2018)

²As the PLM may not generate a quotation mark in the first 40 tokens, we use up to 5 tries to generate the two \mathbf{x}_2 's.

	Model	UD	STS12	STS13	STS14	STS15	STS16	STSb	SICK	Avg.
sup.	InferSent, GloVe	–	52.86	66.75	62.15	72.77	66.87	68.03	65.65	65.01
	USE	–	64.49	67.80	64.61	76.83	73.18	74.92	76.69	71.22
	S-BERT (base)	–	70.97	76.53	73.19	79.09	74.30	77.03	72.91	74.89
	S-RoBERTa (base)	–	<u>71.54</u>	72.49	<u>70.80</u>	78.74	73.69	77.77	<u>74.46</u>	74.21
unsup.	Avg. GloVe	–	55.14	70.66	59.73	68.25	63.66	58.02	53.76	61.32
	Avg. BERT	–	38.78	57.98	57.98	63.15	61.06	46.35	58.40	54.81
	BERT CLS	–	20.16	30.01	20.09	36.88	38.08	16.50	42.63	29.19
	Zhang et al. (2020)	NLI	56.77	69.24	61.21	75.23	70.16	69.21	64.25	66.58
	Li et al. (2020)	NLI	59.54	64.69	64.66	72.92	71.84	58.56	65.44	65.38
	Li et al. (2020)	STS	63.48	72.14	68.42	73.77	75.37	70.72	63.11	69.57
	DINO (STS- \mathcal{L} - $\mathbf{x}_1\mathbf{x}_2$)	–	64.87	78.30	66.38	79.60	76.47	76.51	74.26	73.77
	DINO (STS- \mathcal{L} - \mathbf{x}_2)	STS	70.27	81.26	71.25	80.49	77.18	77.82	68.09	75.20

Table 1: Spearman’s rank correlation on STS12–16, STSb and SICK without finetuning on task-specific examples for models with NLI supervision (“sup.”) and fully unsupervised (“unsup.”) models using the same evaluation setup as Reimers and Gurevych (2019). The second column shows which unlabeled data (“UD”) is used by unsupervised approaches in addition to original pretraining data; the final column shows average performance. Results for all baselines except Zhang et al. (2020) and Li et al. (2020) are from Reimers and Gurevych (2019). The best unsupervised result is shown in bold, the best overall result is underlined. DINO outperforms all unsupervised approaches and, surprisingly, also supervised approaches on four out of six STS datasets.

Model	STS12-16	STSb	SICK
DINO (STS- \mathcal{L} - \mathbf{x}_2)	76.09	77.82	68.09
\downarrow decay constant $\lambda = 0$	65.50	70.71	67.60
\downarrow decay constant $\lambda = 200$	75.40	77.49	66.83
\downarrow no label smoothing	74.50	76.26	66.23
\downarrow no augmentation	70.90	73.81	63.98

Table 2: Effect of removing self-debiasing ($\lambda = 0$) or increasing the decay constant ($\lambda = 200$), using no label smoothing and performing no data augmentation (sampling random \mathbf{x}_2 ’s for each \mathbf{x}_1) on the performance of DINO on STS12-16 (avg), STSb and SICK

and InferSent (Conneau et al., 2017), all of which are trained on hundreds of thousands of labeled text pairs from SNLI (Bowman et al., 2015) and MNLI (Williams et al., 2018). We additionally compare to the following fully unsupervised approaches: averaging word-level GloVe (Pennington et al., 2014) or BERT (Devlin et al., 2019) embeddings, using BERT’s CLS token, and recent methods by Zhang et al. (2020) and Li et al. (2020) based on pretrained BERT models. We do not compare to approaches trained with direct supervision as our focus is on obtaining sentence representations without task-specific labeled examples. As shown in Table 1, training on datasets generated with DINO clearly outperforms the fully unsupervised baselines; on average, training on STS- \mathcal{L} - \mathbf{x}_2 even outperforms all approaches with NLI supervision. STS- \mathcal{L} - \mathbf{x}_2 gives better results than STS- \mathcal{L} - $\mathbf{x}_1\mathbf{x}_2$ on all STS datasets as its examples are – by design – very similar to examples found in these datasets, while the latter gives better results on SICK.

We investigate the importance of self-debiasing (Schick et al., 2021) in Table 2 (top); as can be seen, removing self-debiasing ($\lambda = 0$) dramatically hurts performance. Increasing the decay constant ($\lambda = 200$) leads to slightly worse performance as the overall quality of generated sentences decreases (Schick et al., 2021). Table 2 (bottom) shows that training on STS- \mathcal{L} requires measures to limit the effect of noisy labels: removing label smoothing and performing no data augmentation (i.e., not generating additional pairs $(\mathbf{x}_1, \mathbf{x}_2, 0)$ by sampling random \mathbf{x}_2 ’s for each \mathbf{x}_1) clearly hurts performance.

To further assess the quality of datasets generated with DINO, we additionally perform a small-scale human evaluation. To this end, we consider the exact version of STS- \mathcal{L} - \mathbf{x}_2 used for training S-RoBERTa; that is, we perform label smoothing, augmentation with randomly sampled text pairs, and removal of trivial examples where $\mathbf{x}_1 = \mathbf{x}_2$. From the resulting dataset, we randomly select 100 text pairs $(\mathbf{x}_1, \mathbf{x}_2)$ and annotate them ourselves with similarity scores $y \in \{0, 0.1, 0.5, 0.9\}$, where we assign a score of 0.9 when \mathbf{x}_1 and \mathbf{x}_2 mean (almost) the same thing and a score of 0.1 when they are on different topics, but still show a weak similarity in some aspect.

In Table 3, human annotations are compared to originally assigned scores, yielding some interesting insights. For one, it becomes clear why augmentation with randomly sampled text pairs is important for good downstream task performance: Of the examples generated by DINO that are sup-

DINO Labels →	0.0	0.1	0.5	0.9	
Human Labels	0.0	95%	15%	0%	0%
0.1	0%	44%	11%	12%	
0.5	5%	41%	60%	41%	
0.9	0%	0%	29%	47%	

Table 3: Comparison of similarity scores in STS- \mathbf{x}_1 - \mathbf{x}_2 to human judgments for 100 examples. Examples are chosen randomly from the version of STS- \mathbf{x}_1 - \mathbf{x}_2 used for training (including label smoothing, augmentation with random pairs and removal of examples where $\mathbf{x}_1 = \mathbf{x}_2$). For column i and row j , the value shown is the percentage of examples generated by DINO for similarity score i that were assigned score j in our human evaluation.

posed to be on completely different topics, many (41%) still have a certain similarity according to human judgment. In contrast, randomly sampled pairs are indeed on completely different topics in almost all cases. Moreover, we can see that GPT2-XL has particular difficulty in generating pairs of non-identical sentences that really mean the same thing: Only 47% of all examples that should have the same meaning do actually mean (almost) the same thing. However, the strong performance of S-RoBERTa trained on STS- \mathbf{x}_1 - \mathbf{x}_2 suggests that, despite this noise, there is sufficient signal in this dataset for successful training.

We finally take a qualitative look at both positive examples where DINO is able to create high-quality text pairs and at some typical errors found in many of the generated examples. As shown in Table 4, for $y = 1$ the PLM sometimes comes up with decent paraphrases (e.g. “notches a victory” \mapsto “wins”) or substitutes with very similar meaning (“cutting” \mapsto “slicing”), but more often it generates sentences that either omit or mix up important information, and sometimes it produces sentences with an entirely different meaning. Whereas sentences generated for $y = 0.5$ by and large look reasonable, for $y = 0$ the PLM often simply flips words (“closed” \mapsto “open”, “large” \mapsto “small”) instead of producing sentences on completely different topics.

5 Conclusion

We have introduced DINO, a method for using large PLMs to generate entire datasets of labeled sentence pairs from scratch, requiring no labeled data and no parameter updates. This is achieved by providing instructions in natural language, combined with the self-debiasing method of Schick

	$\mathbf{x}_1 =$ Rick Santorum notches a victory in Kansas caucuses. $\mathbf{x}_2 =$ Rick Santorum wins Kansas caucuses.	✓
	$\mathbf{x}_1 =$ A man is cutting cucumbers. $\mathbf{x}_2 =$ A man is slicing cucumbers.	✓
$y = 1$	$\mathbf{x}_1 =$ US closes embassy in Syria $\mathbf{x}_2 =$ US Embassy in Syria	✗
	$\mathbf{x}_1 =$ A man is playing the cello. $\mathbf{x}_2 =$ The cello is playing the man.	✗
	$\mathbf{x}_1 =$ A plane is taking off. $\mathbf{x}_2 =$ I want to be a pilot.	✗
$y = 0.5$	$\mathbf{x}_1 =$ A woman is seasoning a piece of meat. $\mathbf{x}_2 =$ A man is cooking the meat and adding spices [...]	✓
	$\mathbf{x}_1 =$ Second day of Egyptian presidential election $\mathbf{x}_2 =$ The first night of the election.	✓
	$\mathbf{x}_1 =$ A white bus with the word Julia is near water [...] $\mathbf{x}_2 =$ There is an open beach in my hometown.	✓
	$\mathbf{x}_1 =$ Strong earthquake in Mexico $\mathbf{x}_2 =$ It’s the best time to get a job	✓
$y = 0$	$\mathbf{x}_1 =$ Closed roads in Armenia $\mathbf{x}_2 =$ Open roads in Azerbaijan	✗
	$\mathbf{x}_1 =$ The man is playing the guitar. $\mathbf{x}_2 =$ I’m not a guitar player.	✗
	$\mathbf{x}_1 =$ A man is playing a large flute. $\mathbf{x}_2 =$ A man is listening to a small flute.	✗

Table 4: A selection of high-quality (✓) and low-quality (✗) examples in STS- \mathbf{x}_1 - \mathbf{x}_2 . Many sentence pairs for $y = 1$ are not similar and have quite different meanings. Some sentence pairs for $y = 0$ are not on completely different topics.

et al. (2021). With appropriate measures for handling noisy data, models trained on datasets generated with DINO achieve strong results on several semantic textual similarity datasets.

For future work, it would be interesting to see whether the noise in datasets generated with DINO can further be reduced, e.g., by using different sets of instructions (Jiang et al., 2020; Schick and Schütze, 2021a) or by supplementing our pipeline with some additional filtering steps.

Acknowledgments This work was funded by the European Research Council (ERC #740516). We thank the anonymous reviewers for their helpful comments.

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 Uria, and Janyce Wiebe. 2015. *SemEval-2015 task 2: Semantic textual similarity, English, Spanish and pilot on inter-*

- pretability. 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, 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.
- Eneko Agirre, Mona Diab, Daniel Cer, and Aitor Gonzalez-Agirre. 2012. Semeval-2012 task 6: A pilot on semantic textual similarity. In *Proceedings of 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 '12*, page 385–393, USA. Association for Computational Linguistics.
- Ateret Anaby-Tavor, Boaz Carmeli, Esther Goldbraich, Amir Kantor, George Kour, Segev Shlomov, Naama Tepper, and Naama Zwerdling. 2020. [Do not have enough data? Deep learning to the rescue!](#) *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7383–7390.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- 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.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askeel, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- 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.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. 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.
- Avia Efrat and Omer Levy. 2020. [The turking test: Can language models understand instructions?](#) *Computing Research Repository*, arXiv:2010.11982.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- William Fedus, Barret Zoph, and Noam Shazeer. 2021. [Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity](#). *Computing Research Repository*, arXiv:2101.03961.

- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). 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 3816–3830, Online. Association for Computational Linguistics.
- John M. Giorgi, Osvald Nitski, Gary D. Bader, and Bo Wang. 2020. [DeCLUTR: Deep contrastive learning for unsupervised textual representations](#). *Computing Research Repository*, arXiv:2006.03659.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Ari Holtzman, Jan Buys, Maxwell Forbes, Antoine Bosselut, David Golub, and Yejin Choi. 2018. [Learning to write with cooperative discriminators](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1638–1649, Melbourne, Australia. Association for Computational Linguistics.
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. [How can we know what language models know?](#) *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Skip-thought vectors](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Varun Kumar, Ashutosh Choudhary, and Eunah Cho. 2021. [Data augmentation using pre-trained transformer models](#). *Computing Research Repository*, arXiv:2003.02245.
- Quoc Le and Tomas Mikolov. 2014. [Distributed representations of sentences and documents](#). In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1188–1196, Beijing, China. PMLR.
- 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.
- 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](#). *Computing Research Repository*, arXiv:1907.11692.
- 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).
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). *Computing Research Repository*, arXiv:1301.3781.
- Biswesh Mohapatra, Gaurav Pandey, Danish Contractor, and Sachindra Joshi. 2020. [Simulated chats for task-oriented dialog: Learning to generate conversations from instructions](#). *Computing Research Repository*, arXiv:2010.10216.
- Yannis Papanikolaou and Andrea Pierleoni. 2020. [DARE: Data augmented relation extraction with GPT-2](#). *Computing Research Repository*, arXiv:2004.13845.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. [Automatic differentiation in PyTorch](#). In *NIPS Autodiff Workshop*.
- 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.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). 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 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Nina Pörner and Hinrich Schütze. 2019. [Multi-view domain adapted sentence embeddings for low-resource unsupervised duplicate question detection](#). 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 2019, Hong Kong, China, November 3-7, 2019*, pages 1630–1641. Association for Computational Linguistics.
- Nina Pörner, Ulli Waltinger, and Hinrich Schütze. 2020. [Sentence meta-embeddings for unsupervised semantic textual similarity](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7027–7034. Association for Computational Linguistics.

- Raul Puri and Bryan Catanzaro. 2019. [Zero-shot text classification with generative language models](#). *Computing Research Repository*, arXiv:1912.10165.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving language understanding by generative pre-training](#).
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). Technical report.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- 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.
- Timo Schick and Hinrich Schütze. 2020. [Few-shot text generation with pattern-exploiting training](#). *Computing Research Repository*, arXiv:2012.11926.
- Timo Schick and Hinrich Schütze. 2021a. [Exploiting cloze questions for few shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics*, Kyiv, Ukraine (Online). International Committee on Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021b. [It’s not just size that matters: Small language models are also few-shot learners](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2339–2352, Online. Association for Computational Linguistics.
- Timo Schick, Sahana Udupa, and Hinrich Schütze. 2021. [Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in NLP](#). *Transactions of the Association for Computational Linguistics*.
- C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. 2016. [Rethinking the inception architecture for computer vision](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826.
- Derek Tam, Rakesh R Menon, Mohit Bansal, Shashank Srivastava, and Colin Raffel. 2021. [Improving and simplifying pattern exploiting training](#). *Computing Research Repository*, arXiv:2103.11955.
- Orion Weller, Nicholas Lourie, Matt Gardner, and Matthew Peters. 2020. [Learning from task descriptions](#). *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- John Wieting and Kevin Gimpel. 2018. [ParaNMT-50M: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Melbourne, Australia. Association for Computational Linguistics.
- 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. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Zhuofeng Wu, Sinong Wang, Jiatao Gu, Madian Khabsa, Fei Sun, and Hao Ma. 2020. [CLEAR: Contrastive learning for sentence representation](#). *Computing Research Repository*, arXiv:2012.15466.
- Yiben Yang, Chaitanya Malaviya, Jared Fernandez, Swabha Swayamdipta, Ronan Le Bras, Ji-Ping Wang, Chandra Bhagavatula, Yejin Choi, and Doug Downey. 2020. [Generative data augmentation for commonsense reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1008–1025, 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.

A Experimental Setup

Our implementation is based on the Transformers library (Wolf et al., 2020) and PyTorch (Paszke et al., 2017). All our experiments were conducted using two GPUs with 11GB RAM (NVIDIA GeForce GTX 1080 Ti). Generating STS- \mathcal{L} - x_1x_2 and STS- \mathcal{L} - x_2 using both GPUs took approximately 48 hours per dataset. Training a Sentence Transformer on these datasets took less than 2 hours on average.

B Datasets

Both datasets generated with DINO (STS- \mathcal{L} - x_1x_2 and STS- \mathcal{L} - x_2) are publicly available at <https://github.com/timoschick/dino>. After filtering out examples where the language model did not produce a quotation mark, STS- \mathcal{L} - x_2 contains 121,275 examples and STS- \mathcal{L} - x_1x_2 contains 143,968 examples.

C Additional Results

Our main results do not include scores for DeCLUTR (Giorgi et al., 2020) and CLEAR (Wu et al., 2020) – two recent approaches using contrastive learning – as their evaluation setup differs from that described in Reimers and Gurevych (2019) (and used by all other baselines) in the following respects:

- Both Giorgi et al. (2020) and Wu et al. (2020) treat SICK and STSb as *supervised* tasks, i.e., they use the provided task-specific training sets to perform regular supervised training.
- The STS12–16 datasets each consist of several subsets. Giorgi et al. (2020) and Wu et al. (2020) compute Spearman’s correlation coefficient separately for each of these subsets and report the mean score across all subsets. In contrast, for our main results we follow Reimers and Gurevych (2019) and concatenate all subsets to form one large set on which Spearman’s correlation is computed just once.

As the implementations of both methods are not publicly available as of this writing, we are unable to compute scores for DeCLUTR and CLEAR using the evaluation setup of Reimers and Gurevych (2019) ourselves. Instead, we recompute scores for DINO (both with STS- \mathcal{L} - x_2 and STS- \mathcal{L} - x_1x_2) using the evaluation setup of Giorgi et al. (2020) and

Wu et al. (2020) on STS12–16; results are shown in Table 5.

Model	STS12	STS13	STS14	STS15	STS16	Avg.
CLEAR	49.0	48.9	57.4	63.6	65.6	56.9
DeCLUTR	64.2	70.4	70.0	77.5	75.4	71.5
STS- \mathcal{L} - x_1x_2	65.1	69.9	68.6	76.3	76.6	71.3
STS- \mathcal{L} - x_2	65.3	71.8	72.7	75.9	76.9	72.5

Table 5: Results for CLEAR (Wu et al., 2020), DeCLUTR (Giorgi et al., 2020) and Sentence-RoBERTa (base) trained on STS- \mathcal{L} - x_1x_2 and STS- \mathcal{L} - x_2 using the evaluation setup of Wu et al. (2020) and Giorgi et al. (2020): For each task, we report the mean Spearman correlation of all subtasks in a fully unsupervised setting.