

Virtual Augmentation Supported Contrastive Learning of Sentence Representations

Dejiao Zhang* Wei Xiao Henghui Zhu Xiaofei Ma Andrew O. Arnold
AWS AI Labs, New York

Abstract

Despite profound successes, contrastive representation learning relies on carefully designed data augmentations using domain-specific knowledge. This challenge is magnified in natural language processing, where no general rules exist for data augmentation due to the discrete nature of natural language. We tackle this challenge by presenting a Virtual Augmentation Supported Contrastive Learning of sentence representations (VaSCL). Originating from the interpretation that data augmentation essentially constructs the neighborhoods of each training instance, we in turn utilize the neighborhood to generate effective data augmentations. Leveraging the large training batch size of contrastive learning, we approximate the neighborhood of an instance via its K -nearest in-batch neighbors in the representation space. We then define an instance discrimination task regarding the neighborhood and generate the virtual augmentation in an adversarial training manner. We assess the performance of VaSCL on a wide range of downstream tasks and set a new state-of-the-art for unsupervised sentence representation learning.

1 Introduction

Universal sentence representation learning has been a long-standing problem in Natural Language Processing (NLP). Leveraging the distributed word representations (Bengio et al., 2003; Mikolov et al., 2013; Collobert et al., 2011; Pennington et al., 2014) as the base features to produce sentence representations is a common strategy in the early stage. However, these approaches are tailored to different target tasks, thereby yielding less generic sentence representations (Yessenalina and Cardie, 2011; Socher et al., 2013; Kalchbrenner et al., 2014; Cho et al., 2014).

* The code and pretrained checkpoints can be found at <https://github.com/amazon-research/sentence-representations>. Correspondence to Dejiao Zhang <dejiaoz@amazon.com>.

This issue has motivated more research efforts on designing generic sentence-level learning objectives or tasks. Among them, supervised learning on the Natural Language Inference (NLI) datasets (Bowman et al., 2015a; Williams et al., 2017; Wang et al., 2018) has established benchmark transfer learning performance on various downstream tasks (Conneau et al., 2017; Cer et al., 2018; Reimers and Gurevych, 2019a; Zhang et al., 2021). Despite promising progress, the high cost of collecting annotations precludes its wide applicability, especially when the target domain has scarce annotations but differs significantly from the NLI datasets (Zhang et al., 2020).

On the other hand, unsupervised learning of sentence representations has seen a resurgence of interest with the recent successes in self-supervised contrastive learning. These approaches rely on two main components, data augmentation and an instance-level contrastive loss. The popular contrastive learning objectives Chen et al. (2020); He et al. (2020) and their variants thereof have empirically shown their effectiveness in NLP. However, the discrete nature of the text makes it challenging to establish universal rules for effective text augmentation generation.

Various contrastive learning based approaches have been proposed for sentence representation learning, where the main difference lies in how the augmentations are generated (Fang and Xie, 2020; Giorgi et al., 2020; Wu et al., 2020; Meng et al., 2021; Yan et al., 2021; Kim et al., 2021; Gao et al., 2021). Somewhat surprisingly, a recent work (Gao et al., 2021) shows that Dropout (Srivastava et al., 2014), *i.e.*, *augmentations obtained by feeding the same instance to the encoder twice*, outperforms common data augmentations obtained by operating on the text directly, including cropping, word deletion, or synonym replacement. Again, this observation validates the inherent difficulty of attaining effective data augmentations in NLP.

This paper tackles the challenge by presenting a neighborhood-guided virtual augmentation strategy to support contrastive learning. In a nutshell, data augmentation essentially constructs the neighborhoods of each instance, with the semantic content being preserved. We take this interpretation in the opposite direction by leveraging the neighborhood of an instance to guide augmentation generation. Benefiting from the large training batch of contrastive learning, we approximate the neighborhood of an instance via its K-nearest in-batch neighbors. We then define an instance discrimination task within this neighborhood and generate the virtual augmentation in an adversarial training manner. We run in-depth analyses and show that our VaSCL model leads to a more dispersed representation space with the data semantics at different granularities being better captured. We evaluate our model on a wide range of downstream tasks and show that our model consistently outperforms the previous state-of-the-art results by a large margin.

2 Related Work

Universal Sentence Representation Learning

Arguably, the simplest and most common approaches for attaining sentence representations are bag-of-words (Harris, 1954) and variants thereof. However, bag-of-words suffers from data sparsity and a lack of sensibility to word semantics. In the past two decades, the distributed word representations (Bengio et al., 2003; Mikolov et al., 2013; Collobert et al., 2011; Pennington et al., 2014) have become the more effective base features for producing sentence representations. The downside is that these approaches are tailored to the target tasks (Yessenalina and Cardie, 2011; Socher et al., 2013; Kalchbrenner et al., 2014; Cho et al., 2014), and thereby the resulting sentence representations attain limited transfer learning performance.

More recent efforts focus on directly designing the sentence-level learning objectives or tasks. On the supervised learning regime, Conneau et al. (2017); Cer et al. (2018) empirically show the effectiveness of leveraging the NLI task (Bowman et al., 2015a; Williams et al., 2017) to promote generic sentence representations. The task involves classifying each sentence pair into one of three categories: entailment, contradiction, or neutral. Reimers and Gurevych (2019b) further bolster the performance by using the pre-trained transformer (Devlin et al., 2018; Liu et al., 2019) as backbone.

On the other end of the spectrum, Hill et al. (2016); Bowman et al. (2015b) propose using the denoising or variational autoencoders for sentence representation learning. Kiros et al. (2015); Hill et al. (2016) extend the distributional hypothesis to the sentence level and train an encoder-decoder to construct the surrounding context for each sentence. Alternatively, Logeswaran and Lee (2018) present a model that learns to discriminate the target context sentences from all contrastive ones.

Contrastive Learning Contrastive learning has been the pinnacle of recent successes in sentence representation learning. Gao et al. (2021); Zhang et al. (2021) substantially advance the previous state-of-the-art results by leveraging the entailment sentences in NLI as positive pairs for optimizing the properly designed contrastive loss functions. Nevertheless, we focus on unsupervised contrastive learning and form the positive pairs via data augmentation since such methods are more cost-effective and applicable across different domains and languages. Along this line, several approaches have been proposed recently, where the augmentations are obtained via dropout (Yan et al., 2021; Gao et al., 2021), back-translation (Fang and Xie, 2020), surrounding context sampling (Logeswaran and Lee, 2018; Giorgi et al., 2020), or perturbations conducted at different semantic-level (Wu et al., 2020; Yan et al., 2021; Meng et al., 2021).

Consistency Regularization Our work is also closely related to consistency regularization, which is often used to promote better performance by regularizing the model output to remain unchanged under plausible input variations that are often induced via data augmentations. Bachman et al. (2014); Sajjadi et al. (2016); Samuli and Timo (2017); Tarvainen and Valpola (2017) show randomized data augmentations such as dropout, cropping, rotation, and flipping yield effective regularization. Berthelot et al. (2019, 2020); Verma et al. (2019) improve the performance by applying Mixup (Zhang et al., 2017) and its variants on top of stochastic data augmentations. However, data augmentation has long been a challenge in NLP as there are no general rules for effective text transformations. An alternative that comes to light when considering the violation of consistency regularization can in turn be used to find the most sensitive perturbation for a model. Therefore, we utilize consistency regularization to promote informative virtual augmen-

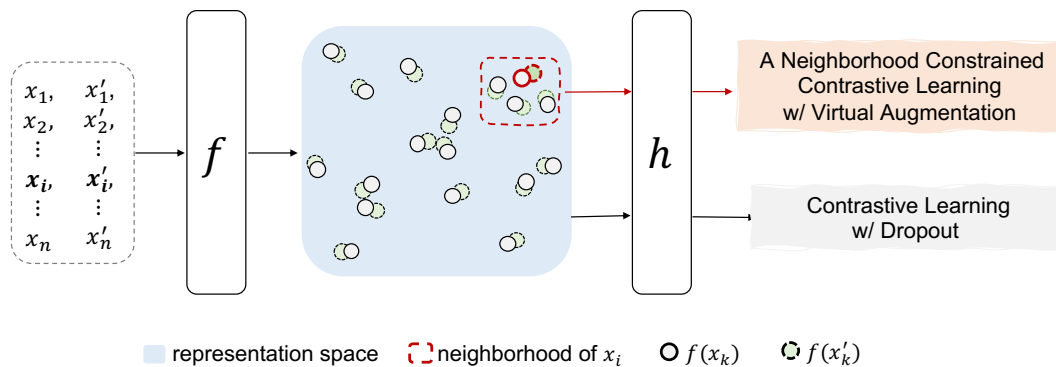


Figure 1: Illustration of VaSCL. For each instance x_i in a randomly sampled batch, we optimize (i) an instance-wise contrastive loss with the dropout induced augmentation obtained by forwarding the same instance twice, *i.e.*, x_i and $x_{i'}$ denote the same text example; and (2) a neighborhood constrained instance discrimination loss with the virtual augmentation proposed in Section 3.2.

tation for a training instance in the representation space while leveraging its approximated neighborhood to regularize the augmentation sharing similar semantic content as its original instance.

3 Method

3.1 Preliminaries

Self-supervised contrastive learning often aims to solve the instance discrimination task. In our scenario, let f denote the transformer encoder that maps the i^{th} input sentence \mathbf{x}_i to its representation vector $\mathbf{e}_i = f(\mathbf{x}_i)$ ¹. Further let h be the contrastive learning head and $\mathbf{z}_i = h(f(\mathbf{x}_i))$ denote the final output for \mathbf{x}_i . Let $\mathcal{B} = \{i, i'\}_{i=1}^M$ denote the indices of a randomly sampled batch of paired examples, where $\mathbf{x}_i, \mathbf{x}_{i'}$ are two independent variations of the i^{th} instance. A popular loss function (Chen et al., 2020) for contrastive learning is defined as follows,

$$\ell_{\mathcal{B}}(\mathbf{z}_i, \mathbf{z}_{i'}) = -\log \frac{e^{\text{sim}(\mathbf{z}_i, \mathbf{z}_{i'})/\tau}}{e^{\text{sim}(\mathbf{z}_i, \mathbf{z}_{i'})/\tau} + \sum_{j \in \mathcal{B} \setminus \{i, i'\}} e^{\text{sim}(\mathbf{z}_i, \mathbf{z}_j)/\tau}}, \quad (1)$$

where τ is the temperature hyper-parameter and $\text{sim}(\cdot)$ denotes the cosine similarity, *i.e.*, $\text{sim}(\cdot) = \mathbf{z}_i^T \mathbf{z}_{i'} / \|\mathbf{z}_i\|_2 \|\mathbf{z}_{i'}\|_2$. Similarly, $\ell_{\mathcal{B}}(\mathbf{z}_{i'}, \mathbf{z}_i)$ is defined by exchanging the roles of \mathbf{z}_i and $\mathbf{z}_{i'}$ in the above equation. Intuitively, Equation (1) defines the log-likelihood of classifying the i^{th} instance as its positive i' among all $2M-1$ candidates within

¹By an abuse of notation, we assume f outputs either the pre-defined sentence representation (a.k.a. [CLS] embedding, (Devlin et al., 2018)), or the mean/max pooling of all tokens' embeddings of that sentence.

the same batch \mathcal{B} . Therefore, minimizing the above log-loss guides the encoder to map each positive pair close in the representation space, and negative pairs further apart.

Dropout based contrastive learning As Equation (1) implies, the success of contrastive learning relies on effective positive pairs construction. However, it is challenging to generate strong and effective data transformations in NLP due to the discrete nature of natural language. This challenge is further demonstrated in a recent work (Gao et al., 2021), which shows that augmentations obtained by Dropout (Srivastava et al., 2014), *i.e.*, $\mathbf{z}_i, \mathbf{z}_{i'}$ obtained by forwarding the same instance \mathbf{x}_i twice, outperforms the common text augmentation strategies such as cropping, word deletion, or synonym replacement. Dropout provides a natural data augmentation by randomly masking its inputs or the hidden layer nodes. The effectiveness of using Dropout as pseudo data augmentations can be traced back to Bachman et al. (2014); Samuli and Timo (2017); Tarvainen and Valpola (2017). Nevertheless, the augmentation strength is weak with Dropout only. There is room for improvement, which we investigate in the following section.

3.2 Neighborhood Constrained Contrastive Learning with Virtual Augmentation

In essence, data augmentation can be interpreted as constructing the neighborhood of a training instance, with the semantic content being preserved. In this section, we take the interpretation in the opposite direction and leverage the neighborhoods of each instance to generate the augmentation. To be more specific, let $\bar{\mathcal{B}} = \{i\}_{i=1}^M$ denote the indices

of a randomly sampled batch with M examples. We first approximate the neighborhood $\mathcal{N}(i)$ of the i^{th} instance as its K -nearest neighbors in the representation space,

$$\mathcal{N}(i) = \{k : \mathbf{e}_k \text{ has the top-}K \text{ similarity with } \mathbf{e}_i \text{ among all other } M-1 \text{ instances in } \bar{\mathcal{B}}\}$$

We then define an instance-level contrastive loss regarding the i^{th} instance and its neighborhood as follows,

$$\ell_{\mathcal{N}(i)}(\mathbf{z}_i^\delta, \mathbf{z}_i) = -\log \frac{e^{\text{sim}(\mathbf{z}_i^\delta, \mathbf{z}_i)/\tau}}{e^{\text{sim}(\mathbf{z}_i^\delta, \mathbf{z}_i)/\tau} + \sum_{k \in \mathcal{N}(i)} e^{\text{sim}(\mathbf{z}_i^\delta, \mathbf{z}_k)/\tau}}. \quad (2)$$

In the above equation, $\mathbf{z}_i^\delta = h(\mathbf{e}_i^\delta)$ denotes the output of the contrastive learning head with the perturbed representation $\mathbf{e}_i^\delta = \mathbf{e}_i + \delta_i$ as input. Here, the initial perturbation δ_i is chosen as isotropic² Gaussian noise. As it implies, Equation (2) shows the negative log-likelihood of classifying the perturbed i^{th} instance as itself rather than its neighbors. Then the augmentation of the i^{th} instance is retained by identifying the optimal perturbation that maximally disturbs its instance-level identity within the neighborhood. That is,

$$\begin{aligned} \delta_i^* &= \arg \max_{\|\delta_i\|_2 \leq \Delta} \ell_{\mathcal{N}(i)}(\mathbf{z}_i^\delta, \mathbf{z}_i), \\ \mathbf{e}_{i^*} &= \mathbf{e}_i + \delta_i^*. \end{aligned} \quad (3)$$

For the i^{th} instance, denote $\mathcal{N}_A(i)$ as the augmented neighborhood that consists of its K nearest neighbors and their associated augmentations. That is, $\mathcal{N}_A(i) = \{k, k^*\}_{k=1}^K$ with \mathbf{e}_k and \mathbf{e}_{k^*} denoting the original representation and the augmented representation of the k^{th} nearest neighbor of instance i , respectively. Here, each augmentation \mathbf{e}_{k^*} is obtained by solving Equations (3) with respect to the neighborhood $\mathcal{N}(k)$ of e_k . We then discriminate the i^{th} instance and its augmentation from the augmented neighborhood $\mathcal{N}_A(i)$,

$$\ell_{\mathcal{N}_A(i)} = \ell_{\mathcal{N}_A(i)}(\mathbf{z}_i^*, \mathbf{z}_i) + \ell_{\mathcal{N}_A(i)}(\mathbf{z}_i, \mathbf{z}_i^*). \quad (4)$$

Here both terms on the right hand side are defined in the same way as Equation (2) with respect to the augmentation \mathbf{e}_i^* and the augmented neighborhood $\mathcal{N}_A(i)$ of the i^{th} instance.

² δ_i is sampled from a multivariate normal distribution with the covariance matrix being a scaled identity matrix.

Putting it all together Therefore, for each randomly sampled minibatch \mathcal{B} with M samples, we minimize the following:

$$\begin{aligned} \mathcal{L}_{\text{VaSCL}} &= \frac{1}{2M} \sum_{i=1}^M \{ \ell_{\bar{\mathcal{B}}}(\mathbf{z}_i, \mathbf{z}_{i'}) + \ell_{\bar{\mathcal{B}}}(\mathbf{z}_{i'}, \mathbf{z}_i) \\ &\quad + \ell_{\mathcal{N}_A(i)}(\mathbf{z}_i, \mathbf{z}_i^*) + \ell_{\mathcal{N}_A(i)}(\mathbf{z}_i^*, \mathbf{z}_i) \} \end{aligned} \quad (5)$$

The last two terms of the right hand side are defined in Equation 4. Notice that, $\ell_{\bar{\mathcal{B}}}(\mathbf{z}_i, \mathbf{z}_{i'})$ is defined in the same way as Equation (1) except that $\mathbf{z}_i, \mathbf{z}_{i'}$ are retained by feeding the i^{th} instance in $\bar{\mathcal{B}}$ to the encoder twice. In summary, two instance discrimination tasks are posed for each training example: i) discriminating each instance and its dropout induced variation from the other in-batch instances; and ii) separating each instance and its virtual augmentation from its K nearest neighbors and their associated virtual augmentations.

4 Experiment

In this section, we mainly evaluate VaSCL against SimCSE (Gao et al., 2021) which leverages the dropout (Srivastava et al., 2014) induced noise as data augmentation. We show that VaSCL consistently outperforms SimCSE on various downstream tasks that involve semantic understanding at different granularities. We carefully study the regularization effects of VaSCL and empirically demonstrate that VaSCL leads to a more dispersed representation space with semantic structure better encoded. Please refer to Appendix A for details of our implementations and the dataset being used.

4.1 Evaluation Datasets

In addition to the popular semantic textual similarity (a.k.a STS) related tasks, we evaluate two additional downstream tasks, short text clustering and few-shot learning based intent classification. Our motivation is twofold. First, these two tasks provide a new evaluation aspect that complements the pairwise similarity-oriented STS evaluation by assessing the high-level categorical semantics encoded in the representations. Second, two desired challenges are posted as short text clustering requires more effective representations due to the weak signal each text example manifests; and intent classification often suffers from data scarcity since the intents can vary significantly over different dialogue systems and the intent examples are costly to collect.

	STS12	STS13	STS14	STS15	STS16	SICK-R	STS-B	Avg.
RoBERTa _{distil}	54.41	46.85	56.96	65.79	64.22	61.10	59.01	58.33
SimCSE _{distil}	65.58	77.42	70.17	79.31	78.45	67.66	77.98	73.79
VaSCL_{distil}	67.68	80.61	72.19	80.92	78.59	68.81	77.32	75.16
RoBERTa _{base}	53.95	47.42	55.87	64.73	63.55	62.94	58.40	58.12
SimCSE _{base}	68.88	80.46	73.54	80.98	80.68	69.54	80.29	76.34
VaSCL_{base}	69.02	82.38	73.93	82.54	80.96	69.40	80.52	76.96
RoBERTa _{large}	55.00	50.14	54.87	62.14	62.99	58.93	54.56	56.95
SimCSE _{large}	69.83	81.29	74.42	83.77	79.79	68.89	80.66	76.95
VaSCL_{large}	73.36	83.55	77.16	83.25	80.66	72.96	82.36	79.04

Table 1: Spearman rank correlation between the cosine similarity of sentence representation pairs and the ground truth similarity scores.

Semantic Textual Similarity The semantic textual similarity (STS) tasks are the most commonly used benchmark for evaluating sentence representations. STS consists of seven tasks, namely STS 2012-2016 (Agirre et al., 2012, 2013, 2014, 2015, 2016), the STS Benchmark (Cer et al., 2017), and the SICK-Relatedness (Marelli et al., 2014). For each sentence pair in these datasets, a fine-grained similarity score ranges from 0 to 5 is provided.

Short Text Clustering Compared with general text clustering, short text clustering has its own challenge due to lack of signal. Nevertheless, texts containing only a few words grow at unprecedented rates from a wide range of popular resources, including Reddit, Stackoverflow, Twitter, and Instagram. Clustering those texts into groups of similar texts plays a crucial role in many real-world applications such as topic discovery (Kim et al., 2013), trend detection (Mathioudakis and Koudas, 2010), and recommendation (Bouras and Tsogkas, 2017). We evaluate six benchmark datasets for short text clustering. As shown in Table 4, the datasets present the desired diversities regarding both the cluster sizes and the number of clusters contained in each dataset.

Intent Classification Intent classification aims to identify the intents of user utterances, which is a critical component of goal-oriented dialog systems. Attaining high intent classification accuracy is an important step towards solving many downstream tasks such as dialogue state tracking (Wu et al., 2019; Zhang et al., 2019) and dialogue management (Gao et al., 2018; Ham et al., 2020). A practical challenge is data scarcity because differ-

ent systems define different sets of intents, and it is costly to obtain enough utterance samples for each intent. Therefore, few-shot learning has attracted much attention under this scenario, which is also our main focus. We evaluate four intent classification datasets originating from different domains. We summarize the data statistics in Appendix B.1.

4.2 Main Results

4.2.1 Evaluation Setup

Semantic Textual Similarity. Same as Reimers and Gurevych (2019b); Gao et al. (2021), in Table 1 we report the Spearman correlation³ between the cosine similarity of the sentence representation pairs and the ground truth similarity scores. **Short Text Clustering.** We evaluate the sentence representations using K-Means (MacQueen et al., 1967; Lloyd, 1982) given its simplicity and report the clustering accuracy⁴ averaged over ten independent runs in Table 2. **Intent Classification.** We freeze the transformer and fine-tune a linear classification layer with the softmax-based cross-entropy loss. We merge the training and validation sets, from which we sample K training and validation samples per class. We report the mean and standard deviation of the testing classification accuracy evaluated over five different splits in Table 3.⁵ We set the learning rate to $1e-04$ and batch size to 32. For each task, we train the model with 1000 iterations

³Same as Reimers and Gurevych (2019b); Gao et al. (2021), we concatenate all the topics and report the overall Spearman’s correlation.

⁴The clustering accuracy is computed by using the Hungarian algorithm (Munkres, 1957).

⁵In each setting, we fix the five different splits for all models.

	Ag News	Search Snippets	Stack Overflow	Bio-medical	Tweet	Google News	Avg
RoBERTa _{distil}	59.32	33.18	14.16	24.69	37.10	58.05	37.75
SimCSE _{distil}	73.33	60.74	66.97	35.69	50.68	67.55	59.16
VaSCL_{distil}	71.71	62.76	73.98	38.82	51.35	67.66	61.05
RoBERTa _{base}	66.50	30.83	15.63	26.98	37.80	58.51	39.38
SimCSE _{base}	65.53	55.97	64.18	38.12	49.16	65.69	56.44
VaSCL_{base}	68.33	47.26	76.15	39.53	51.50	67.10	58.31
RoBERTa _{large}	69.35	53.00	27.89	33.25	46.08	64.04	48.93
SimCSE _{large}	62.93	51.55	54.11	35.39	50.92	67.86	53.79
VaSCL_{large}	66.09	61.57	69.04	42.91	56.74	67.75	60.68

Table 2: Clustering accuracy reported on six short text clustering datasets.

and evaluate the validation set every 100 iterations. We report the testing accuracy on the checkpoint achieving the best validation accuracy.

		SNIPS	BANK77	CLINC150	HWU64
5-Shot	RoBERTa	76.71 \pm 4.84	38.77 \pm 2.29	55.19 \pm 1.99	51.52 \pm 2
	SimCSE	76.94 \pm 2.53	67.48 \pm 1.63	72.84 \pm 1.5	66.1 \pm 1.9
	VaSCL	78.51\pm1.39	70.10\pm1.76	74.23\pm1.17	67.06\pm2.17
10-Shot	RoBERTa	85.63 \pm 2.43	46.55 \pm 1.84	60.55 \pm 1.16	57.47 \pm 0.91
	SimCSE	85.14\pm2.18	72.19 \pm 0.88	77.13 \pm 0.76	70.87 \pm 1.35
	VaSCL	84.83 \pm 1.05	75.25\pm0.81	79.15\pm0.82	72.43\pm1.12
20-Shot	RoBERTa	88.14 \pm 1.54	51.65 \pm 1.42	63.51 \pm 1.08	60.93 \pm 1.27
	SimCSE	88.43 \pm 1.2	75.13 \pm 0.78	78.59 \pm 0.78	74.44 \pm 0.74
	VaSCL	89.11\pm1.29	78.06\pm0.37	81.39\pm0.60	76.39\pm0.26

Table 3: Few-shot learning evaluation of Intent Classification. Each result is aggregated over 5 independent splits. We choose RoBERTa-base as backbone.

4.2.2 Evaluation Results

We report the evaluation results in Tables 1, 2, and 3. As we can see, both SimCSE and VaSCL largely improve the performance of the pre-trained language models, while VaSCL consistently outperforms SimCSE on most tasks. To be more specific, we attain 0.6% – 2.1% averaged absolute improvement over SimCSE on seven STS tasks and 1.8% – 6.9% averaged absolute improvement on six short text clustering tasks. We also achieved considerable improvement over SimCSE on intent classification tasks under different few-shot learning scenarios. We do not include the evaluation on ATIS in Table 3 as this dataset is highly imbalanced with one single class account for more than 73% of the data. Please refer to Appendix C for details.

4.3 Analysis


To better understand what enables the good performance of VaSCL, we carefully analyze the representations at different semantic granularities.

Neighborhood Evaluation on Categorical Data

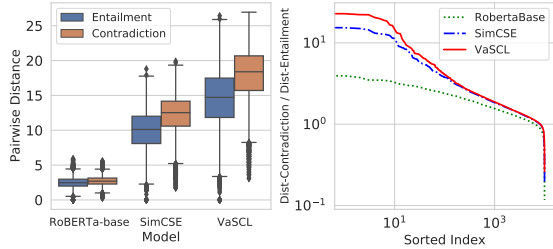
We first evaluate the neighborhood statistics on StackOverflow (Xu et al., 2017) which contains 20 balanced categories, each with 1000 text instances. For each instance, we retrieve its K nearest (top-K) neighbors in the representation space, among which those from the same class as the instance itself as treated as positives. In Figure 2a, we report both the percentage of true positives and the average distance of an instance to its top-K neighbors. For each top-K value, the evaluation is averaged over all 20,000 instances.

As indicated by the small distance values reported in Figure 2a, the representation space of the original RoBERTa model is tighter and is incapable of uncovering the categorical structure of data. In contrast, both VaSCL and SimCSE are capable of scattering representations apart while better capturing the semantic structures. Compared with SimCSE, VaSCL leads to even more dispersed representations with categorical structures being better encoded. This is also demonstrated by the better performance attained on both clustering and few-shot learning reported in Tables 2&3.

Fine-grained Semantic Understanding We then compare VaSCL against SimCSE and RoBERTa on encoding more fine-grained semantic concepts. We randomly sample 20,000 premises from the combined set of SNLI (Bowman et al., 2015a) and MNLI (Williams et al., 2017), where



(a) Neighborhood evaluation on StackOverflow. Instances from the same category are treated as true positives.



(b) Fine-grained semantics encoding evaluation on NLI.

Figure 2: VaSCL leads to more dispersed representation with data structure being better uncovered.

the associated entailment and contradiction hypotheses are also sampled for each premise instance. In Figure 2b, we report both the distributions of the pairwise distances of the entailment or the contradiction pairs (left). While on the right-hand side, we plot the distance of each premise to its entailment hypothesis over that to its contradiction hypothesis (right).

We observe the same trend that both SimCSE and VaSCL well separate different instances apart in the representation space while better discriminating each premise’s entailment hypothesis from the contradiction one. Figure 2b also demonstrates that VaSCL outperforms SimCSE on better capturing the fine-grained semantics when separating different instances apart. This advantage of VaSCL is further validated by Table 1, where VaSCL consistently outperforms SimCSE on the STS tasks that require pairwise semantic inference on an even more fine-grained scale.

4.4 Explicit Data Augmentation

To better evaluate our virtual augmentation-oriented VaSCL model, we compare it against different explicit data augmentation strategies that directly operate on the discrete text. Specifically, we consider the following approaches:⁶ *WDel* (random word deletion) removes words from the input

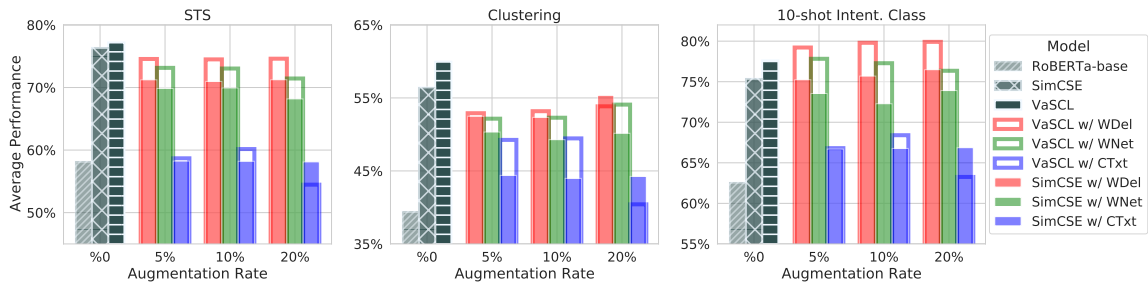
⁶They are implemented using the `nlpaug` library <https://github.com/makcedward/nlpaug>.

text randomly; *WNet* (WordNet synonym substitute) transforms a text instance by replacing its words with the WordNet synonyms (Morris et al., 2020; Ren et al., 2019); and *CTxt* (contextual synonyms substitute) leverages the pre-trained transformers to find top-n suitable words of the input text for substitution (Kobayashi, 2018). For each strategy, we evaluate three augmentation strengths by partially changing 5%, 10%, and 20% words of each text instance. For a positive pair (x_i, x'_i) , x_i denotes the original text and x'_i is the associated augmentation. We also explore the case where both x_i and x'_i are the transformations of the original text, which we find yielding worse performance.

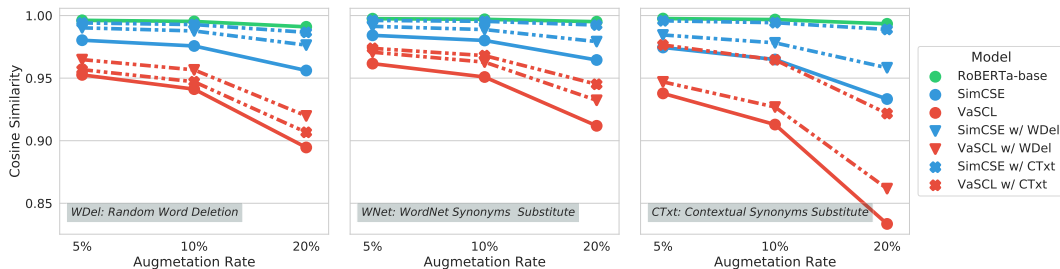
Virtual Augmentation Performs Better The performance of explicit text augmentation is evaluated using the standard dropout for training, *i.e.*, "SimCSE w/ {WDel/WNet/CTxt}" in Figure 3. As Figure 3a shows, contrastive learning with moderate explicit text augmentations, *i.e.*, augmentation strength less than 20%, does yield better sentence representations when compared with the original RoBERTa model. Nevertheless, both virtual augmentation strategies, *i.e.*, SimCSE & VaSCL, substantially outperform all three explicit text augmentation strategies on almost all downstream tasks. Although a bit surprising, especially considering the performance gap between SimCSE and explicit augmentations, this comparison provides a new perspective on interpreting the underlying challenge of designing effective transformations that operate on the discrete text directly.

VaSCL Outperforms SimCSE Figure 3a also empirically demonstrates that VaSCL outperforms SimCSE no matter in the presence of explicit text augmentations or not. The only exception occurs when the explicit augmentation strength is too large, *i.e.*, 20% of the words of each text are perturbed. One possible explanation is that undesired noises are generated by the large perturbations on discrete texts directly, which can violate the coherent semantics maintained by a neighborhood and hence make it hard for VaSCL to generate effective virtual augmentations.

New Linguistic Patterns Are Required Another observation drawn from Figure 3a is that both SimCSE and VaSCL attain worse performance on most downstream tasks when combining with explicit text augmentations. Although VaSCL does improve the performance of explicit augmentations



(a) Virtual augmentation vs. explicit augmentation. For each downstream task, we report the mean performance averaged over all its subtasks. The explicit augmentations are evaluated using SimCSE (dropout) for training, *i.e.*, "SimCSE w/ {WDel/WNet/CTxt}".



(b) Cosine similarity between each original training example and its augmentation evaluated on the representation spaces of different models. From left to right, the augmentations are obtained via *WDel*, *WNet*, and *CTxt*. Each point is averaged over 20,000 randomly sampled training examples and their augmentations. We exclude "SimCSE w/ WNet" and "VaSCL w/ WNet" for better visualization. Please refer to Figure 4 in Appendix for the full plot.

Figure 3: Comparing and combining virtual augmentation with explicit augmentation.

in most cases, this is undesired as we expect a win-win outcome that moderate explicit augmentations could further enhance VaSCL. We hypothesize that new and informative linguistic patterns are missing for the expected performance gain.

To validate our hypothesis, in Figure 3b we report the cosine similarity between each original training example and its augmentation evaluated on the representation spaces of different models. Our observation is twofold. First, the representations induced by RoBERTa and the one trained with contextual synonyms substitution ("SimCSE w/ CTxt") are very similar in all three settings, which also explains why "SimCSE w/ WDel" attains similar performance as RoBERTa on the downstream tasks. We attribute this to the fact that CTxt leverages the transformer itself to generate augmentations which hence carry limited unseen and effective linguistic patterns. Second, as indicated by the comparatively smaller similarity values in Figure 3b, the incorporation of explicit augmentations tightens the representation spaces of both SimCSE and VaSCL, which also results in a worse performance of downstream tasks. One possible explanation is that all the three explicit augmentations are weak and noisy, which harms both the instance discrim-

ination force and the semantic relevance of each neighborhood.

5 Conclusion

In this paper, we present a virtual augmentation-oriented contrastive learning framework for unsupervised sentence representation learning. Our key insight is that data augmentation can be interpreted as constructing the neighborhoods of each training instance, which can, in turn, be leveraged to generate effective data augmentations. We evaluate VaSCL on a wide range of downstream tasks and substantially advance the state-of-the-art results. Moreover, we conduct in-depth analyses and show that VaSCL leads to a more dispersed representation space with the data semantics at different granularities being better encoded.

On the other hand, we observe a performance drop of both SimCSE and VaSCL when combined with the explicit text augmentations. We suspect this is caused by the linguistic patterns generated by explicit augmentations being less informative yet noisy. We hypothesize effective data augmentation operations on the discrete texts could complement our virtual augmentation approach if new and informative linguistic patterns are generated.

References

- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Inigo Lopez-Gazpio, Montse Maritxalar, Rada Mihalcea, et al. 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.
- 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.
- Eneko Agirre, Carmen Banea, Daniel Cer, Mona Diab, Aitor Gonzalez Agirre, Rada Mihalcea, German Rigau Claramunt, and Janyce Wiebe. 2016. Semeval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation. In *SemEval-2016. 10th International Workshop on Semantic Evaluation; 2016 Jun 16-17; San Diego, CA. Stroudsburg (PA): ACL; 2016. p. 497-511. ACL (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.
- 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.
- Philip Bachman, Ouais Alsharif, and Doina Precup. 2014. Learning with pseudo-ensembles. *Advances in neural information processing systems*, 27:3365–3373.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. A neural probabilistic language model. *The journal of machine learning research*, 3:1137–1155.
- David Berthelot, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. 2020. [Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring](#). In *International Conference on Learning Representations*.
- David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. 2019. Mixmatch: A holistic approach to semi-supervised learning. *arXiv preprint arXiv:1905.02249*.
- Christos Bouras and Vassilis Tsogkas. 2017. Improving news articles recommendations via user clustering. *International Journal of Machine Learning and Cybernetics*, 8(1):223–237.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015a. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. 2015b. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349*.
- Inigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. Efficient intent detection with dual sentence encoders. *arXiv preprint arXiv:2003.04807*.
- Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation. *arXiv preprint arXiv:1708.00055*.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Céspedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*.
- Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of machine learning research*, 12(ARTICLE):2493–2537.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loic Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. *arXiv preprint arXiv:1705.02364*.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Calta-girone, Thibaut Lavril, et al. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *arXiv preprint arXiv:1805.10190*.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Hongchao Fang and Pengtao Xie. 2020. Cert: Contrastive self-supervised learning for language understanding. *arXiv preprint arXiv:2005.12766*.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural approaches to conversational ai. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1371–1374.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.
- John M Giorgi, Osvald Nitski, Gary D Bader, and Bo Wang. 2020. Declutr: Deep contrastive learning for unsupervised textual representations. *arXiv preprint arXiv:2006.03659*.
- Donghoon Ham, Jeong-Gwan Lee, Youngsoo Jang, and Kee-Eung Kim. 2020. End-to-end neural pipeline for goal-oriented dialogue systems using gpt-2. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 583–592.
- Zellig S Harris. 1954. Distributional structure. *Word*, 10(2-3):146–162.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738.
- Charles T Hemphill, John J Godfrey, and George R Doddington. 1990. The atis spoken language systems pilot corpus. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- 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.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. *arXiv preprint arXiv:1404.2188*.
- Hwi-Gang Kim, Seongjoo Lee, and Sunghyon Kyeong. 2013. Discovering hot topics using twitter streaming data social topic detection and geographic clustering. In *2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2013)*, pages 1215–1220. IEEE.
- Taeuk Kim, Kang Min Yoo, and Sang-goo Lee. 2021. Self-guided contrastive learning for bert sentence representations. *arXiv preprint arXiv:2106.07345*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2015. Skip-thought vectors. *arXiv preprint arXiv:1506.06726*.
- Sosuke Kobayashi. 2018. Contextual augmentation: Data augmentation by words with paradigmatic relations. *arXiv preprint arXiv:1805.06201*.
- Stefan Larson, Anish Mahendran, Joseph J Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K Kummerfeld, Kevin Leach, Michael A Laurenzano, Lingjia Tang, et al. 2019. An evaluation dataset for intent classification and out-of-scope prediction. *arXiv preprint arXiv:1909.02027*.
- Xingkun Liu, Arash Eshghi, Pawel Swietojanski, and Verena Rieser. 2021. Benchmarking natural language understanding services for building conversational agents. In *Increasing Naturalness and Flexibility in Spoken Dialogue Interaction*, pages 165–183. Springer.
- 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 arXiv:1907.11692*.
- Stuart Lloyd. 1982. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137.
- Lajanugen Logeswaran and Honglak Lee. 2018. An efficient framework for learning sentence representations. *arXiv preprint arXiv:1803.02893*.
- James MacQueen et al. 1967. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, Roberto Zamparelli, et al. 2014. A sick cure for the evaluation of compositional distributional semantic models. In *Lrec*, pages 216–223. Reykjavik.
- Michael Mathioudakis and Nick Koudas. 2010. Twittermonitor: trend detection over the twitter stream. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, pages 1155–1158.

- 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 arXiv:2102.08473*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- John X. Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. 2020. [Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in nlp](#).
- James Munkres. 1957. Algorithms for the assignment and transportation problems. *Journal of the society for industrial and applied mathematics*, 5(1):32–38.
- 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-W*.
- Jeffrey Pennington, Richard Socher, and Christopher D 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.
- Xuan-Hieu Phan, Le-Minh Nguyen, and Susumu Horiguchi. 2008. Learning to classify short and sparse text & web with hidden topics from large-scale data collections. In *Proceedings of the 17th international conference on World Wide Web*, pages 91–100.
- Md Rashadul Hasan Rakib, Norbert Zeh, Magdalena Jankowska, and Evangelos Milios. 2020. Enhancement of short text clustering by iterative classification. *arXiv preprint arXiv:2001.11631*.
- Nils Reimers and Iryna Gurevych. 2019a. Sentencebert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Nils Reimers and Iryna Gurevych. 2019b. [Sentencebert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. Generating natural language adversarial examples through probability weighted word saliency. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 1085–1097.
- Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. 2016. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in neural information processing systems*, 29:1163–1171.
- Laine Samuli and Aila Timo. 2017. Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations (ICLR)*, volume 4, page 6.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- 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.
- Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv preprint arXiv:1703.01780*.
- Vikas Verma, Kenji Kawaguchi, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. 2019. Interpolation consistency training for semi-supervised learning. *arXiv preprint arXiv:1903.03825*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Adina Williams, Nikita Nangia, and Samuel R Bowman. 2017. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426*.
- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019. Transferable multi-domain state generator for task-oriented dialogue systems. *arXiv preprint arXiv:1905.08743*.
- Zhuofeng Wu, Sinong Wang, Jiatao Gu, Madian Khabsa, Fei Sun, and Hao Ma. 2020. Clear: Contrastive learning for sentence representation. *arXiv preprint arXiv:2012.15466*.
- Jiaming Xu, Bo Xu, Peng Wang, Suncong Zheng, Guanhua Tian, and Jun Zhao. 2017. Self-taught convolutional neural networks for short text clustering. *Neural Networks*, 88:22–31.
- Yuanmeng Yan, Rumei Li, Sirui Wang, Fuzheng Zhang, Wei Wu, and Weiran Xu. 2021. Consent: A contrastive framework for self-supervised sentence representation transfer. *arXiv preprint arXiv:2105.11741*.
- Ainur Yessenalina and Claire Cardie. 2011. Compositional matrix-space models for sentiment analysis.

In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 172–182.

Jianhua Yin and Jianyong Wang. 2016. A model-based approach for text clustering with outlier detection. In *2016 IEEE 32nd International Conference on Data Engineering (ICDE)*, pages 625–636. IEEE.

Dejiao Zhang, Shang-Wen Li, Wei Xiao, Henghui Zhu, Ramesh Nallapati, Andrew O Arnold, and Bing Xiang. 2021. Pairwise supervised contrastive learning of sentence representations. *arXiv preprint arXiv:2109.05424*.

Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. 2017. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*.

Jian-Guo Zhang, Kazuma Hashimoto, Chien-Sheng Wu, Yao Wan, Philip S Yu, Richard Socher, and Caiming Xiong. 2019. Find or classify? dual strategy for slot-value predictions on multi-domain dialog state tracking. *arXiv preprint arXiv:1910.03544*.

Xiang Zhang and Yann LeCun. 2015. Text understanding from scratch. *arXiv preprint arXiv:1502.01710*.

Yan Zhang, Ruidan He, Zuozhu Liu, Kwan Hui Lim, and Lidong Bing. 2020. An unsupervised sentence embedding method by mutual information maximization. *arXiv preprint arXiv:2009.12061*.

A Implementation

Same as the original SimCSE work (Gao et al., 2021), we adopted 10^6 randomly sampled sentences from English Wikipedia as training data.⁷

We implement our models with Pytorch (Paszke et al., 2017). We use the pre-trained RoBERTa models as the backbone. We choose a two-layer MLP with size $(d \times d, d \times 128)$ to optimize our contrastive learning losses, where d denotes the dimension of the sentence representations. We use Adam (Kingma and Ba, 2015) as our optimizer with a constant learning rate of $5e-04$, which we scale to $5e-06$ for updating the backbones/transformers. We set the virtual augmentation strength of VaSCL, *i.e.*, Δ in Equation (3), to 15 for both DistilRoBERTa and RoBERTaBase, and 30 for RoBERTaLarge.

We train SimCSE (Gao et al., 2021) using $3e-05$ for optimizing the contrastive learning head and the backbone. We also tried the default learning rate $1e-05$ (suggested in Gao et al. (2021)) as well as our learning rate setup for optimizing the RoBERTa

⁷We download the training data via https://github.com/princeton-nlp/SimCSE/blob/main/data/download_wiki.sh.

models with SimCSE. We found $3e-05$ yields better performance. For both SimCSE and VaSCL, we set the batch size to 1024, train all models over five epochs and evaluate the development set of STS-B every 500 iterations. We report all our evaluations on the downstream tasks with the associated checkpoints attaining the best performance on the validation set of STS-B.

B Dataset Statistics

B.1 Intent Classification Dataset

We evaluate our model on four intent classification datasets: (1) **SNIPS** (Coucke et al., 2018) is a SLU benchmark that consists of 7 distinct intents. (2) **BANKING77** (Casanueva et al., 2020) is a large fine-grained single banking domain intent dataset with 77 intent classes. (3) **HWU64** (Liu et al., 2021) contains 25,716 examples for 64 intents in 21 domains. (4) **CLINC150** (Larson et al., 2019) spans 150 intents and 23,700 examples across 10 domains. As we can see here, SNIPS are limited to only a small number of classes, which oversimplifies the intent detection task and does not emulate the true environment of commercial systems. The remaining three datasets contain much more diversity and are more challenging.


B.2 Short Text Clustering Dataset

Dataset	N	\bar{W}	C	ImN
AgNews	8.0K	23	4	1
SearchSnippets	12.3K	18	8	7
StackOverflow	20K	8	20	1
Biomedical	20K	13	20	1
GoogleNews	11.1K	28	152	143
Tweet	2.5K	8	89	249

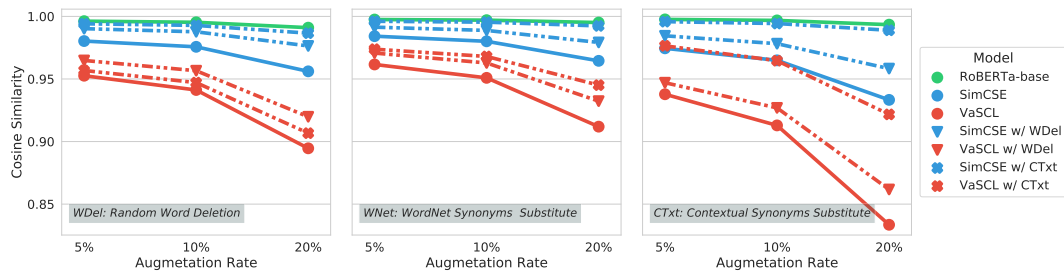
Table 4: Statistics of six short text clustering datasets. N: number of text samples; \bar{W} : average number of words each text example has; C: number of clusters; ImN: imbalance number defined as the size of the largest class divided by that of the smallest class.

- **SearchSnippets** is extracted from web search snippets, which contains 12340 snippets associated with 8 groups Phan et al. (2008).
- **StackOverflow** is a subset of the challenge data published by Kaggle⁸, where 20000 question titles associated with 20 different categories are selected by Xu et al. (2017).

⁸<https://www.kaggle.com/c/predict-closed-questions-on-stackoverflow/download/train.zip>



(a) Evaluating VaSCL in presence of different explicit data augmentation strategies.



(b) Cosine similarity between the representations of each original training example and its augmentation evaluated on different models. From left to right, the augmentations are obtained via *WDel*, *WNet*, and *CTxt*. Each point is averaged over 20,000 randomly sampled training examples.

Figure 4: Comparing and combining virtual augmentation with explicit text augmentations. (Full plot of Figure 3 in Section 4.4.)

- **Biomedical** is a subset of PubMed data distributed by BioASQ⁹, where 20000 paper titles from 20 groups are randomly selected by Xu et al. (2017).
- **AgNews** is a subset of news titles (Zhang and LeCun, 2015), which contains 4 topics selected by Rakib et al. (2020).
- **Tweet** consists of 89 categories with 2472 tweets in total (Yin and Wang, 2016).
- **GoogleNews** contains titles and snippets of 11109 news articles related to 152 events (Yin and Wang, 2016).

C Full Evaluation of Intent Classification

ATIS (Hemphill et al., 1990) is a benchmark for the air travel domain. This dataset is highly imbalanced, with the largest class containing 73% of all the training and validation examples. Moreover, more than 60% classes have less than 20 examples. We thereby exclude this task in our evaluation.

⁹<http://participants-area.bioasq.org>