

Challenging Assumptions in Learning Generic Text Style Embeddings

Phil Ostheimer, Marius Kloft, Sophie Fellenz
RPTU Kaiserslautern-Landau
surname@cs.uni-kl.de

Abstract

Recent advancements in language representation learning primarily emphasize language modeling for deriving meaningful representations, often neglecting style-specific considerations. This study addresses this gap by creating generic, sentence-level style embeddings crucial for style-centric tasks. Our approach is grounded on the premise that low-level text style changes can compose any high-level style. We hypothesize that applying this concept to representation learning enables the development of versatile text style embeddings. By fine-tuning a general-purpose text encoder using contrastive learning and standard cross-entropy loss, we aim to capture these low-level style shifts, anticipating that they offer insights applicable to high-level text styles. The outcomes prompt us to reconsider the underlying assumptions as the results do not always show that the learned style representations capture high-level text styles.

1 Introduction

Learning meaningful representations of text has received much attention recently as models pre-trained on large corpora have become the standard for extracting representations that capture prior knowledge (i.a. Devlin et al., 2019; Radford and Narasimhan, 2018). However, most pre-trained models focus on general aspects, as they are trained with causal or masked language modeling objectives. Therefore, they are suboptimal for NLP tasks focusing on a specific aspect, such as style. There are many notable NLP tasks, such as identifying chatbot-written text based on its style (e.g., Soto et al. (2024)) or style transfer where models rely on dedicated style representations (e.g., John et al. (2019)).

Existing work on style transfer focuses on high-level stylistic aspects such as sentiment (Shen et al., 2017) or formality (Rao and Tetreault, 2018), taking style definitions based on the dataset’s structure.

Lyu et al. (2021) view those high-level text style changes as compositions of more fine-grained, low-level style changes. Generic style representations incorporating low-level aspects such as lexical, syntactic, semantic, and thematic stylistic traits (McDonald and Pustejovsky, 1985; DiMarco and Hirst, 1993) but also high-level, composed stylistic traits might significantly improve style-focused tasks.

Based on the hypothesis that low-level style changes compose high-level style changes, this work explores learning generic, sentence-level style representations. We take a pre-trained encoder model producing general-purpose representations and fine-tune it to distinguish between low-level stylistic changes using contrastive learning and cross-entropy loss. We hypothesize that the resulting text encoder generalizes to high-level styles, applying the view by Lyu et al. (2021) on how low-level stylistic changes compose high-level changes to representation learning for styles.

We train our method using contrastive learning and cross-entropy loss on the StylePTB dataset (Lyu et al., 2021) comprising low-level, fine-grained style changes to obtain high-level, generic style embeddings. We evaluate our method by training a simple classifier on the representations of the learned style. The results show an ambiguous picture of the resulting style embeddings, challenging the underlying assumptions.

2 Related Work

Explicitly Learning Style Representations

Only a few works learn style representations explicitly. StyleDistance (Patel et al., 2024) uses a contrastive triplet loss and synthetic parallel data created by Large Language Models to learn generic style representations, showing strong performance in multiple benchmarks.

Style Representations as a Byproduct Text style transfer is conducted by prominent models by

disentangling content and style to learn separate task-specific representations to control them independently (Fu et al., 2018; Hu et al., 2017; Kim and Sohn, 2020; John et al., 2019; Cheng et al., 2020). Other text style transfer models learn content representations and multiple decoders (one for each style) (Shen et al., 2017; Fu et al., 2018). Another group of models uses a structured style code to enforce a particular style in the decoder, either given as a structured code (Hu et al., 2017; Lample et al., 2019) or learned (Fu et al., 2018; Kim and Sohn, 2020).

For text style classification, TextCNN (Kim, 2014) is the most widely used method (Ostheimer et al., 2023). BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), fine-tuned for style classification, are strong baselines.

In contrast, our proposed method learns style representations at the sentence level from low-level, linguistically motivated style changes. This is much more fine-grained and allows applications to unseen styles.

Contrastive Learning for Text Representations

To learn meaningful content representations on the sentence level from unlabeled text corpora, the QT model (Logeswaran and Lee, 2018) was introduced. The QT model relies on the distributional hypothesis to get meaningful content representations. It uses a contrastive objective to map nearby (context) sentences to similar and distant (non-context) sentences to far-apart representations.

For fine-tuning sentence representations, notable approaches are SimCSE (Gao et al., 2021) and Mirror-BERT (Liu et al., 2021), incorporating minimal data augmentation with dropout. In contrast, SBERT (Reimers and Gurevych, 2019) uses siamese and triplet network structures to generate meaningful sentence representations for calculating semantic sentence similarities using standard measures like cosine distance. Kim et al. (2021) improve the quality of the sentence representations by contrasting the representations of different layers of BERT (Devlin et al., 2019).

However, these approaches focus on improving the sentence-level representations for general language understanding or semantics. We, in contrast, focus on the style of the sentences.

3 Method

This section describes the underlying assumptions of the proposed method to compute (sentence-level)

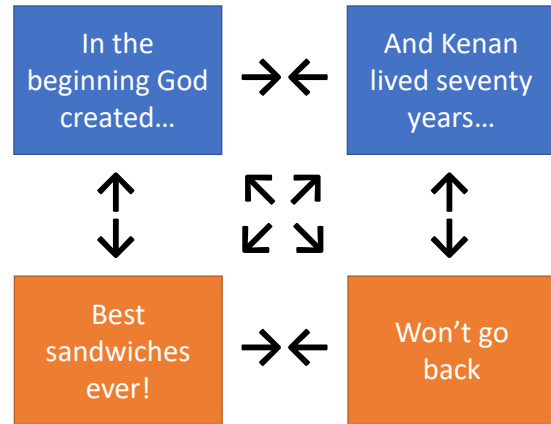


Figure 1: Our training objective pushes sentence representations of the same style close together. In this example, reviews (in orange) are pushed close together, and sentences of one Bible version (in blue) are pushed close together, while the representations of different styles (Bible vs reviews) are pushed to be far apart.

style representations of texts and the method itself.

Assumption 1 *Low-level stylistic changes can be composed to form high-level styles in line with previous work (Lyu et al., 2021).*

Assumption 2 *Learning style embeddings can be achieved by contrasting different styles (Patel et al., 2024).*

We combine the aforementioned assumptions to come up with the following method.

Method The core idea—illustrated in Figure 1—is to embed the data into a space where texts of the same style resemble each other, while texts of different styles are easily distinguishable. Formally, we want to learn a neural encoder f mapping any two input sentences x and \tilde{x} onto vector representations $f(x)$ and $f(\tilde{x})$ such that $f(x) \approx f(\tilde{x})$ if and only if x and \tilde{x} are of the same style: $\text{style}(x) = \text{style}(\tilde{x}) = s$. To achieve this, we first form pairs of sentences, some equal and others of a different style. Then, we use a contrastive objective to push the vector representations of sentences of the same style closer together and ones of different styles far apart.

For a given sentence x with $\text{style}(x) = s$, our set of candidate sentences X_{cand} contains sentences of the same style and different style(s): $X_{\text{cand}} = X_s \cup \overline{X_s}$. We employ a simple architecture to compute the similarity between the outputs $f(x)$ and $f(\tilde{x})$, namely the inner product $f(x)^\top f(\tilde{x})$. Other functions to consider are, e.g., the cosine similarity or kernel functions. We compute the similarity

using a simple architecture to avoid learning a rich similarity measure compensating for the encoder producing poor style representations.

Therefore, we have the following formulation for the probability of each candidate sentence $\tilde{x} \in X_{\text{cand}}$ to have the same style as x :

$$p(\tilde{x}|x, X_{\text{cand}}) = \frac{\exp[f(x)^\top f(\tilde{x})]}{\sum_{\tilde{x}' \in X_{\text{cand}}} \exp[f(x)^\top f(\tilde{x}')] } \quad (1)$$

Our training objective is to maximize the probability of identifying all sentences $\tilde{x} \in X_{\text{cand}}$ where $\text{style}(x) = \text{style}(\tilde{x}) = s$ for each sentence x in the training data D :

$$\sum_{x \in D} \sum_{\tilde{x} \in X_s} \log p(\tilde{x}|x, X_{\text{cand}}) \quad (2)$$

We also experiment with a cross-entropy loss and its combination with a contrastive loss.

4 Experiments

In this section, we describe how we evaluated the effectiveness in learning style representations for multiple styles.

4.1 Experimental Setup

Fine-Tuning Data We use the StylePTB (Lyu et al., 2021) dataset containing 21 individual and 32 compositional fine-grained style changes to train our generic text style embeddings.

Evaluation Data We evaluate the learned style embeddings on datasets that take a data-driven approach, containing high-level stylistic changes, to define a text style. These are used in many recent works on text style transfer. We experiment with the Bible corpus (Carlson et al., 2018) with eight different styles of the Bible (249K sentences), Grammarly’s Yahoo Answers Formality Corpus (GYAFC) (Rao and Tetreault, 2018) (113K sentences) in two styles, a collection of Shakespeare plays in Shakespearean and modern English (Xu et al., 2012) (42K sentences), Amazon (558K) and Yelp (448K) sentiment datasets¹ (two styles each). We follow prior work by using the existing train-dev-test splits.

Training We experiment with both RoBERTa_{Large} (Liu et al., 2019) and BERT_{Large} (Devlin et al., 2019) as pre-trained encoders for

f . The sentence representation is the activation from the last hidden layer for the “CLS” token. We add a linear transformation and a l_2 normalization before applying the objective function. Training stops after a maximum of 10 epochs, and the best model is used based on the loss on the validation dataset. Hyperparameters were chosen using the loss on the validation dataset. We used a batch size of 16, a learning rate of 1e-5 with a linear warmup for the first 10% steps, followed by a linear cooldown for the remaining steps. We used an Adam optimizer and a dropout rate of 0.1. A logistic classifier is trained on the training data and evaluated on the test data using the representations obtained by the encoder f .

Batch Construction We initially experimented with contrasting more than two styles per batch. However, this resulted in no meaningful representations. Therefore, we resorted to two styles per batch. We use a random sampler with replacements to randomly select sentences of each style. We also use two random samplers without replacement: one where we assure for each style pair that only unseen sentence pairs are contrasted and one where we only assure on the corpus level that per epoch, the sentences are only contrasted once. Half of the batch consists of style s while the other half contains style \bar{s} .

Obtaining Generic Style Embeddings To get generic style embeddings, we fine-tune pre-trained encoders f on StylePTB (Lyu et al., 2021) and apply the resulting encoders to the previously mentioned high-level style datasets. The hypothesis is the following: Since StylePTB contains fine-grained (and compositional) style changes, it should also generalize to high-level and unseen styles.

4.2 Results

In Table 1, we summarize our style classification results. Using pre-trained BERT and RoBERTa encoders as f , we apply a logistic regression for classification without fine-tuning, serving as our baseline. Baseline accuracies for datasets like Yelp, Amazon, GYAFC, and Shakespeare (two styles each) are already nearly 80% or higher. The Bible dataset, with eight styles, exhibits lower accuracy, as expected. Generally, BERT outperforms RoBERTa in style classification. Our approach considers three sampling strategies (Section 4.1), crucial for contrastive loss performance.

¹<https://github.com/lijuncen/Sentiment-and-Style-Transfer>

	Dataset	PT	CEL	Random Sampler		Pairwise Sampler		Corpus Sampler	
				CEL+CL	CL	CEL+CL	CL	CEL+CL	CL
BERT	Yelp	94.3	92.9	90.5	90.1	72.9	87.0	71.4	88.1
	Amazon	77.4	74.9	72.1	74.5	59.6	68.3	58.1	69.3
	GYAFC	88.4	88.0	82.1	82.1	71.5	86.2	74.3	78.3
	Shakespeare	85.7	86.4	83.4	83.5	75.1	83.7	73.7	81.7
	Bible	50.9	52.8	42.9	46.7	31.1	52.4	33.9	41.2
RoBERTa	Yelp	86.2	87.5	85.3	74.2	55.1	50.8	56.0	71.9
	Amazon	75.2	67.1	64.4	60.3	54.1	65.4	55.9	58.1
	GYAFC	79.4	85.5	80.2	76.3	55.5	59.0	60.4	66.9
	Shakespeare	80.1	85.1	81.7	76.1	64.8	74.0	67.4	77.6
	Bible	55.9	60.8	53.8	44.4	18.9	24.2	19.1	38.5

Table 1: Shown is the style classification accuracy of a logistic regression fitted to the two text encoders f BERT and RoBERTa as a pre-trained (PT) encoder, fine-tuned on StylePTB and applied to the mentioned datasets using a Cross-Entropy Loss (CEL), Contrastive Loss (CL), or both (CEL+CL) with the three mentioned sampling strategies.

Fine-tuning BERT and RoBERTa with the cross-entropy loss yields slight accuracy improvements for the Shakespeare and Bible corpora. Moreover, fine-tuning RoBERTa further enhances accuracies on Yelp and GYAFC. However, incorporating the contrastive loss reduces accuracy compared to cross-entropy fine-tuning. Solely fine-tuning with a contrastive loss also leads to less accuracy.

4.3 Discussion

While fine-tuning an encoder f on the StylePTB dataset using cross-entropy loss slightly improves some settings’ accuracy, contrastive learning does not. These findings of our study question the representational capacity of the learned style embeddings, especially for contrastive learning.

Contrastive Objective Is Too Aggressive Contrary to our expectations, applying the contrastive loss does not improve the accuracy compared to the cross-entropy loss across various settings we explored. One possible reason is that the contrastive loss might push dissimilar styles too far apart. Contrary to previous work (Patel et al., 2024), we do not use synthetic parallel data but contrast non-parallel data. This might hamper the model’s ability to learn the relationship between different styles and style levels.

RoBERTa’s “CLS” Tokens Need Fine-Tuning

The improvements with the cross-entropy loss can be attributed to the fact that the “CLS” token is not pre-trained using the next sentence prediction task and, therefore, any fine-tuning might improve the “CLS” token representations.

Differentiated Picture for Cross-Entropy Loss

One possible reason for the differentiated results can be seen in the examples from StylePTB in Table 2. While some low-level changes, like info addition, relate directly to a formality change, others, such as tense changes, do not align with the investigated styles. Mixing low-level styles may confuse the encoding mechanisms that distinguish higher-level styles.

Aspect	Original	Transferred
Info addition	Morgan Freeman did the new one	Morgan Freeman did <i>perform</i> the new one.
To future tense	It is also planning another night of original series.	It will be also planning another night of original series.

Table 2: Examples from the StylePTB dataset

5 Conclusion

Learning generic, high-level text style representations from low-level, linguistically motivated changes and generalizing to high-level styles according to Assumption 1 by Lyu et al. (2021) using contrastive learning according to Assumption 2 by Patel et al. (2024) presents challenges. Although the approach does not yield the expected results using contrastive learning, cross-entropy loss shows improvements for some settings. However, our approach does not yield the expected results compared to previous work (Patel et al., 2024) using

contrastive learning to learn generic style representations.

Limitations

One limitation of this study is the reliance on the StylePTB dataset, which, to our knowledge, is the only available dataset containing low-level and composed stylistic changes. The dataset is restricted to English, limiting the generalizability of our findings to other languages. As style may manifest differently across languages, a more diverse, multilingual dataset would allow for broader application and a more comprehensive evaluation of the proposed method.

Additionally, our study focuses on contrastive learning to capture text style representations. This choice was made because it intuitively aligns with the underlying assumptions of the task, but it may not be the optimal approach for all settings.

Acknowledgements

Part of this work was conducted within the DFG Research Unit FOR 5359 on Deep Learning on Sparse Chemical Process Data (BU 4042/2-1, KL 2698/6-1, and KL 2698/7-1) and the DFG TRR 375 (no. 511263698). SF and MK acknowledge support by the Carl-Zeiss Foundation (AI-Care, Process Engineering 4.0), the BMBF award 01IS2407A, and the DFG awards FE 2282/6-1, FE 2282/1-2, KL 2698/5-2, and KL 2698/11-1.

References

- Keith Carlson, Allen Riddell, and Daniel Rockmore. 2018. [Evaluating prose style transfer with the Bible](#). *Royal Society Open Science*, 5(10):171920.
- Pengyu Cheng, Martin Renqiang Min, Dinghan Shen, Christopher Malon, Yizhe Zhang, Yitong Li, and Lawrence Carin. 2020. [Improving disentangled text representation learning with information-theoretic guidance](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7530–7541. 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Chrysanne DiMarco and Graeme Hirst. 1993. A Computational Theory of Goal-Directed Style in Syntax. *Computational Linguistics*, 19(3):50.
- Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2018. [Style transfer in text: Exploration and evaluation](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 663–670. AAAI Press.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple Contrastive Learning of Sentence Embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*. Association for Computational Linguistics.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. [Toward controlled generation of text](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1587–1596. PMLR.
- Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. 2019. [Disentangled representation learning for non-parallel text style transfer](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 424–434. Association for Computational Linguistics.
- Heejin Kim and Kyung-Ah Sohn. 2020. [How positive are you: Text style transfer using adaptive style embedding](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 2115–2125. International Committee on Computational Linguistics.
- Taeuk Kim, Kang Min Yoo, and Sang-goo Lee. 2021. [Self-guided contrastive learning for BERT sentence representations](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 2528–2540. Association for Computational Linguistics.
- Yoon Kim. 2014. [Convolutional Neural Networks for Sentence Classification](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.
- Guillaume Lample, Sandeep Subramanian, Eric Michael Smith, Ludovic Denoyer, Marc’Aurelio

- Ranzato, and Y-Lan Boureau. 2019. [Multiple-attribute text rewriting](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Fangyu Liu, Ivan Vulić, Anna Korhonen, and Nigel Collier. 2021. [Fast, Effective, and Self-Supervised: Transforming Masked Language Models into Universal Lexical and Sentence Encoders](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1442–1459, Online and Punta Cana, Dominican Republic. 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](#). *CoRR*, abs/1907.11692.
- Lajanugen Logeswaran and Honglak Lee. 2018. [An efficient framework for learning sentence representations](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Yiwei Lyu, Paul Pu Liang, Hai Pham, Eduard H. Hovy, Barnabás Póczos, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2021. [Styleptb: A compositional benchmark for fine-grained controllable text style transfer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 2116–2138. Association for Computational Linguistics.
- David D McDonald and James Pustejovsky. 1985. A computational theory of prose style for natural language generation. In *Second Conference of the European Chapter of the Association for Computational Linguistics*.
- Phil Ostheimer, Mayank Kumar Nagda, Marius Kloft, and Sophie Fellenz. 2023. [A call for standardization and validation of text style transfer evaluation](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 10791–10815. Association for Computational Linguistics.
- Ajay Patel, Jiacheng Zhu, Justin Qiu, Zachary Horvitz, Marianna Apidianaki, Kathleen McKeown, and Chris Callison-Burch. 2024. [Styledistance: Stronger content-independent style embeddings with synthetic parallel examples](#). *arXiv preprint arXiv:2410.12757*.
- Alec Radford and Karthik Narasimhan. 2018. Improving language understanding by generative pre-training.
- Sudha Rao and Joel R. Tetreault. 2018. [Dear sir or madam, may I introduce the GYAFC dataset: Corpus, benchmarks and metrics for formality style transfer](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 129–140. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3980–3990. Association for Computational Linguistics.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi S. Jaakkola. 2017. [Style transfer from non-parallel text by cross-alignment](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6830–6841.
- Rafael Alberto Rivera Soto, Kailin Koch, Aleem Khan, Barry Y. Chen, Marcus Bishop, and Nicholas Andrews. 2024. [Few-shot detection of machine-generated text using style representations](#). In *The Twelfth International Conference on Learning Representations*.
- Wei Xu, Alan Ritter, Bill Dolan, Ralph Grishman, and Colin Cherry. 2012. [Paraphrasing for style](#). In *COLING 2012, 24th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, 8-15 December 2012, Mumbai, India*, pages 2899–2914. Indian Institute of Technology Bombay.