

Self-Vocabularizing Training for Neural Machine Translation

Pin-Jie Lin[♣] Ernie Chang[♦] Yangyang Shi[♦] Vikas Chandra[♦]

[♣]Virginia Tech

[♦]Meta

pinjie@vt.edu, erniechyc@meta.com

Abstract

Past vocabulary learning techniques identify relevant vocabulary before training, relying on statistical and entropy-based assumptions that largely neglect the role of model training. Empirically, we observe that trained translation models are induced to use a byte-pair encoding (BPE) vocabulary subset distinct from the original BPE vocabulary, leading to performance improvements when retrained with the induced vocabulary. In this paper, we analyze this discrepancy in neural machine translation by examining vocabulary and entropy shifts during self-training—where each iteration generates a labeled dataset by pairing source sentences with the model’s predictions to define a new vocabulary. Building on these insights, we propose *self-vocabularizing training*, an iterative method that self-selects a smaller, more optimal vocabulary, yielding up to a 1.49 BLEU improvement. Moreover, we find that deeper model architectures lead to both an increase in unique token usage and a 6–8% reduction in vocabulary size.

1 Introduction

Vocabulary construction, also known as vocabularization, is essential for many natural language processing tasks that involve neural networks, including neural machine translation (MT), as highlighted in various studies (Mikolov et al., 2013; Vaswani et al., 2017; Gehrmann et al., 2018; Zhang et al., 2018; Devlin et al., 2019). However, past vocabulary learning techniques rely on corpus statistics such as entropy (Xu et al., 2020) or frequency counts (Sennrich et al., 2016), without considering contextual information or the model’s ability to represent it.

Despite the success of vocabularization in improving MT model efficiency (Xu et al., 2020), we observe a discrepancy between the original byte-pair encoding (BPE) vocabulary (Gage, 1994) (V_0), derived from the initial training data, and the BPE

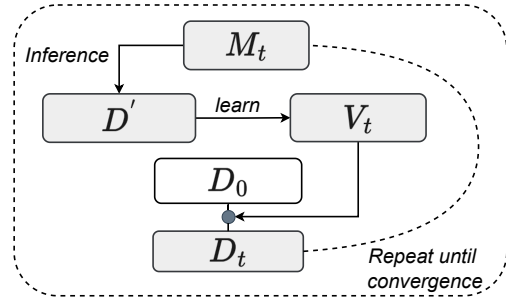


Figure 1: Illustration of self-vocabularizing training: At each iteration, the original dataset D_0 is segmented using vocabulary V_t to form the training set D_t . D_t is then used to train model M_t , which generates a pseudo dataset D' . A new vocabulary set V_{t+1} is derived from D' , completing the training loop. This process repeats until no further improvements are observed.

vocabulary induced from pseudo-labeled data (V_1) (see Figure 1). This discrepancy is surprising, as it suggests that MT models implicitly learn a pseudo-“optimal” vocabulary (V_1) that is substantially different from the original vocabulary (V_0) and is also smaller in size. For instance, on the IWSLT14 DE-EN dataset, $|V_1|$ is approximately 20% smaller than $|V_0|$. Moreover, MT models retrained with the pseudo vocabulary V_1 outperform those trained on the original vocabulary set. This suggests a need to re-examine the assumptions underlying vocabulary learning techniques such as byte-pair encoding and the marginal utility of vocabularization (Xu et al., 2020), as existing methods may overlook model-data interactions, leaving key optimization factors unaccounted for.

In this paper, we aim to understand this discrepancy in neural machine translation models by analyzing shifts in vocabulary and entropy during self-training. To this end, we conduct experiments on two language tasks, comparing the vocabulary sets learned from the original training data and the pseudo-labeled data. Our results suggest limited overlap between the two vocabularies and that pseudo data induces a more optimal vocabulary,

enabling further improvements. Furthermore, our findings indicate that the decoder has a limited impact on this vocabulary shift, whereas encoder-based interactions play a crucial role in entropy reduction. This suggests that future vocabulary induction methods should focus more on the cross-attention module. Finally, our study has implications for defining an optimal vocabulary set in language generation.

This preliminary study also introduces a simple yet effective technique: iterative self-training to self-select a more optimal vocabulary set for performance gains.

In summary, this paper makes the following contributions: (1) We identify a discrepancy between the optimal and pseudo-labeled vocabulary derived from MT models. (2) We analyze shifts in vocabulary and entropy during self-training. (3) We propose a simple approach to obtain a competitive vocabulary set and introduce a self-vocabularizing training algorithm that improves performance.

2 Iterative Self-Vocabularization

Current vocabularization techniques adopt two contrasting perspectives: (1) Focusing on frequency or entropy statistics to avoid the computational cost of trial training (Xu et al., 2020), which often neglects important parameters and interactions in the process. (2) Obtaining a more optimal vocabulary set through training (Salesky et al., 2020), but at a higher computational cost. This work combines both perspectives by adopting an entropy-based vocabularization approach while utilizing self-training.

In self-training, a base model M_t is trained on the dataset to generate predictions for input sequences, which are then used to update the next iteration of the base model M_{t+1} . This process is repeated iteratively with the supervised loss \mathcal{L} from labeled instances (He et al., 2019), where x and y are the source and target texts, respectively:

$$\mathcal{L} = -\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} \mathbb{E}_{\mathbf{y} \sim p_{\theta^*}(\mathbf{y}|\mathbf{x})} \log p_{\theta}(\mathbf{y}|\mathbf{x}), \quad (1)$$

where $p(\mathbf{x})$ is the empirical data distribution approximated with samples from D , and $p_{\theta}(\mathbf{y}|\mathbf{x})$ is the conditional distribution defined by the model. The parameter θ^* is randomly initialized at every iteration. For each training iteration at $t + 1$, we relearn the BPE vocabulary using the *original source* and the *pseudo target* generated with M_t

(i.e., $D' = \{(\mathbf{x}, f_{\theta}(\mathbf{x})) | \mathbf{x} \in U\}$). Then, M_{t+1} is trained on D_{t+1} , segmented with the newly derived vocabulary V_{t+1} .

Measuring Vocabulary Shifts Subword-based approaches like byte-pair encoding are widely used and have demonstrated strong empirical performance (Sennrich et al., 2016; Al-Rfou et al., 2019; Costa-jussà and Fonollosa, 2016; Lee et al., 2017; Ding et al., 2019; Liu et al., 2020; Ott et al., 2018; Kudo and Richardson, 2018; Wang et al., 2020). These methods construct vocabulary by selecting high-probability subword units.

Following Xu et al. (2020), we define the vocabulary shift as the negative change in entropy normalized by vocabulary size:

$$\frac{-(\mathcal{H}_{M_{t+1}(\mathbf{x})} - \mathcal{H}_{M_t(\mathbf{x})})}{|V_t|}, \quad (2)$$

where $M_{t+1}(\mathbf{x}) \rightarrow V_{t+1}$ and $M_t(\mathbf{x}) \rightarrow V_t$ represent vocabularies from two consecutive training iterations, with sizes $|V_{t+1}|$ and $|V_t|$. The ratio $|V_{t+1}|/|V_t|$ reflects compression in vocabulary size.

Corpus entropy \mathcal{H}_v with vocabulary V is defined as the sum of token entropy, normalized by the average token length:

$$\mathcal{H}_v = -\frac{1}{l_v} \sum_{j \in V} p(j) \log p(j), \quad (3)$$

where $p(j)$ represents the relative frequency of token j in the training corpus, and l_v is the average token length (i.e., the number of characters per token).

3 Experimental Settings

For our experiments, we used the IWSLT14 German-English parallel corpus for both German-to-English (DE-EN) and English-to-German (EN-DE) translation tasks. We preprocessed the data using MOSES (Bollmann et al., 2021) and applied byte-pair encoding (BPE) (Sennrich et al., 2016) to construct the vocabulary set.

We trained a transformer-based NMT model using the fairseq library (Ott et al., 2019), with six layers, four attention heads, and a hidden size of 1024 dimensions. The Adam optimizer was used with a learning rate of 0.0002 and a batch size of 64. Training lasted for 50 epochs, with exponential learning rate decay and early stopping based on the validation set.




Figure 2: Entropy and performance across self-vocabularizing training iterations. **(Left)** BLEU score (blue) consistently improves across iterations. Meanwhile, the self-learned vocabulary reduces corpus entropy (teal), indicating a better estimation of token distribution. **(Right)** Vocabulary shift measured by vocabulary overlap (orange) between consecutive vocabularies V_t and V_{t-1} , showing that the model initially selects a broad set of subwords before consolidating onto a subset of V_{t-1} ¹. The type-token ratio (TTR) (purple) reflects the diversity of learned semantic units, reported on the training corpus scaled by 1000.

We evaluated model performance using BLEU scores on the test set, comparing against the baseline and other comparable models. For self-training, we ran each iteration until performance converged. In each iteration i , the model was trained from scratch in the self-training step ($ST-i$) for analysis purposes. Results were averaged over three initialization runs with different θ^* .

3.1 Main Results

We first compared the performance of MT models trained with two different approaches: one using a fixed output vocabulary and the other refining the output vocabulary through self-training iterations. As shown in Table 1 and Figure 3, the MT model trained with the self-trained output vocabulary gradually improves with each newly derived vocabulary, achieving up to a 1.3 BLEU point increase after a single iteration. Table 1 further confirms a consistent trend across both language tasks: self-training improves model performance and reduces vocabulary entropy, leading to enhanced fluency and correctness while decreasing vocabulary size.

Beyond translation quality, we also observe lower overall corpus entropy and a smaller vocabulary in the self-trained model (see Figure 2). This suggests that self-training not only enhances translation accuracy but also results in a more efficient model with a compact, more targeted vocabulary—potentially enabling faster and more memory-efficient deployment.

¹Vocabulary overlaps at first iteration leverages the identical vocabulary where $V_0 = V_1$.

	BLEU	IVI	Overlap (%)	Fluency	Adequacy
ST-0	34.62	10000	-	2.89	3.21
ST-1	35.92	8950	66.42	3.13	3.46
ST-5	36.01	8892	88.27	3.42	3.87
ST-10	36.11	8702	96.19	3.95	4.21

Table 1: Performance comparison of BLEU, vocabulary size ($|V|$), vocabulary overlap (%), fluency, and adequacy on the IWSLT14 DE-EN translation task for **ST-0**, **ST-1**, **ST-5**, and **ST-10** models. Fluency and adequacy scores are segment-level averages on 100 random outputs, rated on a 1-5 scale (5 being the most fluent or correct) (Koehn and Monz, 2006; Freitag et al., 2021). Scores were assigned by three raters and then averaged. Detailed results for IWSLT14 EN-DE are provided in Appendix A.1.

4 Ablations of Self-Vocabularization

4.1 Shifts Across Iterations

In text generation, self-training can enhance the quality of the generated output. However, the impact of the number of self-training iterations on output entropy (i.e., the randomness or unpredictability of the generated text) is not straightforward and depends on the specifics of the model and training data. We therefore examine: (1) corpus entropy and (2) subword-based overlap between the original and self-trained BPE vocabulary.

Corpus Entropy. Increasing the number of self-training iterations allows the model to learn from a progressively smaller set of labeled examples, potentially leading to more coherent and accurate outputs with reduced diversity. In the left plot of Figure 3, we observe that, in general, as entropy gradually decreases, self-training performance improves until the rate of change in both entropy and BLEU slows. Surprisingly, even at the 10th iteration, the model continues to improve its BLEU




Figure 3: Performance and vocabulary overlap across models with different encoder and decoder depths. **(Left)** As the number of encoder (teal -) or decoder (teal - -) layers increases, BLEU scores consistently improve. However, vocabulary overlap decreases for deeper encoder (blue -) or decoder (blue - -) layers, indicating that deeper models tend to use more unique tokens. **(Right)** Vocabulary compression (VC) across models with varying depths. All models trained with self-vocabularizing training effectively compress the token set. Notably, deeper encoder models (purple) exhibit a smoother reduction in VC rates, whereas deeper decoder models (orange) require more tokens for inference. VC is reported on the test set using models of different depths in either the encoder or decoder, with a single round of self-vocabularizing training.

score.

Vocabulary Overlap ($V_O \cap V_P$). The original BPE vocabulary consists of subword units created by applying BPE to the training data, serving as a fixed vocabulary during training and inference. While the number of self-training iterations does not directly alter the BPE vocabulary (as it is predefined before training), fine-tuning on additional labeled examples can improve model performance, leading to more accurate and diverse outputs that better align with the original BPE vocabulary. Additionally, vocabulary size consistently decreases across iterations. We observe an initial sharp drop of approximately 10% after the first iteration, followed by a gradual reduction in BPE vocabulary size until the 5th iteration (see Figure 2).

4.2 Ablations on Model Architecture

Deeper Model Depth Contributes to Lower Vocabulary Overlap. The number of encoder and decoder layers in a neural network plays a crucial role in determining output coherence and accuracy, which in turn affects the model’s output token set. As shown in Figure 3, increasing encoder or decoder layers generally improves BLEU scores. Additionally, vocabulary overlap gradually decreases to approximately 93% as the number of layers increases, following a similar trend observed in Figure 2 at the first iteration. This suggests that deeper architectures allow the model to implicitly select more unique tokens compared to shallow models, with encoders playing a particularly important role

in vocabulary selection.

Vocabulary Compression. *Vocabulary compression (VC)* is defined as the ratio of the number of tokens used in the inference output to the number of tokens in the original test set, i.e., $\frac{|V|^{inf}}{|V|^{test}}$. Figure 3 illustrates the relationship between the number of encoder/decoder layers and VC. All models achieve significant token set compression, reducing vocabulary size by 6% to 8%. Notably, increasing encoder depth results in a smaller token set, whereas increasing decoder depth leads to a larger token set. We conjecture that deeper encoders have a stronger ability to process source sentences and represent them as fixed-length context vectors, enabling the decoder to use fewer subword units for translation.

5 Conclusions and Findings

In this paper, we investigated the discrepancy between the "optimal" vocabulary set identified prior to training a translation model and the vocabulary actually used by the trained model. We found that the trained model diverged from the original BPE vocabulary and that a single iteration of self-training was sufficient to generate a competitive vocabulary set. Additionally, we examined the relationship between the self-vocabularizing process and the encoder-decoder architecture, demonstrating that deeper models favor the selection of rarer tokens while reducing vocabulary size, whereas decoders have a lesser influence on vocabularization.

Limitations

While self-vocabularizing training is simple and provides significant improvements over baseline training, it remains time-consuming. Moreover, further analysis is needed to better understand vocabulary shifts and how to efficiently determine the optimal set without requiring costly training iterations. This analysis should include an examination of token types and subword granularity, such as how subword segmentation evolves across training iterations.

In addition, our findings have yet to be verified across multiple language pairs, leaving this as an avenue for future work. Overall, this study highlights the need to incorporate vocabulary relearning during self-training and suggests that new vocabulary construction techniques could bridge the gap between model training and text interactions.

Ethics Statement

Vocabularization with model training has significantly improved machine translation performance. To minimize potential negative impacts, we conduct our experiments on publicly available datasets commonly used in machine translation research. However, if this method is applied to sensitive data, such as medical records, privacy-preserving policies should be strictly considered.

Additionally, while deeper model architectures promote the use of unique tokens, they also increase computational demands. The potential environmental impact of large-scale model training should be carefully evaluated when scaling this approach.

References

- Rami Al-Rfou, Dokook Choe, Noah Constant, Mandy Guo, and Llion Jones. 2019. Character-level language modeling with deeper self-attention. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 3159–3166. AAAI Press.
- Marcel Bollmann, Rahul Aralikatte, Héctor Murrieta Bello, Daniel Hershcovich, Miryam de Lhoneux, and Anders Søgaard. 2021. [Moses and the character-based random babbling baseline: CoAStL at AmericasNLP 2021 shared task](#). In *Proceedings of the First Workshop on Natural Language Processing for Indigenous Languages of the Americas*, pages 248–254. Online. Association for Computational Linguistics.
- Marta R. Costa-jussà and José A. R. Fonollosa. 2016. Character-based neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 2: Short Papers*. The Association for Computer 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.
- Shuoyang Ding, Adithya Renduchintala, and Kevin Duh. 2019. A call for prudent choice of subword merge operations in neural machine translation. In *Proceedings of Machine Translation Summit XVII Volume 1: Research Track, MTSummit 2019, Dublin, Ireland, August 19-23, 2019*, pages 204–213. European Association for Machine Translation.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, errors, and context: A large-scale study of human evaluation for machine translation](#). *Transactions of the Association for Computational Linguistics*, 9:1460–1474.
- Philip Gage. 1994. A new algorithm for data compression. *C Users Journal*, 12(2):23–38.
- Sebastian Gehrmann, Yuntian Deng, and Alexander M. Rush. 2018. Bottom-up abstractive summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 4098–4109. Association for Computational Linguistics.
- Junxian He, Jiatao Gu, Jiajun Shen, and Marc’Aurelio Ranzato. 2019. Revisiting self-training for neural sequence generation. In *International Conference on Learning Representations*.
- Philipp Koehn and Christof Monz. 2006. [Manual and automatic evaluation of machine translation between European languages](#). In *Proceedings on the Workshop on Statistical Machine Translation*, pages 102–121, New York City. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP*

- 2018: *System Demonstrations, Brussels, Belgium, October 31 - November 4, 2018*, pages 66–71. Association for Computational Linguistics.
- Jason Lee, Kyunghyun Cho, and Thomas Hofmann. 2017. Fully character-level neural machine translation without explicit segmentation. *Transactions of the Association for Computational Linguistics*, 5:365–378.
- Xiaodong Liu, Kevin Duh, Liyuan Liu, and Jianfeng Gao. 2020. Very deep transformers for neural machine translation. *CoRR*, abs/2008.07772.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3111–3119.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. Scaling neural machine translation. In *Proceedings of the Third Conference on Machine Translation: Research Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pages 1–9. Association for Computational Linguistics.
- Elizabeth Salesky, Andrew Runge, Alex Coda, Jan Niehues, and Graham Neubig. 2020. Optimizing segmentation granularity for neural machine translation. *Machine Translation*, pages 1–19.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008.
- Changhan Wang, Kyunghyun Cho, and Jiatao Gu. 2020. Neural machine translation with byte-level subwords. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 9154–9160. AAAI Press.
- Jingjing Xu, Hao Zhou, Chun Gan, Zaixiang Zheng, and Lei Li. 2020. Vocabulary learning via optimal transport for neural machine translation. *arXiv preprint arXiv:2012.15671*.
- Lei Zhang, Shuai Wang, and Bing Liu. 2018. Deep learning for sentiment analysis: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(4).




Figure 4: Impact of self-vocabularizing training on IWSLT14 EN-DE. **(Left)** BLEU scores improve consistently across iterations, while corpus entropy decreases, indicating more stable and predictable token distributions. **(Right)** Vocabulary overlap reduces as the model gradually refines its subword selection, while the type-token ratio (TTR) reflects evolving semantic diversity.

A Detailed Results

A.1 IWSLT14 EN-DE

We present results for MT models trained with self-vocabularizing on the IWSLT14 EN-DE dataset over 10 iterations. We report performance for the baseline model (**ST-0**) and models trained with self-vocabularizing at iterations 1, 5, and 10. As shown in Table 2, all self-trained models outperform the fixed-vocabulary baseline (**ST-0**), with a 2.4-point increase in BLEU after 10 iterations. Additionally, vocabulary size decreases with each iteration, leading to a more compact vocabulary.

	BLEU	IVI	Overlap (%)
ST-0	28.64	10000	-
ST-1	29.63	8969	66.42
ST-5	29.66	8892	88.27
ST-10	30.14	8863	93.91

Table 2: BLEU scores, vocabulary size ($|V|$), and overlap (%) on the IWSLT14 EN-DE translation task for **ST-0**, **ST-1**, **ST-5**, and **ST-10** models.

Figure 4 illustrates the impact of self-vocabularizing training on corpus entropy, performance, vocabulary overlap, and diversity for the IWSLT14 EN-DE dataset. We observe a consistent decrease in corpus entropy and vocabulary overlap, alongside performance improvements with increasing training iterations. This confirms the effectiveness of self-vocabularizing training on the EN-DE translation task. Notably, the EN-DE translation exhibits lower diversity than the DE-EN task, which aligns with expectations since German shares more semantic units than English.