

Breaking the Representation Bottleneck of Chinese Characters: Neural Machine Translation with Stroke Sequence Modeling

Zhijun Wang Xuebo Liu* Min Zhang

Institute of Computing and Intelligence, Harbin Institute of Technology, Shenzhen, China
1190303311@stu.hit.edu.cn, {liuxuebo, zhangmin2021}@hit.edu.cn

Abstract

Existing research generally treats Chinese character as a minimum unit for representation. However, such Chinese character representation will suffer two bottlenecks: 1) Learning bottleneck, the learning cannot benefit from its rich internal features (e.g., radicals and strokes); and 2) Parameter bottleneck, each individual character has to be represented by a unique vector. In this paper, we introduce a novel representation method for Chinese characters to break the bottlenecks, namely StrokeNet, which represents a Chinese character by a Latinized stroke sequence (e.g., “凹(concave)” to “ajaie” and “凸(convex)” to “aeaqe”). Specifically, StrokeNet maps each stroke to a specific Latin character, thus allowing similar Chinese characters to have similar Latin representations. With the introduction of StrokeNet to neural machine translation (NMT), many powerful but not applicable techniques to non-Latin languages (e.g., shared subword vocabulary learning and ciphertext-based data augmentation) can now be perfectly implemented. Experiments on the widely-used NIST Chinese-English, WMT17 Chinese-English and IWSLT17 Japanese-English NMT tasks show that StrokeNet can provide a significant performance boost over the strong baselines with fewer model parameters, achieving 26.5 BLEU on the WMT17 Chinese-English task which is better than any previously reported results without using monolingual data. Code and scripts are freely available at <https://github.com/zjwang21/StrokeNet>.

1 Introduction

Neural machine translation (NMT) has become the mainstream paradigm in machine translation recently (Sutskever et al., 2014; Cho et al., 2014; Bahdanau et al., 2015). With rich bilingual parallel corpora, NMT achieves state-of-the-art performance on multiple translation benchmarks. In

*Co-first and Corresponding Author

Zh	布 什和沙龙举行了 会谈
En	Bush held a talk with Sharon
Zh (Stroke)	etasa taea teatoaie oodatot etcto ootetnea ttaeer hr tneelo oyottoottn

Table 1: StrokeNet represents a Chinese character by a Latinized stroke sequence. For example, “布” to “etasa” and “了” to “hr”.

Chinese NMT tasks, the Chinese character has been the minimum representation unit for a long time. However, such representation perhaps might not be the best choice for Chinese NMT due to the two following representation bottlenecks.

The first is the learning bottleneck. The representation learning of Chinese does not fully utilize its rich internal features. Latin languages have rich information in words like affixes. Actually, Chinese also has this kind of internal information. A Chinese character usually contains one radical (rarely has two) and several other radical-like components (Li et al., 2015). Characters with the radical “扌” commonly are verbs. The characters “扎(tie)”, “拉(pull)”, “打(hit)”, “扔(throw)” and “提(carry)” all have the meaning of acting with hands because they have the same radical “扌”. Latin languages can easily learn this internal information by subword modeling while Chinese cannot if just taking character as the minimum unit into consideration, which limits the representation capability of NMT models.

The second is the parameter bottleneck. In Chinese NLP models, the parameters used for Chinese word representation can be a huge number. In large-scale cross-lingual pre-trained language models like XLM-R (Conneau et al., 2020), mBART (Liu et al., 2020b) and mT5 (Xue et al., 2021), Chinese tokens account for a very unbalanced proportion of the vocabulary. For instance, the vocabulary of XLM-R and mBART is learned from the corpus of 100 languages, resulting in 250K subword tokens,

of which Chinese tokens account for nearly 20K. Besides, in these models, non-Latin languages like Chinese have to learn embedding individually. It is difficult for them to compress the size of vocabulary by sharing subwords as Latin languages do. In such an overparameterized scenario, NMT models might meet a serious over-fitting issue.

To break the above two representation bottlenecks of Chinese, in this paper we introduce StrokeNet, a novel representation method for Chinese characters. Specifically, StrokeNet transforms each Chinese character to its corresponding stroke sequence. Then StrokeNet transforms each stroke to a lowercased Latin character by a predefined rule. With this transformation, a Chinese character is represented by a Latinized stroke sequence, looking like an English word. The similar Chinese characters that share the same radical will also share the same affixes in their Latinized representation.

As StrokeNet represents Chinese as Latinized stroke sequence, we now can implement several powerful methods, which are previously not applicable to non-Latin languages, on Chinese NMT tasks. StrokeNet can learn subword vocabulary (Sennrich et al., 2016) on the Latinized stroke sequence to break the learning bottleneck. Through this technique, StrokeNet can benefit from internal features such as radicals to enhance representation learning. Furthermore, StrokeNet can share subword embedding between source and target languages (Press and Wolf, 2017), leading to a significant parameter reduction and overcoming model over-fitting. Finally, the ciphertext-based data augmentation (Kambhatla et al., 2022), a powerful technique in Latin languages can be added into StrokeNet to achieve performance boost in several NMT tasks. These techniques can take full advantage of the rich internal information brought by Latinized strokes.

We conduct experiments on widely-used NIST Chinese-English, WMT17 Chinese-English and IWSLT17 Japanese-English NMT tasks. The results show that StrokeNet provides a significant performance boost over strong baselines using fewer model parameters. We achieve a new state-of-the-art result of 26.5 BLEU on the WMT17 Chinese-English task, with an increment of 2.1 BLEU over the scaling Transformer baseline (Ott et al., 2018).

Our main contributions are as follows:

- We propose StrokeNet to break the representation bottleneck of Chinese characters by cap-

turing their rich internal features.

- We incorporate StrokeNet to Chinese NMT and make it possible to include the previously inapplicable methods in non-Latin languages.
- Our NMT models trained with StrokeNet outperform strong baselines by fewer model parameters, achieving a new state-of-the-art result on the WMT17 Zh-En task.

2 Related Work

2.1 Chinese Character Representation

The research on Chinese character representation mainly focuses at character level and sub-character level. Character level representation is a natural and powerful approach. Chinese word segmentation (CWS) is the mainstream paradigm in character level representation which cuts text into words consisting of at least one character. Existing research pays much attention to CWS tasks with neural network architecture. Ma et al. (2018) use Bi-LSTMs to conduct CWS, leveraging both previous and future information in a sentence while Gan and Zhang (2020) show that self-attention network gives more competitive results. These techniques perform well on Chinese NLP tasks but suffer from the representation learning bottleneck of the internal features.

The need to segment Chinese characters into smaller units and leverage their internal features arises in Chinese NLP tasks. Sub-character level representation is another promising approach. In Chinese, sub-characters contain internal features, such as radicals or strokes. Several researches focus on radical level information. Nguyen et al. (2020) believe that Chinese characters have a recursive structure and use treeLSTM to build hierarchical character embedding. Saunders et al. (2020) leverage radical level data during training, which proves that applying radical decomposition improves Chinese-Japanese translation and performs well on translating unseen words.

Another line focuses on stroke level information. Cao et al. (2018) propose stroke n-gram for learning Chinese character embedding with stroke n-gram information. Zhang and Komachi (2019); Han et al. (2021) focus on leveraging varying degrees of sub-character data, which points out that the stroke level system performs better than the ideograph level systems. Zhang and Komachi (2018) decompose each Chinese character

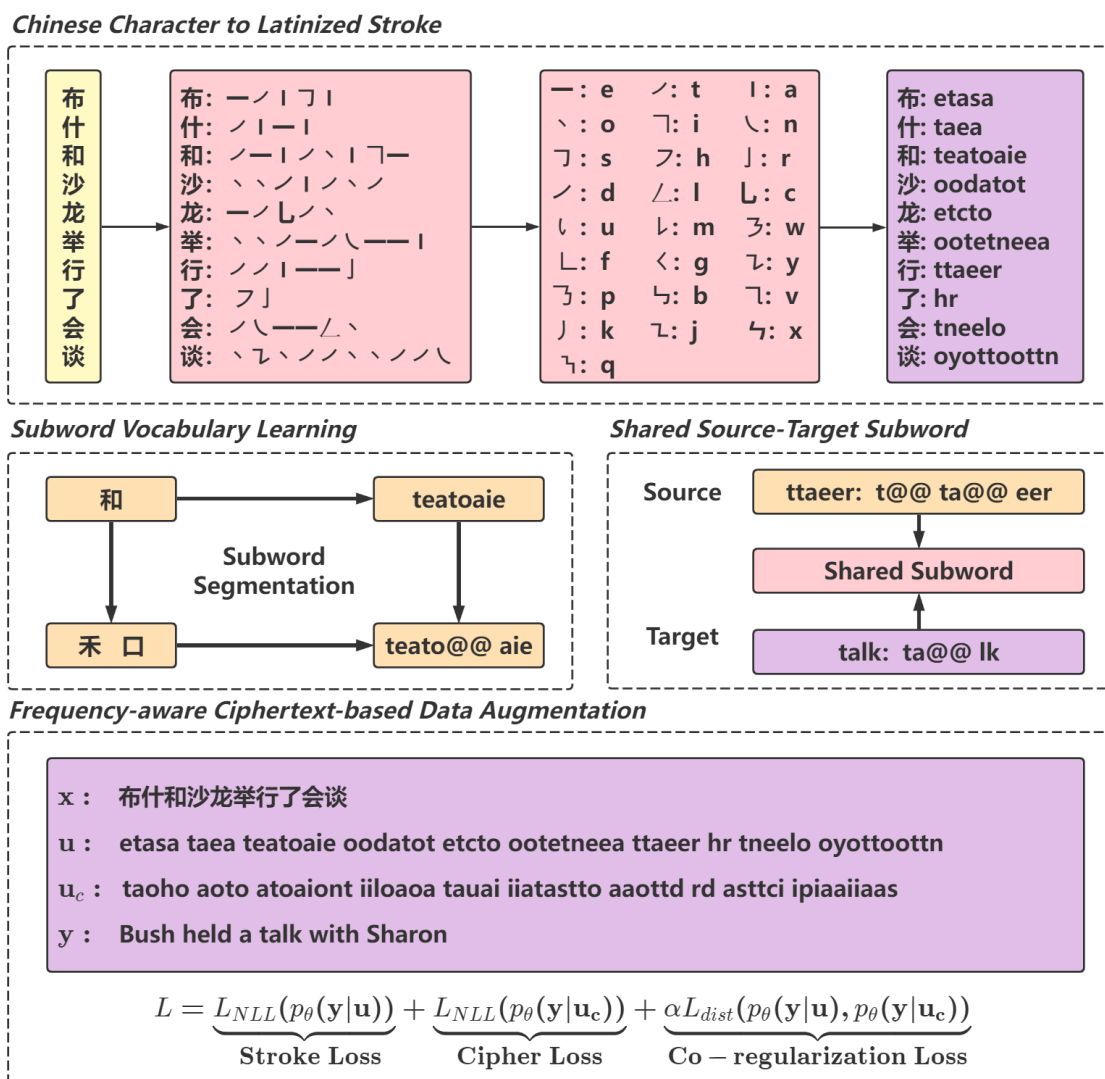


Figure 1: Overall framework of StrokeNet. Each Chinese character is mapped to a sequence of Latin characters, like an English word. Many powerful techniques inapplicable to Chinese now can be easily applied to NMT tasks.

into ideographs and each kind of ideograph is then decomposed into 33 kinds of predefined strokes.

Other works try to utilize the glyph or pronunciation information in Chinese NLP tasks. Dai and Cai (2017) propose glyph-aware embedding of Chinese characters. Sun et al. (2021) propose ChineseBERT, which fuses the glyph and pinyin embeddings with the original character embedding to enhance Chinese representation. Si et al. (2021) propose sub-character tokenization to encode a Chinese character into a sequence of its glyph or pronunciation and learn a new vocabulary for Chinese language model pretraining.

These researches mainly focus on learning internal features to enhance Chinese language understanding tasks. However, the learning bottleneck still exists in Chinese NMT tasks.

2.2 Subword Learning for NMT

Subword learning is widely used to address the limited vocabulary problem in NMT and has been proved powerful (Sennrich et al., 2016). Several researches leverage different segmentation as augmented data or a noisy term during training. Kudo (2018) propose subword regularization by integrating different segmentation of words to NMT models by probability. Provilkov et al. (2020) propose the BPE-dropout technique to stochastically corrupt the segmentation procedure of BPE. Wang et al. (2021) propose multi-view subword regularization to make full use of different kinds of segmentation. Manghat et al. (2022) propose a hybrid subword segmentation algorithm to deal with out-of-vocabulary words. Tay et al. (2021) propose a soft gradient-based subword tokenization

algorithm to learn subword representation in data-driven fashion. Ács et al. (2021) investigate how different strategies of subword pooling affect the downstream performance.

Shared embedding is a popular and powerful technique jointly used with subword learning in NMT (Press and Wolf, 2017; Pappas et al., 2018; Liu et al., 2019a,b). For NMT tasks in Latin languages, shared subword learning has become the de facto standard to improve the performance of NMT and compress the vocabulary size (Vaswani et al., 2017; Joshi et al., 2020; Dai et al., 2019). This technique reduces the size of NMT models greatly without harming their performance. However, this technique is difficult to implement in Chinese because of the differences between Chinese and Latin languages, making it hard to break the above introduced parameter bottleneck.

3 StrokeNet

To break the learning and parameter bottlenecks of Chinese character representation, we propose StrokeNet that maps a Chinese character into a Latinized sequence, and apply it to NMT tasks. Figure 1 shows the overall framework of StrokeNet.

3.1 Chinese Character to Latinized Stroke

Chinese Character to Strokes Mapping To learn more internal information in Chinese characters, we first need to map Chinese character to its corresponding stroke sequence. Formally, given a Chinese sentence $\mathbf{x} = (x_1, x_2, x_3, \dots, x_n)$, StrokeNet maps it into $\mathbf{s} = (s_1, s_2, s_3, \dots, s_n)$, where s_i represents the corresponding stroke sequence of x_i . As shown in Figure 1, the Chinese word “什” can be transferred to the stroke sequence “丨丨丨”. Besides, since a small number of Chinese characters have the same sequence of strokes, we follow Zhang and Komachi (2018) to make them distinguishable by adding a different digit at the end of the stroke sequence. For instance, “井” and “开” have the same stroke sequence “一一丨”. In StrokeNet, the corresponding stroke sequence of “井” is “一一丨0” and “开” is “一一丨1”. Without loss of generality, we follow the implementation¹ of the previous work (Cao et al., 2018) to define strokes, which is the most widely-used criterion consisting of 25 kinds of strokes.² Through stroke

level representation, more internal information is easier to learn for NMT models.

Frequency Mapping To make Latin language techniques applicable in StrokeNet, we then map the stroke vocabulary, which consists of 25 kinds of strokes, to the lowercased Latin alphabet of 26 characters. Lexical marker is an important part of information composition. Frequent words are low-information words because they have few lexical markers (Finn, 1977). Inspired by this information theory, we construct the mapping rule by the frequency of character occurrence. For instance, the Latin character “e” has the highest frequency of 12.7% in English while the stroke “一” has the highest frequency of 27.9% in Chinese. We map “一” to “e” and follow the frequency order to map the other strokes to Latin characters. We leave the character “z” not mapped because we only define 25 kinds of strokes and “z” has the lowest frequency in English. Finally, we get the Latinized stroke sequence of the Chinese sentence. We use $\mathbf{u} = (u_1, u_2, u_3, \dots, u_n)$ to represent the corresponding Latinized stroke sequence of \mathbf{x} . Appendix A.1 shows the mapping dictionary.

3.2 Application to NMT tasks

We apply StrokeNet to Chinese NMT tasks and introduce how to combine it with other popular techniques of NMT. It is noted that the techniques are inapplicable to Chinese NMT without StrokeNet.

Subword Vocabulary Learning We use the subword vocabulary learning (Sennrich et al., 2016) technique to break the learning bottleneck of internal information. After mapping Chinese characters to Latinized stroke sequences, characters are decomposed into smaller units. We conduct byte pair encoding (BPE) algorithm on the corpus of Latinized strokes. BPE segments Chinese characters into smaller components like subwords in English. During training, NMT models utilize this segmentation to learn better representation. For instance, the character “和” can be cut into “禾” and “口” because its corresponding Latinized stroke sequence “teatoaie” can be cut into “teato@@” and “aie”. According to Li et al. (2015), simple characters in Chinese account for less than 20% which cannot be split into other components. The others are compound characters. So more than 80% of Chinese characters can benefit from our stroke-based representation. With this advantage, we can learn stronger Chinese representation.

¹<https://github.com/bamtercelboo/cw2vec>

²Si et al. (2021) represent Chinese characters into Latinized stroke sequence with 5 kinds of strokes, but do not seem to get clear improvements.

Shared Source-Target Representation We then use shared subword embedding (Press and Wolf, 2017) between Latinized strokes and English to break the parameter bottleneck. NMT models with shared embedding can benefit from shared source-target representation and parameter reduction. After transforming Chinese characters to Latinized stroke sequences, Chinese can jointly learn BPE merge operations with Latin Languages. For example, if we cut the Latinized strokes of “行” into “t@@ ta@@ eer”, and the English word “talk” into “ta@@ lk”, the “ta@@” could be a shared subword in both Latinized Chinese and English. The shared source-target representation can work as a regularization term in the model training, smoothing the learning process. Besides, the difference in parameters between NMT models with the same architecture mainly comes from the vocabulary size. The shared subword vocabulary can also lead to a great parameter reduction.

Frequency-aware Ciphertext-based Data Augmentation As a powerful technique in NMT of Latin languages, the ciphertext-based data augmentation (CDA) is difficult to implement in Chinese NMT tasks due to the huge character list (Kambhatla et al., 2022). StrokeNet addresses the limitation and now it can be well implemented. CDA is a character substitution method that replaces a character in the text with the k th character after it in the alphabet. For cipher- I , the character “e” is replaced by “f” to produce the pseudo source text and other characters follow the same rule. The last character “z” in the alphabet is replaced to “a”. k represents the distance between the source character and the target replaced character.

In StrokeNet, we further propose a frequency-aware ciphertext-based data augmentation (FCDA). FCDA replaces a character with the k th character after it by the frequency order instead of alphabet order. For cipher- I , the character “e” is replaced by “t” instead of “f” because “e” has the highest frequency and “t” has the second highest frequency. We apply FCDA to the Latinized stroke sequence, producing the pseudo sources of the same semantic meaning and performing the multi-source learning for NMT as follows:

$$L = \underbrace{L_{NLL}(p_{\theta}(\mathbf{y}|\mathbf{u}))}_{\text{Stroke Loss}} + \underbrace{L_{NLL}(p_{\theta}(\mathbf{y}|\mathbf{u}_{\mathbf{c}}))}_{\text{Cipher Loss}} + \underbrace{\alpha L_{dist}(p_{\theta}(\mathbf{y}|\mathbf{u}), p_{\theta}(\mathbf{y}|\mathbf{u}_{\mathbf{c}}))}_{\text{Co-regularization Loss}}$$

We follow Kambhatla et al. (2022) to minimize three losses in training, i.e., the stroke loss for the Latinized strokes, the cipher loss for the ciphered Latinized strokes, and the co-regularization loss. This method can reduce the impact of rare words and significantly improve performance in NMT.

4 Evaluation

We aim to answer the research questions through the following experiments:

- Can StrokeNet improve the performance of Chinese NMT tasks?
- Can StrokeNet reduce the scale of parameters of NMT models?

4.1 Experimental Setup

Data We conduct experiments on the NIST Zh-En and WMT17 Zh-En benchmarks. For the NIST Zh-En, the training data contains 1.25M sentence pairs. We use MT06 as the validation set and report results on MT02, MT03, MT04, and MT08 test sets, with each consists of four references. For the WMT17 Zh-En, the training data contains 20M sentence pairs. The development set is newsdev2017 and the test set is newstest2017. We use the scripts in Moses (Koehn et al., 2007) to tokenize and true-case the data.³ We use jieba⁴ to conduct CWS. Then we apply the BPE algorithm to Chinese and English separately. For the NIST Zh-En, we execute 30K BPE merge operations on Chinese and English separately in the baseline and 30K joint-BPE operations on Chinese and English together in StrokeNet. For the WMT17 Zh-En, we conduct 32K BPE operations on Chinese and English separately in the baseline, and 50K joint-BPE operations in StrokeNet.

Besides, to verify the validity of StrokeNet in other non-Latin languages, we also conduct experiments on the IWSLT17 Ja-En, which contains 223K training sentence pairs. The preprocessing keeps the same with the other two benchmarks. We use mecab⁵ to conduct Japanese word segmentation. We make a statistic of the composition of Japanese. Chinese characters account for about 26%. Japanese pseudonyms account for about 52% and others appear to be special characters. There

³<https://github.com/amos-sm/amos-sm-decoder>

⁴<https://github.com/fxsjy/jieba>

⁵<https://github.com/taku910/mecab>

Model	Parameters		Performance (BLEU)					
	Total	Emb.	Valid	MT02	MT03	MT04	MT08	ALL
Shared-Private (Liu et al., 2019b)	63M	19M	42.6	43.7	42.0	44.5	33.8	41.6
AdvGen (Cheng et al., 2019)	95M	39M	47.0	47.0	46.5	47.4	37.4	-
AdvAug (Cheng et al., 2020)	95M	39M	49.2	49.0	48.0	48.9	39.6	-
Manifold (Chen et al., 2021a)	83M	46M	49.4	49.6	50.3	49.5	39.2	-
Vanilla	80M	36M	47.4	47.7	47.5	47.7	38.1	45.5
StrokeNet	59M	15M	49.7	50.7	49.8	50.2	41.3	48.1

Table 2: Model parameters and performance (BLEU) on the NIST Zh-En translation task. ‘‘Emb.’’ denotes the parameters used for the embedding layer. StrokeNet provides a significant performance boost over the strong baselines with dramatically fewer model parameters.

Model		N. Zh-En	W. Zh-En	Ja-En
Vanilla	Src	40K	50K	33K
	Trg	30K	37K	30K
StrokeNet		29K	50K	28K

Table 3: The vocabulary sizes of the three corpora. The vanilla baseline learns BPE operations separately on the source and target text. StrokeNet learns joint vocabulary and shares all embeddings of the model.

are 190 kinds of Japanese pseudonyms in total. Based on this fact, we simply apply StrokeNet to those Chinese characters in Japanese after converting them to simplified Chinese. The rest characters keep unchanged. The BPE merge operations are 30K for source and target separately in the baseline and 30K for joint vocabulary in StrokeNet. The final vocabulary sizes of these three corpora are detailed in Table 3.

Model We use three kinds of Transformer architectures in the experiments. For the small-scale Ja-En dataset, we use the small architecture with hidden size 288, FFN size 507, 4 heads, and 5 encoder/decoder layers. For the medium-scale NIST Zh-En dataset, we use the base architecture with hidden size 512, FFN size 2048, 8 heads, and 6 encoder/decoder layers. For the large-scale WMT17 Zh-En dataset, we use the large architecture with hidden size 1024, FFN size 4096, 16 heads, and 6 encoder/decoder layers.

Settings For training vanilla NMT models, the decoder input and output embeddings are shared. For the NIST Zh-En, we use Adam to train for 100K steps, with 32K max tokens per batch, the learning rate 0.0005, $\beta_1 = 0.9$, $\beta_2 = 0.98$, weight decay of 0.0001 and dropout ratio 0.3. We warm up

the learning rate for the first 4K steps and then use the inverse square root scheduler. For the WMT17 Zh-En, the number of max tokens per batch is 288K and others keep the same with the NIST Zh-En. For the Ja-En, we train for 100K steps with 16K max tokens per batch and the learning rate is 0.0003 with a cosine scheduler.

For training NMT models with StrokeNet, each Chinese character is mapped to the corresponding Latinized stroke sequence. Joint vocabulary is learned from both source and target texts together. During training, all the embeddings and softmax weights are shared. For the Ja-En, NIST Zh-En, and WMT17 Zh-En, the max tokens per batch are 16K, 96K, and 192K respectively. The rest hyperparameters keep the same with the vanilla models.

During testing, for the NIST Zh-En and Ja-En, we generate with length penalty 1.0 and beam size 5. For the WMT17 Zh-En, we generate with length penalty 1.4 and beam size 5. The BLEU scores are evaluated with *multi-bleu.perl* provided by Moses (Koehn et al., 2007). We use the checkpoint with the best validation BLEU for testing.

4.2 Main Results

Parameter Reduction To verify whether StrokeNet can reduce the parameter of NMT models, we look into the parameter and vocabulary size in each experiment. Table 3 and Table 2 show the parameters and vocabulary sizes in all experiments. StrokeNet decreases the vocabulary size and the parameters obviously. For the NIST Zh-En, the embedding layer parameters are 15M, which is the smallest. The model parameters of StrokeNet are 59M, while the vanilla baseline and previous methods are over 80M generally. Compared with the prior work (Chen et al., 2021a) with competi-

Model	Ja-En	W. Zh-En
Shared-Private (Liu et al., 2019b)	12.4	-
Norm (Liu et al., 2020a)	-	25.3
Prior (Chen et al., 2021b)	-	25.5
Vanilla	12.0	24.4
StrokeNet	13.1	26.5

Table 4: Performance (BLEU) on the Ja-En and large-scale WMT17 Zh-En benchmarks.

tive performance, StrokeNet has a 24M parameter reduction. This parameter reduction can greatly decrease the pressure of computational complexity while achieving more competitive results.

For the WMT17 Zh-En and Ja-En, we list their vocabulary sizes in Table 3. StrokeNet gains 37K, and 35K vocabulary reduction respectively on these two corpora. With the smaller vocabulary, StrokeNet reduces the representation redundancy in Chinese characters and learns shared representation. As the parameter reduction of NMT models mainly comes from the vocabulary size reduction, StrokeNet on these two datasets also has an obvious parameter reduction. These results show that StrokeNet breaks the parameter bottleneck, which can alleviate the over-fitting problem.

Performance Boost To verify whether StrokeNet can improve performance on Chinese NMT tasks, we look into the results of the three benchmarks. Table 2 shows the translation performance of the validation and test sets on the NIST Zh-En benchmark. StrokeNet obtains a BLEU of 48.1 on the collection of all test sets, an improvement of 2.6 BLEU over the vanilla baseline. We also see improvements over prior work (Chen et al., 2021a) on every subset except MT03.

Table 4 shows the translation performance on the test sets for the Ja-En and large-scale WMT17 Zh-En. For the Ja-En, StrokeNet improves translation quality by 1.1 BLEU over the vanilla baseline and 0.7 BLEU over the prior work (Liu et al., 2019b). Although there are only about 26% Chinese characters in Japanese, StrokeNet can still gain 1.1 BLEU improvement on the Ja-En task. For the WMT17 Zh-En, StrokeNet achieves a new state-of-the-art result of 26.5 BLEU, obtaining +2.1 BLEU over the vanilla baseline, +1.2 BLEU over Liu et al. (2020a), and +1.0 BLEU over Chen et al. (2021b). In the future, we will further enhance the NMT models with pretrained knowledge (Liu et al., 2021a,b).




Figure 2: BLEU on the valid and test sets for different numbers of BPE merge operations.




Figure 3: Learning curves for the vanilla baseline and StrokeNet.

The above results confirm the effectiveness of StrokeNet, which can be applied to different scales of training data and different languages.

5 Analysis

Effect of BPE Merge Operations StrokeNet benefits from the subword modeling technique. To explore how it works in StrokeNet, we conduct experiments on the NIST Zh-En benchmark, applying different numbers of BPE merge operations. We conduct experiments on 20K, 30K, 40K, and 50K merge operations. Results are detailed in Figure 2. 30K merge operations appear to be the best choice. For the validation set, the translation performance reaches the highest at 30K. For the test set, it gradually decreases as the number of merge operations increases. We see variations of less than 0.7 BLEU in the dev set and less than 0.4 BLEU in the test set as the number of BPE merge operations changes. And large improvements over the vanilla baseline are observed regardless of the number of BPE merge operations with at least +2.2 BLEU for both the validation and test sets. The results show that StrokeNet is robust to the number of BPE merge operations.




Figure 4: BLEU on the NIST Zh-En for different scales of training data. We randomly extract the data subsets from the 1.25M sentence pairs.

Learning Curves To explore how StrokeNet performs better than the vanilla baseline, we draw the learning curves during training. We train StrokeNet and the vanilla baseline for both 50 epochs on the NIST Zh-En. The validation BLEU during training is shown in Figure 3.⁶ In StrokeNet, the BLEU on the validation set rises faster in the early period and finally achieves higher BLEU than the vanilla baseline. With the Latinized stroke level representation and the application of powerful techniques in Latin languages, NMT models with StrokeNet learn faster and better than vanilla models. The rich internal information in Chinese characters becomes readily available for StrokeNet to learn. The results prove that StrokeNet successfully breaks the representation learning bottleneck, showing its positive effects on Chinese NMT tasks.

Effect of Data Scale To further illustrate the effects of different data scales in StrokeNet, we randomly extract four subsets from the original 1.25M source sentence pairs in the NIST Zh-En. We control their size to be 100k, 300k, 600k, and 900k. The results on the test set are given in Figure 4. With each data scale, StrokeNet yields large improvements over the vanilla baseline by 2.1-5.7 BLEU. Furthermore, with the decrement of data scale, the performance margin between StrokeNet and the vanilla baseline becomes larger. In particular, the improvement is much larger under the 100K setting (+5.7 BLEU) than that under the 1.25M setting (+2.6 BLEU). Besides, we also see great improvements on the low-scale Ja-En and large-scale WMT17 Zh-En. StrokeNet is proved powerful on NMT tasks of varying data sizes.

Ablation Analysis To explore which part of StrokeNet makes a difference, we conduct several ab-

⁶The BLEU here is calculated with one reference.

Model	Emb.	BLEU
Vanilla	36M	45.5
StrokeNet w/ Rand. Mapping	15M	47.5
StrokeNet w/ Freq. Mapping	15M	48.1
StrokeNet w/o Freq. CDA	15M	45.9
StrokeNet w/o Shared Voc.	21M	47.7

Table 5: Performance (BLEU) of different model variants on the NIST Zh-En benchmark.

Model	Ratio	Length
StrokeNet w/ Rand. Mapping	37.9	5.7
StrokeNet w/ Freq. Mapping	39.1	5.8

Table 6: Statistics of shared subwords in the NIST Zh-En data obtained using frequency mapping and random mapping with the same 30K BPE merge operations.

lation experiments on the NIST Zh-En benchmark. First, we explore the effect of the frequency mapping technique, which is inspired by the information theory (Finn, 1977). The theory states that more frequent words are lower information words because they have fewer lexical markers. We compare StrokeNet with frequency mapping to StrokeNet with randomly mapping, which means that each kind of stroke is mapped to a unique Latin character randomly. Table 5 shows that the performance of StrokeNet with mapping randomly is 0.6 BLEU worse than frequency mapping, which proves that our application of the information theory is reasonable. Mapping units in two languages by their frequency reduces information loss.

To further explore how frequency mapping improves performance, we conduct statistics of shared subwords in the NIST Zh-En data obtained using frequency mapping and random mapping with the same BPE merge operations. Table 6 gives the results. Ratio refers to the ratio of shared subwords over the whole subword in the training data. Length refers to the weighted average of the length of shared subwords. BPE is also a mapping algorithm based on frequency and can benefit from the proposed frequency mapping. The results show that frequency mapping produces more shared subwords and longer subword units between the Latinized stroke sequences and English, resulting in shorter sequences, which can lead to stronger memorization in Transformer models (Kharitonov et al., 2021), and thus better translation quality.

Second, we explore the effect of FCDA. We con-

Model	Low	Medium	High	Total
Vanilla	38.9	47.0	63.2	59.5
StrokeNet	39.2	48.5	64.2	60.6

Table 7: Prediction accuracy for words of different frequencies. StrokeNet performs well on medium- and high-frequency words.

duct experiments on StrokeNet without FCDA and keep the other settings unchanged. The result in Table 5 still yields an improvement over the vanilla baseline. StrokeNet enables this strong data augmentation technique in Latin languages to be implemented for Chinese NMT tasks.

Finally, we explore the effect of shared vocabulary and embedding. We use this technique to achieve shared representation and parameter reduction. We conduct experiments in StrokeNet without sharing vocabulary. The results are detailed in Table 5. Without sharing vocabulary, the performance decreases by 0.4 BLEU but still gains large improvements over the vanilla baseline. The parameters of StrokeNet are 6M fewer than StrokeNet without sharing vocabulary, which is consistent with intuition. Through sharing vocabulary, StrokeNet achieves parameter reduction and better performance by learning shared source-target representation. This means that the shared subword learning technique works well in StrokeNet.

Translation of Word of Different Frequency

To explore the translation quality difference between the vanilla baseline and StrokeNet, we conduct an analysis by comparing accuracy on different frequency words in the test set on the NIST Zh-En benchmark. The frequency of words is based on the training set. As shown in Table 7, StrokeNet achieves pretty good translation accuracy on medium and high-frequency words. For words of medium frequency between 200 and 2,000, StrokeNet achieves 48.5 and shows an improvement of 1.5 BLEU over the vanilla baseline. For words of high frequency over 2,000, it achieves 64.2 while the baseline achieves only 63.2. Words of low frequency, also known as rare words, still get an improvement of 0.3 over the vanilla baseline. For all the words, StrokeNet improves the prediction accuracy from 59.5 to 60.6. The results show that representation learning has been improved by learning more internal features through stroke modeling.

6 Conclusion

In this paper we introduce StrokeNet, a novel technique for Chinese NMT tasks using Latinized stroke sequence of Chinese characters. StrokeNet breaks the representation learning bottleneck and the parameter bottleneck in Chinese NMT tasks, which requires no external data and significantly outperforms several strong prior works. We show that representing Chinese characters in stroke level works well on NMT tasks to bring more internal structure information. We demonstrate that it is possible to implement popular and powerful techniques designed for Latin languages in Chinese NMT tasks. We conduct several analyses on the effects of these Latin language techniques, proving they bring an obvious performance boost to StrokeNet. Overall, StrokeNet is a simple and effective approach for Chinese NMT tasks and yields strong results in both high-source and low-source settings. Future work includes applying StrokeNet to other language generation tasks (Liu et al., 2021c).

Limitations

Challenges remain in StrokeNet. As shown in Table 7, even with the best BPE merge operations, the translation accuracy of low-frequency words gains a minor boost over the baseline by just 0.3, which is not as good as middle and high-frequency words. We speculate that low-frequency Chinese characters might be hurt when they are cut into subwords. For example, the low-frequency Chinese character “劑 (medicament)”, whose corresponding Latinized stroke sequence is “oetasttmntaeear”, is segmented into “oet@@@ a@@ stt@@ m@@ n@@ ta@@ eea@@ r”. It is too chopped up and its semantic information becomes incomplete. How to handle this kind of segmentation and improve the translation quality of low-frequency Chinese characters remains to be explored.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (Grant No. 62206076) and Shenzhen College Stability Support Plan (Grant No. GXWD20220811173340003 and GXWD20220817123150002). The computational resources were supported by Education Center of Experiments and Innovations at Harbin Institute of Technology, Shenzhen. We would like to thank the anonymous reviewers and meta-reviewer for their insightful comments and suggestions.

References

- Judit Ács, Ákos Kádár, and Andras Kornai. 2021. [Subword pooling makes a difference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2284–2295, Online. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Shaosheng Cao, Wei Lu, Jun Zhou, and Xiaolong Li. 2018. [cw2vec: Learning chinese word embeddings with stroke n-gram information](#). 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 5053–5061. AAAI Press.
- Guandan Chen, Kai Fan, Kaibo Zhang, Boxing Chen, and Zhongqiang Huang. 2021a. [Manifold adversarial augmentation for neural machine translation](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3184–3189, Online. Association for Computational Linguistics.
- Kehai Chen, Rui Wang, Masao Utiyama, and Eiichiro Sumita. 2021b. [Integrating prior translation knowledge into neural machine translation](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:330–339.
- Yong Cheng, Lu Jiang, and Wolfgang Macherey. 2019. [Robust neural machine translation with doubly adversarial inputs](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4324–4333, Florence, Italy. Association for Computational Linguistics.
- Yong Cheng, Lu Jiang, Wolfgang Macherey, and Jacob Eisenstein. 2020. [AdvAug: Robust adversarial augmentation for neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5961–5970, Online. Association for Computational Linguistics.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Falcon Dai and Zheng Cai. 2017. [Glyph-aware embedding of Chinese characters](#). In *Proceedings of the First Workshop on Subword and Character Level Models in NLP*, pages 64–69, Copenhagen, Denmark. Association for Computational Linguistics.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. 2019. [Transformer-xl: Attentive language models beyond a fixed-length context](#). *CoRR*, abs/1901.02860.
- Patrick J Finn. 1977. [Word frequency, information theory, and cloze performance: A transfer feature theory of processing in reading](#). *Reading Research Quarterly*, pages 508–537.
- Leilei Gan and Yue Zhang. 2020. [Investigating self-attention network for chinese word segmentation](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2933–2941.
- Lifeng Han, Gareth Jones, Alan Smeaton, and Paolo Bolzoni. 2021. [Chinese character decomposition for neural MT with multi-word expressions](#). In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 336–344, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. 2020. [Spanbert: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Nishant Kambhatla, Logan Born, and Anoop Sarkar. 2022. [CipherDAug: Ciphertext based data augmentation for neural machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 201–218, Dublin, Ireland. Association for Computational Linguistics.
- Eugene Kharitonov, Marco Baroni, and Dieuwke Hupkes. 2021. [How BPE affects memorization in transformers](#). *CoRR*, abs/2110.02782.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.

- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. Association for Computational Linguistics.
- Yanran Li, Wenjie Li, Fei Sun, and Sujian Li. 2015. [Component-enhanced Chinese character embeddings](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 829–834, Lisbon, Portugal. Association for Computational Linguistics.
- Xuebo Liu, Houtim Lai, Derek F. Wong, and Lidia S. Chao. 2020a. [Norm-based curriculum learning for neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 427–436, Online. Association for Computational Linguistics.
- Xuebo Liu, Longyue Wang, Derek F. Wong, Liang Ding, Lidia S. Chao, Shuming Shi, and Zhaopeng Tu. 2021a. [On the complementarity between pre-training and back-translation for neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2900–2907, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xuebo Liu, Longyue Wang, Derek F. Wong, Liang Ding, Lidia S. Chao, Shuming Shi, and Zhaopeng Tu. 2021b. [On the copying behaviors of pre-training for neural machine translation](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4265–4275, Online. Association for Computational Linguistics.
- Xuebo Liu, Longyue Wang, Derek F. Wong, Liang Ding, Lidia S. Chao, and Zhaopeng Tu. 2021c. [Understanding and improving encoder layer fusion in sequence-to-sequence learning](#). In *International Conference on Learning Representations*.
- Xuebo Liu, Derek F. Wong, Lidia S. Chao, and Yang Liu. 2019a. [Latent attribute based hierarchical decoder for neural machine translation](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(12):2103–2112.
- Xuebo Liu, Derek F. Wong, Yang Liu, Lidia S. Chao, Tong Xiao, and Jingbo Zhu. 2019b. [Shared-private bilingual word embeddings for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3613–3622, Florence, Italy. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020b. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Ji Ma, Kuzman Ganchev, and David Weiss. 2018. [State-of-the-art Chinese word segmentation with Bi-LSTMs](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4902–4908, Brussels, Belgium. Association for Computational Linguistics.
- Sreeja Manghat, Sreeram Manghat, and Tanja Schultz. 2022. [Hybrid sub-word segmentation for handling long tail in morphologically rich low resource languages](#). In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6122–6126. IEEE.
- Minh Nguyen, Gia H. Ngo, and Nancy F. Chen. 2020. [Hierarchical character embeddings: Learning phonological and semantic representations in languages of logographic origin using recursive neural networks](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:461–473.
- 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*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.
- Nikolaos Pappas, Lesly Miculicich, and James Henderson. 2018. [Beyond weight tying: Learning joint input-output embeddings for neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 73–83, Brussels, Belgium. Association for Computational Linguistics.
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.
- Ivan Provilkov, Dmitrii Emelianenko, and Elena Voita. 2020. [BPE-dropout: Simple and effective subword regularization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1882–1892, Online. Association for Computational Linguistics.
- Danielle Saunders, Weston Feely, and Bill Byrne. 2020. [Inference-only sub-character decomposition improves translation of unseen logographic characters](#). In *Proceedings of the 7th Workshop on Asian Translation*, pages 170–177, Suzhou, China. Association for Computational Linguistics.
- 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 (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.

Chenglei Si, Zhengyan Zhang, Yingfa Chen, Fanchao Qi, Xiaozhi Wang, Zhiyuan Liu, and Maosong Sun. 2021. [SHUOWEN-JIEZI: linguistically informed tokenizers for chinese language model pretraining](#). *CoRR*, abs/2106.00400.

Zijun Sun, Xiaoya Li, Xiaofei Sun, Yuxian Meng, Xiang Ao, Qing He, Fei Wu, and Jiwei Li. 2021. [Chinese-BERT: Chinese pretraining enhanced by glyph and Pinyin information](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2065–2075, Online. Association for Computational Linguistics.

Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3104–3112.

Yi Tay, Vinh Q. Tran, Sebastian Ruder, Jai Prakash Gupta, Hyung Won Chung, Dara Bahri, Zhen Qin, Simon Baumgartner, Cong Yu, and Donald Metzler. 2021. [Charformer: Fast character transformers via gradient-based subword tokenization](#). *ArXiv preprint*, abs/2106.12672.

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, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.

Xinyi Wang, Sebastian Ruder, and Graham Neubig. 2021. [Multi-view subword regularization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 473–482, Online. Association for Computational Linguistics.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Longtu Zhang and Mamoru Komachi. 2018. [Neural machine translation of logographic language using sub-character level information](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 17–25, Brussels, Belgium. Association for Computational Linguistics.

Longtu Zhang and Mamoru Komachi. 2019. [Chinese-japanese unsupervised neural machine translation using sub-character level information](#). *CoRR*, abs/1903.00149.

A Appendix

A.1 Statistical Data of Character and Stroke Frequencies

Figure 5 shows the frequency of occurrence of each lowercased Latin character and each Chinese stroke. The frequency of each lowercased Latin character is from Wikipedia.⁷ The frequency of each stroke is from WMT17 Zh-En data, which contains 20M Chinese sentences. We order them by frequency and each stroke is mapped to the Latin lowercased character in the same row. “z” has no corresponding stroke because we only define 25 kinds of strokes and it has the minimum frequency.

Letter	Freq	Stroke	Freq
e	12.702	一	27.940
t	9.056	丿	16.869
a	8.167	丨	16.618
o	7.507	丶	13.223
i	6.966	冫	6.060
n	6.749	㇇	3.873
s	6.327	冫	2.917
h	6.094	フ	2.399
r	5.987	丿	2.214
d	4.253	丿	2.112
l	4.025	厶	1.507
c	2.782	㇇	0.983
u	2.758	㇇	0.513
m	2.406	㇇	0.485
w	2.36	㇇	0.474
f	2.228	㇇	0.402
g	2.015	㇇	0.327
y	1.974	㇇	0.313
p	1.929	㇇	0.227
b	1.492	㇇	0.218
v	0.978	㇇	0.134
k	0.772	㇇	0.095
j	0.153	㇇	0.068
x	0.150	㇇	0.028
q	0.095	㇇	0.00012
z	0.074		

Figure 5: Frequencies of Latin lowercased characters and Chinese strokes.

⁷https://en.wikipedia.org/wiki/Letter_frequency