

Scaling Back-Translation with Domain Text Generation for Sign Language Gloss Translation

Jinhui Ye^{1*} Wenxiang Jiao^{2*} Xing Wang^{2†} Zhaopeng Tu²

¹The Hong Kong University of Science and Technology (Guangzhou)

jye624@connect.hkust-gz.edu.cn

²Tencent AI Lab

{joelwxjiao,brightxwang,zptu}@tencent.com

Abstract

Sign language gloss translation aims to translate the sign glosses into spoken language texts, which is challenging due to the scarcity of labeled gloss-text parallel data. Back-translation (BT), which generates pseudo parallel data by translating in-domain spoken language texts into sign glosses, has been applied to alleviate the data scarcity problem. However, the lack of large-scale high-quality in-domain spoken language text data limits the effect of BT. In this paper, to overcome the limitation, we propose a **P**rompt based domain text **G**eneration (PGEN) approach to produce the large-scale in-domain spoken language text data. Specifically, PGEN randomly concatenates sentences from the original in-domain spoken language text data as prompts to induce a pre-trained language model (i.e., GPT-2) to generate spoken language texts in similar style. Experimental results on three benchmarks of sign language gloss translation in varied languages demonstrate that BT with spoken language texts generated by PGEN significantly outperforms the compared methods. In addition, as the scale of spoken language texts generated by PGEN increases, the BT technique can achieve further improvements, demonstrating the effectiveness of our approach. We release the code and data for facilitating future research in this field¹.

1 Introduction

Sign language is the dominant form of communication for the deaf and hearing impaired community. Sign language processing has received substantial attention in the last few years and achieved significant progress (Bragg et al., 2019; Yin et al., 2021; Shterionov, 2021; De Coster et al., 2022). Among

*Jinhui Ye and Wenxiang Jiao contributed equally to this work. Work was mainly done when Jinhui Ye was interning at Tencent AI Lab.

†Xing Wang is the corresponding author.

¹Code and data are available at <https://github.com/Atrewin/PGen>.




Figure 1: The translation performance of back-translation when scaling the generated spoken language data from 1) our PGEN which uses the large pretrained language model, and 2) selected in-domain data from large-scale spoken language corpus. The red dashed line denotes the baseline model without back-translation. Best viewed in color.

them, sign language translation (SLT) aims to transform continuous sign language videos into natural spoken language texts (Bungeroth and Ney, 2004; Camgoz et al., 2018). SLT consists of two sub-tasks: (1) sign language understanding task that recognizes the continued videos to the sign glosses; and (2) spoken language gloss translation (SLGT) task that generates the spoken language text of the given sign glosses. In this work we focus on the second sub-task, i.e., SLGT.

Data scarcity has been considered the major limitation of sign language gloss translation (Moryossef et al., 2021; Zhang and Duh, 2021). To alleviate the data scarcity problem, back-translation (Sennrich et al., 2016), which translates in-domain spoken language texts into sign glosses to construct synthetic parallel data, has been adopted and achieved certain success in SLGT. However, the lack of large-scale high-quality in-domain spoken language text data limits the capability of back-translation for the SLGT task (Zhang and Duh, 2021). The common practice is adopt-

ing data selection (Axelrod et al., 2011), or data mining (Jiang et al., 2009) approaches to obtain in-domain data. But these approaches assume the availability of enough in-domain data or expert knowledge (e.g., language background to mine the in-domain data in a specific language), which prevents them from applying to gloss translation task that has broader scenarios (e.g., more languages and domains).

In this work, we propose a **Prompt-based domain text Generation (PGEN)** to generate large-scale high-quality in-domain spoken language text data, motivated by the advances in data augmentation with pretrained language models (PLMs). The main idea is to induce the large PLMs to mimic the style of original spoken language texts with prompt-based learning techniques (Radford et al., 2019; Liu et al., 2021). Our PGEN approach is able to generate large-scale in-domain spoken language texts based on the small-scale original monolingual texts and maintains diversity (§2.1). Besides, our approach can be performed without requiring large-scale in-domain data or expert knowledge of the sign language domain and maintain the high quality of generated in-domain spoken language texts. Finally, we employ a sequence-sequence pretrained model (e.g, mT5) to translate spoken language texts generated by PGEN into sign glosses and synthesize gloss-text pseudo-parallel data (§2.2).

We conduct extensive analyses of the spoken language texts generated by PGEN. We find that the generated and original spoken language texts share a similar word distribution (§3.2). To further verify the effectiveness of PGEN, we also conduct back-translation experiments for the SLGT task with large-scale in-domain spoken language texts generated by PGEN (§3.3). Experimental results on three widely used benchmark datasets across languages and domains show that back-translation with spoken language texts generated by PGEN significantly outperforms the compared methods. Most importantly, as shown in Figure 1, when scaling the spoken language texts generated by PGEN in BT approach, the performance of the gloss-to-text translation task can achieve constant improvement while conventional data selection approach failed. The contributions of our work are summarized as follows:

- We propose a novel text generation approach, i.e., PGEN, to produce large-scale in-domain spoken language texts which share similar lin-

guistic properties as the original spoken language texts.

- We scale back-translation with the proposed PGEN approach and achieve significant and consistent improvements on three benchmark SLGT datasets.
- We release the code and the large-scale synthetic gloss-text datasets produced by the proposed approach to promote the research in sign language gloss translation field.

2 Methodology

The whole framework of this work includes two components, i.e., the PGEN method for in-domain text generation and the BT model (i.e., text-to-gloss translation) for constructing synthetic data. We will introduce more details for these two components in this section. For clarity, we provide the definition of SLGT (i.e., gloss-to-text translation task) and notations used throughout the paper as below.

Task Definition. Let X and Y denote the gloss annotations and spoken language texts, and \mathcal{X} and \mathcal{Y} represent the sentence sets of corresponding languages. The dataset of gloss-text pairs can be expressed as $\mathcal{D}_{g2t} = \{(\mathbf{x}^i, \mathbf{y}^i)\}_{i=1}^N$, where $\mathbf{x}^i \in \mathcal{X}$ is the annotations, $\mathbf{y}^i \in \mathcal{Y}$ is the spoken language sentence and N is the number of pairs. Given a sequence of gloss annotations, the task is to output the corresponding fluent and semantically equivalent sentence.

2.1 Prompt-Based Domain Text Generation

We exploit the large PLMs for domain text generation. Large PLMs have been successfully applied for text generation in NLP, such as text classification (Kumar et al., 2020) and medical dialogue summarization (Chintagunta et al., 2021), for two advantages: (1) PLMs are demonstrated to memorize the knowledge of their training data (Carlini et al., 2021), which usually covers different domains. With proper guidance (e.g., prompts), we can export the memorized sentences that belong to the same domain as the sign language text data. (2) Large PLMs are also able to generate abundant new sentences rather than only the memorized sentence in training data (Qiu et al., 2020). Therefore, we propose the **Prompt-based domain text Generation (PGEN)** approach to produce large-scale in-domain spoken language text

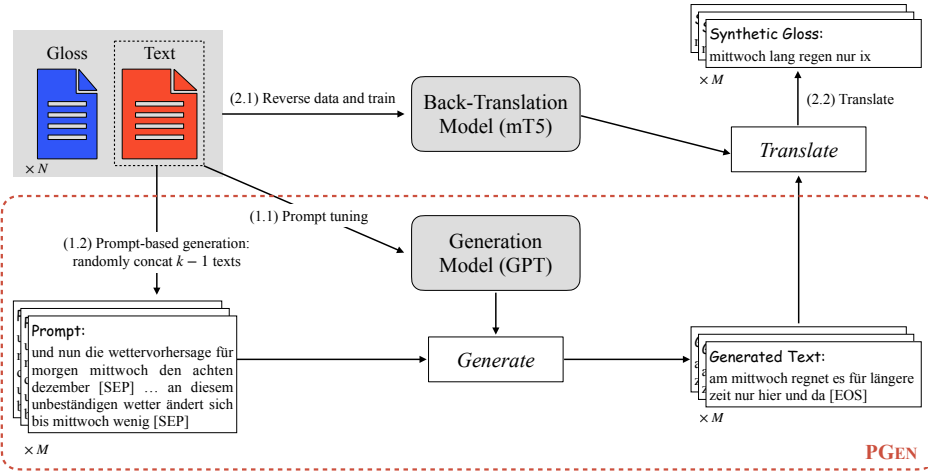


Figure 2: The overall framework of back-translation for gloss-to-text translation in this work. Procedures framed in the red dashed box corresponds to our prompt-based domain text generation (PGEN) approach.

corpora based on the text part of the original small-scale dataset \mathcal{D}_{g2t} .

For the original text $\mathcal{D}_t = \{\mathbf{y}^i\}_{i=1}^N$ from \mathcal{D}_{g2t} , and a PLM \mathcal{M}_G , we attempt to generate an in-domain spoken language text dataset $\hat{\mathcal{D}}_t = \{\hat{\mathbf{y}}^i\}_{i=1}^M$ with a much larger data size than \mathcal{D}_t (i.e., $M \gg N$). As shown in Figure 2, our PGEN approach includes two phases:

- **Prompt Tuning:** Following (Kumar et al., 2020), we finetune the PLM on the original small-scale spoken text dataset \mathcal{D}_t with artificial prompts. Specifically, we randomly concatenate k sentences from \mathcal{D}_t as a training sample, i.e., $[\mathbf{y}^{i_1}; [\text{SEP}]; \mathbf{y}^{i_2}; [\text{SEP}]; \dots; \mathbf{y}^{i_k}; [\text{EOS}]]$, where $[\text{SEP}]$ and $[\text{EOS}]$ represent the delimiter and the end-of-sentence tokens, respectively. We denote the finetuned PLM as $\mathcal{M}_{G_{\text{FT}}}$.
- **Prompt-Based Generation:** In the generation phase, we randomly select $k-1$ sentences from \mathcal{D}_t to form a prompt, i.e., $\text{prompt} = [\mathbf{y}^{j_1}; [\text{SEP}]; \mathbf{y}^{j_2}; [\text{SEP}]; \dots; \mathbf{y}^{j_{k-1}}; [\text{SEP}]]$. Then, we input the prompt into $\mathcal{M}_{G_{\text{FT}}}$ to generate the k -th sentence $\hat{\mathbf{y}}^{j_k} = \mathcal{M}_{G_{\text{FT}}}(\text{prompt})$. We complete the text generation process when the model produces an $[\text{EOS}]$ token.

According to the design of prompts, the number of permutations for any $k-1$ sentences from the full set with N sentences is $A_N^{k-1} = \frac{N!}{(N-k+1)!} \gg N$, which allows us to generate a large number of in-domain sentences and maintain the diversity of $\hat{\mathcal{D}}_t$.

2.2 Back-Translation

Generally, the BT model is trained on the same dataset for the gloss-to-text task but in the opposite direction, i.e., $\mathcal{D}_{t2g} = \{(\mathbf{y}^i, \mathbf{x}^i)\}_{i=1}^N$. However, the data scale of \mathcal{D}_{t2g} is too small to develop a well-performing text-to-gloss translation model. Previous study (Hoang et al., 2018) on machine translation also suggests that the quality of BT models heavily affects the performance of the final models. Therefore, we take advantage of pretrained sequence-to-sequence models by finetuning them on \mathcal{D}_{t2g} to improve the performance of the BT model. Specifically, we utilize a multilingual pretrained model, i.e., mT5 (Xue et al., 2020), to support the different languages (i.e., German, Chinese and English) of the SLT benchmarks.

2.3 Overall Framework

The workflow of our approach is illustrated in Figure 2, which is divided into four steps: (1.1) finetune the pretrained GPT-2 model on the original small-scale spoken language text dataset; (1.2) apply the finetuned GPT-2 model for in-domain spoken language texts generation with artificial prompts; (2.1) finetune the pretrained sequence-to-sequence model (i.e., mT5) on the original text-gloss dataset \mathcal{D}_{t2g} to obtain the BT model; (2.2) adopt the BT model to translate the generated in-domain spoken language texts into glosses to synthesize a large-scale pseudo-parallel data, which are combined with the original small-scale dataset \mathcal{D}_{g2t} to train the final gloss-to-text translation model.

Dataset	Language Pair	Gloss-Text Pairs
Phoenix2014T	DSL-German	7,086 / 519 / 642
CSL-Daily	CSL-Chinese	18,401 / 1,077 / 1,176
ASLG-PC12	ASL-English	82,709 / 4,000 / 1,000

Table 1: Statistics of the three benchmark datasets for gloss-to-text translation used in this work. The third column presents the number of gloss-text pairs in the training, validation and test sets, respectively.

3 Experiments

In this section, we conduct both intrinsic and extrinsic evaluations (Kumar et al., 2020) for the proposed PGEN approach. For intrinsic evaluation (§3.2), we perform analyses on the domain properties of spoken language texts generated by PGEN. As for extrinsic evaluation (§3.3), we conduct sign gloss translation experiments using back-translation approach with the generated texts. The performance of the downstream task can indirectly reflect the effectiveness of PGEN.

3.1 Experimental Setup

Dataset. We employ three widely used benchmark datasets for sign language translation, namely, Phoenix2014T (Camgoz et al., 2018), CSL-Daily² (Zhou et al., 2021a), and ASLG-PC12³ (Othman and Jemni, 2012), which are in German, Chinese and English, respectively. Statistics of the datasets are presented in Table 1.

Model. As shown in Figure 2, there are three kinds of models involved in this work for 1) in-domain text generation, 2) gloss-to-text translation and 3) back-translation, respectively. Details of the model and training settings can be found in §3.2, §3.3 and Appendix A.1.

3.2 Domain Text Generation

We perform intrinsic evaluation by analyzing the domain properties of spoken language texts generated by PGEN. Unless otherwise stated, we primarily conduct the analyses on the German Phoenix2014T dataset. More results on ASLG-PC12 and CSL-Daily datasets can be found in Appendix A.3.

We first adopt the pre-trained GPT-2 model (e.g., German GPT-2⁴ for Phoenix2014T) and finetune

²<http://home.ustc.edu.cn/~zhouh156/dataset/csl-daily/>

³<https://github.com/kayoyin/transformer-slt>

⁴<https://huggingface.co/dbmdz/german-gpt2>




Figure 3: The word frequency distribution on different types of spoken language texts. The X-axis represents different words, while the Y-axis represents the normalized word frequency. Best viewed in color. Analyses is conducted on the German Phoenix2014T

the model on the artificial prompts created from the text part of SLT dataset. Specifically, the PLM is finetuned to predict the next token in the exact way that GPT-2 was pretrained, with the same training procedure and hyper-parameters. Then, we use the finetuned PLM to generate in-domain spoken language texts for each SLT task (§2.1). By default, we set the hyper-parameter k to 20.

For comparison, we consider four types of spoken language texts:

- **TEXT-AUTHENTIC:** The text side of the SLT dataset.
- **TEXT-PGEN:** The spoken language texts generated by our PGEN approach with $k = 20$.
- **TEXT-SELECTED:** We use the cross-entropy difference selection method (Moore and Lewis, 2010) to collect in-domain texts data from IWSLT17 Multilingual Task⁵ based on the AUTHENTIC texts.
- **TEXT-GENERAL:** We randomly sample sentences from the target text part of IWSLT17 Multilingual Task. TEXT-GENERAL can be considered as the general domain spoken language texts.

We measure the similarity of the other three spoken language texts to the TEXT-AUTHENTIC at both word level and sentence level.

Word Distribution. Recent study by Wang et al. (2022) suggests that the word frequency distributions can reflect the domain difference of datasets.

⁵<https://sites.google.com/site/iwslt2017/TED-tasks>

Data	JS↓
TEXT-PGEN vs. TEXT-AUTHENTIC	0.01
TEXT-SELECTED vs. TEXT-AUTHENTIC	0.18
TEXT-GENERAL vs. TEXT-AUTHENTIC	0.26

Table 2: The JS divergence between different types of spoken language texts for Phoenix2014T dataset.

Thus, we count the word frequencies for the above four spoken language texts and present the results by both visualization and the Jensen-Shannon (JS) divergence (Lin, 1991). For a fair comparison, we sample the same number of texts from the other three spoken language texts as TEXT-AUTHENTIC (i.e., 7086 for Phoenix2014T).

Figure 3 visualizes the word frequency distributions of top 10000, in which the words are ranked according to their frequencies in the TEXT-AUTHENTIC corpus. We can observe that TEXT-PGEN (blue line) shows a similar distribution as TEXT-AUTHENTIC (red line) while TEXT-SELECTED and TEXT-GENERAL differ from TEXT-AUTHENTIC significantly. This result qualitatively shows that the spoken language texts generated by our PGEN approach are more close to the domain of TEXT-AUTHENTIC (e.g. the domain of sign language).

To quantitatively measure the distance between these distributions, we compute the JS divergence expressed as:

$$JS(P||Q) = \frac{1}{2} \left(KL(P||\frac{P+Q}{2}) + KL(Q||\frac{P+Q}{2}) \right),$$

where $KL(\cdot||\cdot)$ denotes the Kullback–Leibler divergence (Kullback and Leibler, 1951) of two distributions (i.e., P and Q). Table 2 lists the JS divergence from the other three corpora to TEXT-AUTHENTIC. We find that the JS divergence from TEXT-PGEN to TEXT-AUTHENTIC is much smaller than the others, further demonstrating that the spoken language texts generated by PGEN are closer to the domain of TEXT-AUTHENTIC. These demonstrate the effectiveness and generalizability of the proposed PGEN approach.

Domain Classifier. The word frequency distribution only characterizes one aspect of the domain. More features, for example, the styles of texts, can not be explicitly modeled. Thus, we follow Du et al. (2020) to train a domain classifier to distinguish the in-domain and the general domain data with

Test Data	In-domain	General
TEXT-AUTHENTIC	99.38%	0.62%
TEXT-PGEN	98.60%	1.40%
TEXT-SELECTED	56.23%	43.77%
TEXT-GENERAL	0.31%	99.69%

Table 3: The domain classification results of different spoken language texts on Phoenix2014T dataset.

the consideration of all potential features implicitly. We perform a binary classification task with equal examples (i.e., train/valid as 7086/519) from TEXT-GENERAL and TEXT-AUTHENTIC, respectively. To train the domain classifier, we finetune the German BERT⁶ on the above dataset (Sun et al., 2019). Specifically, we use the German BERT to encode the input sentence and feed the [CLS] token vector as a reasonable sentence embedding to the domain discriminator. For testing, we also sample the same number of examples from the other three spoken language text corpora as the test set of TEXT-AUTHENTIC (i.e., 642) and predict their domains (i.e., general or in-domain).

The results are listed in Table 3. We observe that the domain classifier successfully predicts the true labels of TEXT-AUTHENTIC and TEXT-GENERAL, indicating the significant domain differences between sign language spoken texts and general spoken language texts. As for TEXT-PGEN, the examples are categorized into authentic texts with very high accuracy (i.e., 98.60%), while the value is much lower for TEXT-SELECTED (i.e., 56.23%). These results again demonstrate that our PGEN approach can produce better in-domain spoken language texts than the compared methods.

3.3 Sign Language Gloss Translation

We perform extrinsic evaluation for the proposed PGEN by applying back-translation with the generated spoken language texts to the sign language gloss translation task. For the different SLT tasks in 1, we adopt the corresponding German⁷, Chinese⁸ and English⁹ GPT-2 models to generate in-domain spoken language texts with PGEN.

For the gloss-to-text translation task, we follow

⁶<https://huggingface.co/bert-base-german-cased>

⁷<https://huggingface.co/dbmdz/german-gpt2>

⁸<https://huggingface.co/uer/gpt2-chinese-cluecorpussmall>

⁹<https://huggingface.co/gpt2>

	Dev Set				Test Set			
	BLEU-1	BLEU-2	BLEU-3	BLEU-4	BLEU-1	BLEU-2	BLEU-3	BLEU-4
<i>Phoenix2014T</i>								
Camgoz et al. (2018)	44.40	31.83	24.61	20.16	44.13	31.47	23.89	19.26
Camgoz et al. (2020b)	50.69	38.16	30.53	25.35	48.90	36.88	29.45	24.54
Yin and Read (2020)	49.05	36.20	28.53	23.52	47.69	35.52	28.17	23.32
Transformer	43.05	32.57	25.50	20.81	43.71	33.40	26.45	21.73
+ Scaling BT(General)	43.52	32.42	25.30	20.62	42.68	32.06	25.09	20.40
+ Scaling BT(Selected)	44.20	33.02	25.86	21.06	44.29	33.25	25.97	21.09
+ Scaling BT(PGen)	48.68	37.94	30.58	25.56	48.30	37.59	30.32	25.54
<i>ASLG-PC12</i>								
Yin and Read (2020)	92.67	88.72	85.22	81.93	92.88	89.22	85.95	82.87
Transformer	91.85	87.53	83.73	80.19	92.04	88.07	84.56	81.25
+ Scaling BT(General)	91.59	87.43	83.81	80.45	91.97	88.19	84.89	81.79
+ Scaling BT(PGen)	93.23	88.91	85.63	82.04	93.51	89.74	86.55	83.35
<i>CSL-Daily</i>								
Transformer	49.63	35.62	25.52	18.64	49.41	35.57	25.55	18.72
+ Scaling BT(General)	54.66	39.80	29.23	21.78	54.07	39.34	28.99	21.75
+ Scaling BT(PGen)	60.48	46.92	36.95	29.72	60.21	46.76	36.90	29.75

Table 4: Gloss-to-text translation performance on Phoenix2014T, ASLG-PC12 and CSL-daily. "+ Scaling BT(PGen)" represents that training data is increased by 40 times with BT, in which the monolingual is generated by our PGen approach.

Yin and Read (2020) to train a Transformer model with 2 encoder layers and 2 decoder layers. For back-translation, we first finetune the mT5 pre-trained model on the authentic text-gloss parallel data and then use it to translate the collected spoken language texts (sampled from TEXT-PGEN, TEXT-SELECTED or TEXT-GENERAL) into glosses. Following Wu et al. (2019), we train the Transformer model on the combination of the authentic and synthetic parallel data, and then finetune it on the authentic gloss-to-text parallel data.

For evaluation, we follow previous studies (Camgoz et al., 2018, 2020b; Zhou et al., 2021a) to evaluate the performance of gloss-to-text translation with BLEU score (Papineni et al., 2002), ROUGE-L (Lin, 2004), and METEOR (Banerjee and Lavie, 2005) scores. Specifically, we report the BLEU-1,2,3,4 scores to reflect the translation quality at different phrase levels.

Table 4 lists the main results of the gloss-to-text translation performance on the Phoenix2014T, ASLG-PC12 and CSL-daily datasets. By scaling the BT synthetic data to 40 times of the authentic parallel data with the different types of spoken text data, TEXT-PGEN (i.e., "+ Scaling BT(PGen)") improves the performance over the baseline Transformer model significantly and consistently (e.g.,

up to 11.03 BLEU-4 points on CSL-daily), while TEXT-GENERAL and TEXT-SELECTED only improve slightly or even hurt the performances (e.g., down to 1.33 BLEU-4 points on Phoenix2014). We also report the translation performance in terms of ROUGE-L and METEOR in Appendix A.4. These demonstrate the effectiveness and generalizability of the proposed PGEN approach.

4 Analysis

To gain a deeper understanding on our PGEN approach, we conduct extensive analyses in terms of ablation study and translation outputs.

4.1 Ablation Study

We conduct three ablation studies regarding the prompt length, the back-translation model and the scale of synthetic data on the Phoenix2014T dataset. We introduce more details as below:

Prompt Length. We first investigate the impact of the prompt length k on the gloss-to-text translation task, which decides how many sentences are concatenated as prompts for generation via PGEN. Specifically, we increase k from 0 to 35 to induce the PLM to generate in-domain spoken language texts with a data size 5 times of the authentic par-

	Text-to-Gloss Dev Set				Gloss-to-Text Dev Set			
	BLEU-1	BLEU-2	BLEU-3	BLEU-4	BLEU-1	BLEU-2	BLEU-3	BLEU-4
Transformer	–	–	–	–	43.05	32.57	25.50	20.81
+ BT Model SC	57.72	39.41	27.86	19.66	44.80	34.33	27.47	22.86
+ BT Model mT5	59.78	43.70	32.58	25.01	46.87	35.95	28.59	23.54

Table 5: Translation performance with different back-translation models on text-to-gloss and gloss-to-text tasks.




Figure 4: Gloss-to-text translation performance with respect to different prompt sizes. The X-axis represents the number of sentences in one prefix, while the Y-axis represents the BLEU-4 score.

allel data. Then, we perform back-translation with the generated spoken language texts on the gloss-to-text translation task, individually. The results are shown in Figure 4, in which we observe that the performance of gloss-to-text translation constantly improves with the increase of prompt length. This is because a larger prompt length can provide more domain signals so as to encourage the generation of higher-quality spoken language texts with closer domain and higher diversity (see Appendix A.2). However, the larger prompt length requires more computation memory, slowing down the generation process. Therefore, we set the prompt length to 20 throughout the work for a good tradeoff between the generation quality and the computation costs.

Back-Translation Model. In §2.2, we state that the quality of BT models heavily affects the performance of final models. To validate this claim, we compare two back-translation models: a Transformer model trained from scratch (i.e., “+ BT Model SC”) and the finetuned mT5 (Xue et al., 2020) model (i.e., “+ BT Model mT5”). As shown in the left of Table 5, the finetuned mT5 model produces higher-quality pseudo parallel data according to BLEU scores on the validation set of the text-to-

gloss translation (i.e., back-translation) task. Consequently, the performance of gloss-to-text translation is considerably improved when synthesizing data by “+ BT Model mT5”, which reconfirms our claim. Therefore, throughout this work, we adopt the finetuned mT5 model for back-translation.

Scale of Synthetic Data. In §2.1, we show the potential of PGEN in generating large-scale in-domain spoken language texts. Let us recap Figure 1, where we increase the scale of spoken language texts used for back-translation. We observe that scaling the spoken language texts generated by PGEN can improve the performance of the gloss-to-text translation task consistently while that by retrieval degrades the performance. It suggests that our approach can scale the BT technique to play its maximum effect for gloss-to-text translation, which has not been achieved in previous studies.

4.2 Translation Output

We conduct further analyses to understand how the proposed approach improves the gloss-to-text translation quality. Specifically, we analyze the translation outputs of Phoenix2014 in Table 4 by the compare-*mt*¹⁰ toolkit in terms of word frequency and sentence length.

Words Frequency. Previous study (Fadaee and Monz, 2018) shows that the back-translation improves the translation performance by improving the low-frequency word predictions. Meanwhile, our analyses in §3.2 suggest that the spoken language texts generated by our PGEN shows a similar word frequency distribution as TEXT-AUTHENTIC. We wonder how such consistency benefits the prediction of low-frequency words in gloss-to-text translation. Specifically, we first categorize the vocabulary into three groups based on the word frequency in the training data, including High: frequency $\in [2000, +\infty)$; Medium: frequency $\in [100, 2000)$; Low: frequency $\in (0, 100]$. Then,

¹⁰<https://github.com/neulab/compare-mt>

Data	Word Frequency		
	Low	Medium	High
TEXT-AUTHENTIC	28.86	49.94	58.27
+ TEXT-PGEN	33.23	52.98	60.17
+ TEXT-SELECTED	24.65	49.87	58.58
+ TEXT-GENERAL	25.45	49.12	58.12

Table 6: Prediction accuracy (F1 score) of target words in the test set with respect to word frequency. As the higher F1 score, the better, we mark the improvement by green and degradation by red background.

Data	Sentence Length		
	Short	Medium	Long
TEXT-AUTHENTIC	21.73	22.61	10.00
+ TEXT-PGEN	21.87	27.29	20.15
+ TEXT-SELECTED	16.24	23.93	12.80
+ TEXT-GENERAL	16.68	23.97	9.47

Table 7: Translation quality (BLEU score) of examples in the test set with respect to sentence length.

we utilize compare-mt to calculate the prediction accuracy of target words in the test set with respect to the three groups.

Table 6 lists the results for different models with scaling back-translation (e.g. 40 times synthetic data). As seen, scaling back-translation with in-domain spoken language texts of TEXT-PGEN improves the prediction of words in all three groups, especially for low-frequency words. However, the situation is much different for TEXT-SELECTED and TEXT-GENERAL such that they bring little improvement for high-frequency words and inversely harm the performance on low-frequency words. It indicates that back-translation becomes ineffective when the domains of spoken language texts and authentic parallel data are mismatched, which implies the importance of our PGEN approach.

Sentence Length. We investigate the translation quality of examples with varied lengths, which can be biased during generating or retrieving spoken language texts. Similar to word frequency, we also categorize the examples of test set into three groups based on the sentence length, including Long: (20, $+\infty$) tokens; Medium: (10, 20] tokens; Short: (0, 10] tokens.

Table 7 lists the corresponding results. Clearly, long sentences are more difficult to translate (Zheng et al., 2020) and our PGEN can par-

ticularly improve the translation quality of medium and long sentences. In contrast, the other methods show little improvement on medium and long sentences and degrade the performance on short sentences significantly. This demonstrates the better stability of our approach regarding the distribution of sentence length over the compared methods.

5 Related Work

Sign Language Gloss Translation. SLGT translates sign gloss to spoken language texts, which has attracted more attention in recent years with the development of neural machine translation (NMT). For example, Camgoz et al. (2018) released the PHOENIX14T and for the first time proposed a neural SLT model to translate from spatial representations or sign glosses. Recent studies attempt to improve both SLR and gloss-to-text translation for the better performance of SLT. Yin and Read (2020) proposed the STMC-Transformer network (Zhou et al., 2020) to improve SLR and exploited Transformer for gloss-to-text translation. Camgoz et al. (2020b) formulated SLR and gloss-to-text translation in the multi-task form while Li et al. (2020) explored the hierarchical structure for learning sign video representations. More recently, multi-cue characteristics of sign language have also been utilized for improving SLT (Camgoz et al., 2020a; Zhou et al., 2021b; Kan et al., 2022). Different from these works, we improve SLT by focusing on the gloss-to-text translation task in the perspective of spoken language generation.

Data Augmentation. Data augmentation has been proposed and proven valuable and effective in machine translation research (Sennrich et al., 2016; Zhang and Zong, 2016; Wang et al., 2018; Jiao et al., 2020, 2021). To address the data scarcity issue in gloss-to-text translation, there have been studies on producing synthetic gloss-text pairs for data augmentation. One stream is to extract discrete phrases from natural texts based on linguistic rules (Moryossef et al., 2021). Another is to adopt BT technique (Sennrich et al., 2016) to generate glosses from natural texts by a pretrained text-to-gloss translation model. However, the limited in-domain spoken language texts prevent BT from playing its maximum effect for the gloss-to-text translation task. While we may collect in-domain texts from public websites as Zhou et al. (2021a), it is unreliable due to both the requirement for domain knowledge and the accessibility

of websites. Our approach exploits the knowledge and generalization capability of large PLMs to produce in-domain texts with prompt learning. While large PLMs have been successfully applied for text generation in NLP (Kumar et al., 2020; Chintagunta et al., 2021), we craft the prompts to theoretically guarantee that we can produce large-scale in-domain texts based on the small-scale original spoken language texts data.

6 Conclusion

In this paper, we propose the PGEN approach to produce large-scale in-domain monolingual texts based on the original small-scale texts of the gloss-text parallel data. With PGEN, we scale back-translation and achieve significant and consistent improvements on three benchmark datasets for SLGT task. Extensive analyses suggest that our approach generates monolingual texts with similar linguistic properties as the original monolingual texts, thus outperforms the compared methods in terms of both low-frequency word prediction and long sentence translation. Future work includes exploring ChatGPT for sign language translation task by using proper prompts (Jiao et al., 2023).

7 Limitations

We identify two limitations of our PGEN approach:

- We rely on the seed in-domain spoken language texts and the pretrained models. Therefore, a bit of expert knowledge is still required for the selection of the seed dataset, and the pretrained models need to be carefully selected to ensure a large data coverage and a strong generalization ability.
- We only investigate the sign language translation without considering the sign language understanding part. Since the latter part may introduce errors into glosses, it requires the sign gloss translation models to be robust to noisy inputs. We leave this for the future work.

References

Amittai Axelrod, Xiaodong He, and Jianfeng Gao. 2011. Domain adaptation via pseudo in-domain data selection. In *Proceedings of the 2011 conference on empirical methods in natural language processing*.

Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved

correlation with human judgments. In *Proceedings of the ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*.

Danielle Bragg, Oscar Koller, Mary Bellard, Larwan Berke, Patrick Boudreault, Annelies Braffort, Naomi Caselli, Matt Huenerfauth, Hernisa Kacorri, Tessa Verhoef, et al. 2019. Sign language recognition, generation, and translation: An interdisciplinary perspective. In *Proceedings of the 21st international ACM SIGACCESS conference on computers and accessibility*.

Jan Bungeroth and Hermann Ney. 2004. Statistical sign language translation. In *Proceedings of the Workshop on representation and processing of sign languages, Language Resources and Evaluation Conference*.

Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, Hermann Ney, and Richard Bowden. 2018. Neural sign language translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.

Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. 2020a. Multi-channel transformers for multi-articulatory sign language translation. In *Proceedings of the European Conference on Computer Vision*. Springer.

Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. 2020b. Sign language transformers: Joint end-to-end sign language recognition and translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.

Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom B Brown, Dawn Song, Ulfar Erlingsson, et al. 2021. Extracting training data from large language models. In *Proceedings of the USENIX Security Symposium*.

Bharath Chintagunta, Namit Katariya, Xavier Amatriain, and Anitha Kannan. 2021. Medically aware gpt-3 as a data generator for medical dialogue summarization. In *Proceedings of Machine Learning for Healthcare Conference*.

Mathieu De Coster, Dimitar Shterionov, Mieke Van Herreweghe, and Joni Dambre. 2022. Machine translation from signed to spoken languages: State of the art and challenges. *arXiv*.

Chunning Du, Haifeng Sun, Jingyu Wang, Qi Qi, and Jianxin Liao. 2020. Adversarial and domain-aware bert for cross-domain sentiment analysis. In *Proceedings of the 58th annual meeting of the Association for Computational Linguistics*.

Marzieh Fadaee and Christof Monz. 2018. Back-translation sampling by targeting difficult words in neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural*

- Language Processing*, Brussels, Belgium. Association for Computational Linguistics.
- Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. Iterative back-translation for neural machine translation. In *Proceedings of the 2nd workshop on neural machine translation and generation*.
- Long Jiang, Shiquan Yang, Ming Zhou, Xiaohua Liu, and Qingsheng Zhu. 2009. Mining bilingual data from the web with adaptively learnt patterns. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*.
- Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. 2023. Is ChatGPT a good translator? A preliminary study. *arXiv*.
- Wenxiang Jiao, Xing Wang, Shilin He, Irwin King, Michael Lyu, and Zhaopeng Tu. 2020. Data rejuvenation: Exploiting inactive training examples for neural machine translation. In *Proceedings of the Empirical Methods in Natural Language Processing*.
- Wenxiang Jiao, Xing Wang, Zhaopeng Tu, Shuming Shi, Michael Lyu, and Irwin King. 2021. Self-training sampling with monolingual data uncertainty for neural machine translation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*.
- Jichao Kan, Kun Hu, Markus Hagenbuchner, Ah Chung Tsoi, Mohammed Bennamoun, and Zhiyong Wang. 2022. Sign language translation with hierarchical spatio-temporal graph neural network. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representation*.
- Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics*.
- Varun Kumar, Ashutosh Choudhary, and Eunah Cho. 2020. Data augmentation using pre-trained transformer models. In *Proceedings of the 2nd Workshop on Life-long Learning for Spoken Language Systems*.
- Dongxu Li, Chenchen Xu, Xin Yu, Kaihao Zhang, Benjamin Swift, Hanna Suominen, and Hongdong Li. 2020. Tspnet: Hierarchical feature learning via temporal semantic pyramid for sign language translation. *Advances in Neural Information Processing Systems*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Proceedings of the Text summarization branches out*.
- Jianhua Lin. 1991. Divergence measures based on the shannon entropy. *IEEE Transactions on Information theory*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv*.
- Robert C Moore and William Lewis. 2010. Intelligent selection of language model training data. In *Proceedings of the Association for Computational Linguistics*.
- Amit Moryossef, Kayo Yin, Graham Neubig, and Yoav Goldberg. 2021. Data augmentation for sign language gloss translation. In *Proceedings of the 1st International Workshop on Automatic Translation for Signed and Spoken Languages (AT4SSL)*.
- Achraf Othman and Mohamed Jemni. 2012. English-asl gloss parallel corpus 2012: Aslg-pc12. In *Proceedings of the 5th Workshop on the Representation and Processing of Sign Languages: Interactions between Corpus and Lexicon LREC*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*.
- Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. 2020. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Dimitar Shterionov, editor. 2021. *Proceedings of the 1st International Workshop on Automatic Translation for Signed and Spoken Languages (AT4SSL)*. Association for Machine Translation in the Americas, Virtual.
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to fine-tune bert for text classification? In *Proceedings of the China national conference on Chinese computational linguistics*. Springer.
- Wenxuan Wang, Wenxiang Jiao, Yongchang Hao, Xing Wang, Shuming Shi, Zhaopeng Tu, and Michael Lyu. 2022. Understanding and improving sequence-to-sequence pretraining for neural machine translation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*.

- Xinyi Wang, Hieu Pham, Zihang Dai, and Graham Neubig. 2018. Switchout: an efficient data augmentation algorithm for neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Lijun Wu, Yiren Wang, Yingce Xia, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2019. Exploiting monolingual data at scale for neural machine translation. 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)*.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2020. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv*.
- Kayo Yin, Amit Moryossef, Julie Hochgesang, Yoav Goldberg, and Malihe Alikhani. 2021. Including signed languages in natural language processing. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*.
- Kayo Yin and Jesse Read. 2020. Better sign language translation with stmc-transformer. In *Proceedings of the 28th International Conference on Computational Linguistics*.
- Jiajun Zhang and Chengqing Zong. 2016. Exploiting source-side monolingual data in neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Xuan Zhang and Kevin Duh. 2021. Approaching sign language gloss translation as a low-resource machine translation task. In *Proceedings of the 1st International Workshop on Automatic Translation for Signed and Spoken Languages (AT4SSL)*.
- Jiangbin Zheng, Zheng Zhao, Min Chen, Jing Chen, Chong Wu, Yidong Chen, Xiaodong Shi, and Yiqi Tong. 2020. An improved sign language translation model with explainable adaptations for processing long sign sentences. *Computational Intelligence and Neuroscience*.
- Hao Zhou, Wengang Zhou, Weizhen Qi, Junfu Pu, and Houqiang Li. 2021a. Improving sign language translation with monolingual data by sign back-translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Hao Zhou, Wengang Zhou, Yun Zhou, and Houqiang Li. 2020. Spatial-temporal multi-cue network for continuous sign language recognition. In *Proceedings of the Association for the Advancement of Artificial Intelligence*.
- Hao Zhou, Wengang Zhou, Yun Zhou, and Houqiang Li. 2021b. Spatial-temporal multi-cue network for sign language recognition and translation. *IEEE Transactions on Multimedia*.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. Texus: A benchmarking platform for text generation models. In *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*.

A Appendix

A.1 Model Training

We perform extrinsic evaluation for the proposed PGEN by applying back-translation with the generated spoken language texts to the sign language gloss translation task. For the Chinese and English SLT tasks, we adopt the corresponding English¹¹ and Chinese¹² GPT-2 models to generate in-domain spoken language texts with PGEN.

For the gloss-to-text translation task, we follow Yin and Read (2020) to train a Transformer model with 2 encoder layers and 2 decoder layers. We use Adam (Kingma and Ba, 2015) with $\beta = (0.9, 0.98)$ and $\epsilon = 10^{-6}$ for optimization. We adopt the warm-up learning rate scheduler, which linearly increases from 1.0×10^{-4} to a peak of 5.0×10^{-4} within 2000 steps, and then decays with the inverse square root schedule. The dropout rate is 0.3 and the label smoothing is 0.1.

For back-translation, we first finetune the mT5 pre-trained model on the authentic text-gloss parallel data and then use it to translate the collected spoken language texts (sampled from TEXT-PGEN or TEXT-SELECTED) into glosses. Following Wu et al. (2019), we train the Transformer model on the combination of the authentic and synthetic parallel data, and then finetune it on the authentic gloss-to-text parallel data.

Table 8 presents the hyper-parameters of different transformer models used in this work.




Figure 5: Diversity and domain closeness of sign language texts generated by PGEN measured by Self-BLEU and JS with respect to different prompt sizes.

¹¹<https://huggingface.co/gpt2>

¹²<https://huggingface.co/uer/gpt2-chinese-cluecorpussmall>

A.2 Diversity and Quality

Figure 5 shows the diversity and domain closeness of sign language texts generated by PGEN measured by Self-BLEU and JS scores with respect to prompt length. Lower Self-BLEU (Zhu et al., 2018) and JS scores indicate higher diversity and closer domain, respectively.

A.3 Intrinsic Analysis Results for ASLG-PC12 and CSL-Daily Datasets

We extend the intrinsic analyses to both ASLG-PC12 and CSL-Daily datasets. Table 9 shows the JS divergence between different types of spoken language texts on the two SLT datasets. As seen, the JS divergence from TEXT-PGEN to TEXT-AUTHENTIC is much smaller than TEXT-GENERAL to TEXT-AUTHENTIC. Table 10 lists the domain classification results on the ASLG-PC12 and CSL-Daily datasets. The results indicate the significant domain differences between sign language spoken texts and general spoken language texts, and our PGEN approach can produce in-domain spoken language texts.

Clearly, the results on both ASLG-PC12 and CSL-Daily datasets are consistent with that in section 3.2, which demonstrates the effectiveness and generalizability of the proposed PGEN approach.

A.4 Other Metrics

Table 11 presents the ROUGE-L and METEOR scores of the gloss-to-text translation performance on the Phoenix2014T, ASLG-PC12 and CSL-daily datasets.

Parameter	Text-to-Gloss		Gloss-to-Text	
	Transformer	mT5	Pretrain	Finetune
encoder-layers	2	12	6	6
decoder-layers	2	12	6	6
learning rate	$7 \cdot 10^{-4}$	$7 \cdot 10^{-4}$	$7 \cdot 10^{-4}$	$7 \cdot 10^{-4}$
learning rate scheduler	inverse_sqrt	inverse_sqrt	inverse_sqrt	inverse_sqrt
Adam β	(0.9, 0.98)	-	(0.9, 0.98)	(0.9, 0.98)
warmup-updates	2000	-	4000	1000
label-smoothing	0.1	0.1	0.1	0.1
dropout	0.3	0.1	0.1	0.5
batch-size	2048	2048	2048	2048

Table 8: Hyperparameters of translation models.

Data		ASLG-PC12	CSL-Daily
TEXT-PGEN	vs. TEXT-AUTHENTIC	0.02	0.08
TEXT-GENERAL	vs. TEXT-AUTHENTIC	0.18	0.14

Table 9: The JS divergence results on the ASLG-PC12 and CSL-Daily datasets.

Test Data	In-domain	General	Test Data	In-domain	General
TEXT-AUTHENTIC	98.63%	1.37%	TEXT-AUTHENTIC	99.86%	0.14%
TEXT-PGEN	99.44%	0.56%	TEXT-PGEN	98.57%	1.43%
TEXT-GENERAL	1.58%	98.42%	TEXT-GENERAL	0.24%	99.76%

ASLG-PC12

CSL-Daily

Table 10: The domain classification results on the ASLG-PC12 and CSL-Daily datasets.

	Dev Set		Test Set	
	ROUGE-L	METEOR	ROUGE-L	METEOR
<i>Phoenix2014T</i>				
Camgoz et al. (2018)	46.02	-	45.45	-
Yin and Read (2020)	47.36	46.09	46.58	44.85
Transformer	47.77	43.46	47.48	42.36
+ Scaling BT(General)	46.43	42.35	46.12	42.24
+ Scaling BT(Selected)	48.16	43.35	47.95	42.96
+ Scaling BT(PGen)	50.89	45.50	49.25	44.78
<i>ASLG-PC12</i>				
Yin and Read (2020)	82.41	95.93	95.87	96.46
Transformer	91.45	92.85	94.74	95.30
+ Scaling BT(General)	91.65	92.95	94.96	95.50
+ Scaling BT(PGen)	94.82	96.79	96.43	96.79
<i>CSL-Daily</i>				
Transformer	40.94	23.78	40.87	23.53
+ Scaling BT(General)	55.76	36.28	44.64	35.62
+ Scaling BT(PGen)	62.31	50.48	60.54	50.35

Table 11: Gloss-to-text translation performance on Phoenix2014T, ASLG-PC12 and CSL-daily. "+ Scaling BT(PGen)" represents that training data is increased by 40 times with BT, in which the monolingual is generated by our PGen approach.