

LLMs are Good Sign Language Translators

Jia Gong^{1†} Lin Geng Foo^{1†} Yixuan He^{1†} Hossein Rahmani² Jun Liu^{1‡}

¹Singapore University of Technology and Design ²Lancaster University

{jia.gong, lingeng.foo, yixuan.he}@mymail.sutd.edu.sg,
h.rahmani@lancaster.ac.uk, jun.liu@sutd.edu.sg

Abstract

Sign Language Translation (SLT) is a challenging task that aims to translate sign videos into spoken language. Inspired by the strong translation capabilities of large language models (LLMs) that are trained on extensive multilingual text corpora, we aim to harness off-the-shelf LLMs to handle SLT. In this paper, we regularize the sign videos to embody linguistic characteristics of spoken language, and propose a novel SignLLM framework to transform sign videos into a language-like representation for improved readability by off-the-shelf LLMs. SignLLM comprises two key modules: (1) The Vector-Quantized Visual Sign module converts sign videos into a sequence of discrete character-level sign tokens, and (2) the Codebook Reconstruction and Alignment module converts these character-level tokens into word-level sign representations using an optimal transport formulation. A sign-text alignment loss further bridges the gap between sign and text tokens, enhancing semantic compatibility. We achieve state-of-the-art gloss-free results on two widely-used SLT benchmarks.

1. Introduction

Sign languages, which are visual signals expressed through hand, body, and facial movements, serve as the primary means of communication within the hearing-impaired community. In an effort to facilitate effective communication with this community, much attention has been directed towards developing techniques to tackle the Sign Language Translation (SLT) task [10, 11, 69, 72, 73], where the goal is to translate sign videos into spoken language. SLT is a challenging task that requires cross-modality understanding of visual and linguistic cues [69, 72, 73], and the challenge is exacerbated by the limited availability of paired sign-text data [11, 13, 68, 69, 72]. Despite the notable advancements in terms of network architectures [10, 11, 37], visual sign representations (e.g., with keypoint estimators [14, 74]),

and training methods [22, 69, 72], how to effectively tackle the challenging cross-modal SLT task with limited paired sign-text data largely remains an open question.

On the other hand, large language models (LLMs) – referring to language models that have been trained on a large web-scale text corpus – have recently received a lot of attention. Since LLMs are trained over a very large corpus across multiple languages with distinct syntax and lexicon, they possess rich semantic understanding and powerful linguistic abilities [9, 15, 58]. At the same time, LLMs have also demonstrated an impressive capability to translate across multiple languages [9, 15], even showing a strong potential for translating languages with limited data [66, 76]. The foundation for this translation proficiency lies in the shared linguistic properties of syntax, lexicon, and morphology that many languages hold, which is particularly evident within language families [3]. Therefore, when faced with a new language with limited data, LLMs can draw upon the wealth of knowledge acquired from previously learned languages, leveraging any shared properties of syntax, lexicon and morphology with previous languages to effectively generate translations for new languages with remarkable accuracy and fluency [66, 76].

Inspired by the impressive translation capabilities of LLMs, we aim to harness off-the-shelf LLMs to handle the challenging SLT task. However, training LLMs directly on the relatively small SLT dataset can potentially lead to forgetting of their rich knowledge [12, 29] and a decline in performance, thus we follow previous LLM-based works [24, 43, 54] to keep the off-the-shelf LLM frozen, which preserves the rich knowledge acquired during its pre-training on a vast multilingual corpus. Consequently, our focus shifts towards making the sign videos compatible and readable for the off-the-shelf and frozen LLM to perform SLT. Specifically, in this paper, we explore the following question: *Can we treat the sign video as a form of language, and leverage an off-the-shelf and frozen LLM to translate them?* Notably, this is not a straightforward task because directly encoding features from sign videos with a pre-trained feature extractor [28, 64] will result in a large gap between

† Equal contribution; ‡ Corresponding author

the sign video features and text tokens, making it difficult for off-the-shelf LLMs to understand them.

Based on the observation that LLMs can effectively handle new languages by leveraging shared commonalities with previously learned languages, we aim to introduce designs that transform our sign videos into a language-like format which are readable and friendly to LLMs. Specifically, we hypothesize that providing language-like representations of sign videos to the LLM improves the LLM’s understanding of the sign videos and facilitates greater exploitation of shared properties with previously learned languages, thus resulting in better SLT performance by the LLM. To obtain language-like sign video representations, we draw inspiration from linguistic studies and analyses on LLMs [30, 63] and regularize the sign video to embody two fundamental language-like characteristics: **Discrete Characteristics:** Spoken languages are inherently discrete, since each language contains a finite set of words (and subwords) that convey distinct concepts, allowing them to be naturally represented through a discrete vocabulary with distinct tokens [8, 59]. **Hierarchical Structure:** Most spoken languages exhibit three hierarchical semantic levels – the sentence, word, and character levels [36, 52]. This hierarchical structure enables languages to express a wide range of words with a limited set of characters, and convey diverse sentences with a limited number of words.

In this paper, we present SignLLM, a novel framework designed to regularize input sign videos to produce sign token representations with language-like characteristics that are compatible and friendly to LLMs. Our proposed SignLLM includes two key designs to impart discrete characteristics and a hierarchical structure to the produced sign tokens. Firstly, we introduce the Vector-Quantized Visual Sign (**VQ-Sign**) module that facilitates the conversion of sign videos into a *sequence of discrete character-level sign tokens*. To achieve this, the VQ-Sign module consists of a discrete character-level sign codebook which is optimized through a self-supervised context prediction task. Next, we introduce the Codebook Reconstruction and Alignment (**CRA**) module that converts the character-level sign tokens into *word-level sign tokens*, facilitated by an optimal transport formulation. Moreover, we employ a sign-text alignment loss to further narrow the gap between the sign tokens and text tokens. These designs enable SignLLM to produce sign sentences that embody two key characteristics of spoken languages: discrete characteristics and a hierarchical structure, which enhances their compatibility with LLMs and makes them more readily interpretable by LLMs.

After producing the language-like sign sentences, we feed them into an off-the-shelf and frozen LLM along with a text prompt that instructs the LLM to generate translations in the desired language. We empirically observe that, by employing our SignLLM’s designs to align sign videos with

languages, we can already leverage a frozen LLM to attain state-of-the-art SLT performance. These findings suggest that our proposed SignLLM framework is a promising first step towards effectively harnessing LLMs for SLT. We hope our initial explorations can inspire future work within the community to leverage LLMs for SLT.

In summary, our main contributions are: (1) We propose a novel SignLLM framework that is the first to harness the power of off-the-shelf and frozen LLMs for SLT. (2) To make the input sign video compatible with LLMs, our SignLLM framework incorporates two designs: a VQ-Sign module to quantize the sign video into a sequence of discrete character-level sign tokens and a CRA module that transforms the character-level sign tokens to word-level sign tokens. (3) Through our proposed designs, we achieve state-of-the-art gloss-free results on two popular SLT datasets.

2. Related Work

Sign Language Translation (SLT) aims to transform sign videos into natural language sentences. It is a challenging task that requires understanding of both visual and linguistic cues [69, 72, 73], and the challenge is exacerbated by the limited availability of paired sign-text data [11, 13, 68, 69, 72] which limits the performance of SLT methods. To improve SLT performance, many previous works [10, 11, 14, 22, 34, 35, 37, 56, 68, 69, 72, 74] aim to enhance the visual sign representations and text decoding capabilities of SLT methods. Some works propose deep architectures based on RNNs [10, 35], GCNs [34], and Transformers [11, 37, 62, 68]. Other approaches include introducing a keypoint estimator to enhance the visual sign representation [14, 35, 56, 74], introducing pre-training tasks [22, 72], or jointly modelling several SLT-related tasks [69]. Some works also introduce larger datasets (e.g., How2Sign [20] and BOBSL [1]), which present a huge challenge with their large sign and text vocabularies. Besides, some recent works [67, 72] focus on the gloss-free setting – these works do not use sign gloss annotations for training, which reduces the cost of training SLT models, and our work also falls into this category. In contrast to existing works, we aim to harness the capabilities of off-the-shelf and frozen LLMs to perform SLT, by regularizing the sign videos into a language-like representation and prompting the LLM to generate text of the desired language.

Large Language Models (LLMs) refer to language models that have been trained extensively on a very large web-scale text corpus. LLMs have shown impressive text generation capabilities, attracting a lot of attention recently [21, 49, 71]. In particular, since LLMs have been trained on large amounts of text data, they demonstrate strong generalization capabilities across various text-based tasks, including code generation [51], open-domain question answering [75], and multilingual translation [9, 15]. Inspired by the re-

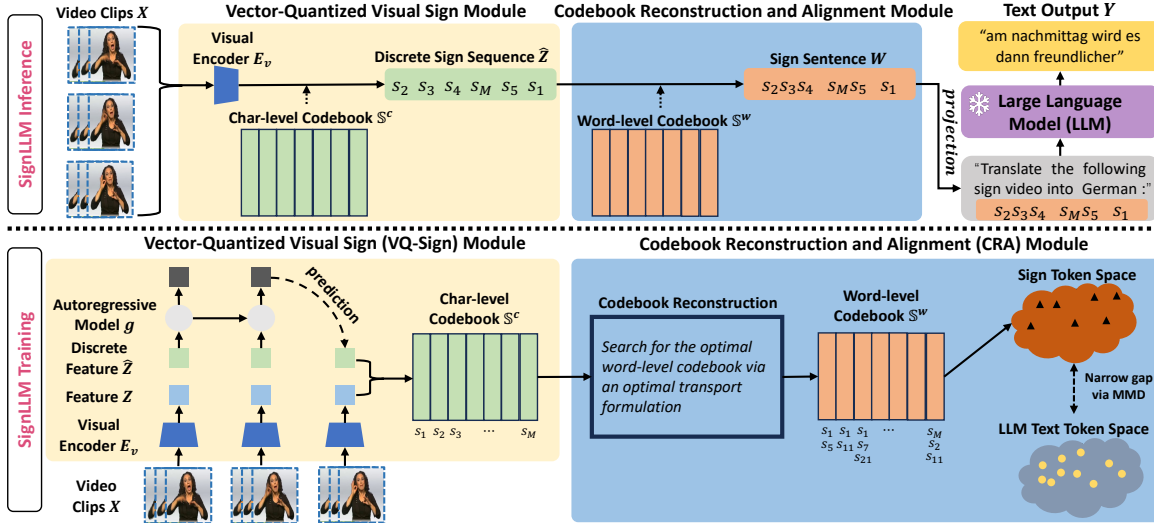


Figure 1. An overview of our SignLLM framework. **During inference (top):** Given an input sign video X , we first pass it through our VQ-Sign module to obtain a sequence of discrete character-level sign tokens \hat{Z} . Our VQ-Sign consists of a visual encoder E_v to extract compact features and a character-level sign codebook \mathbb{S}^c for quantization to obtain \hat{Z} . Next, we feed \hat{Z} into our CRA module, which reorganizes \hat{Z} by replacing short sequences of character tokens with word-level tokens via the word-level codebook, e.g., character sequence $[s_2, s_3, s_4]$ to word $s_2s_3s_4$. This transforms the sign video data to a language-like sign sentence W , which is fed into the LLM along with a text prompt which guides the LLM to generate translations in the desired language. **During training (bottom):** We optimize VQ-Sign and its discrete sign codebook via a context prediction task, which seeks to recognize the future time steps based on the current context information. Next, for our CRA module, we construct the optimal word-level codebook by considering two aspects: entropy and size, which we address using optimal transport techniques. Then, we narrow the gap between the sign token space and LLM’s text token space via minimizing the MMD loss, which improves the semantic compatibility between them.

cent advancements in LLMs, we explore harnessing LLMs for translation of sign videos, by converting the sign videos to a sequence of language-like sign tokens via our SignLLM framework, and treating the sign tokens as a form of language that can be translated by an LLM. To the best of our knowledge, we are the first work to leverage an off-the-shelf and frozen LLM to tackle SLT.

3. Method

In this section, we first introduce the overview of our SignLLM in Sec. 3.1. Then, we describe two main components of our SignLLM framework: the VQ-Sign and CRA modules in Sec. 3.2 and Sec. 3.3, respectively. Finally, we list the training and inference details in Sec. 3.4.

3.1. SignLLM Overview

In order to effectively handle SLT, in this paper we draw inspiration from LLMs’ remarkable capabilities in generating translations across multiple languages [9, 15]. In particular, LLMs have been extensively trained on a large web-scale multilingual text corpus and have learned diverse knowledge regarding the properties of many languages, thus they are able to draw on shared commonalities with previously learned languages to effectively handle new languages with limited data [66, 76].

Therefore, to leverage the strong translation capabilities

of LLMs to handle SLT, we introduce a novel SignLLM framework. SignLLM converts the input sign video X into a language-like sign sentence W that aligns with the linguistic characteristics of spoken languages and is friendly and compatible with LLMs. Then, to perform SLT, the language-like sign sentence W can be fed into an off-the-shelf and frozen LLM along with a text prompt that guides the LLM to generate translations in the desired language.

Specifically, to produce sign sentences W that are friendly and understandable by LLMs, we aim to regularize our sign sentences W to embody two core linguistic characteristics: **Discrete Characteristics:** Spoken languages are naturally discrete and consist of distinct words or sub-words with corresponding discrete tokens in a vocabulary [8, 59]. **Hierarchical Structure:** Most spoken languages exhibit three hierarchical semantic levels – the sentence, word, and character levels [36, 52], where words are composed from characters and sentences are composed from various words.

To achieve the above, our SignLLM framework comprises of three parts, as shown in Fig. 1: (1) The **VQ-Sign module** converts the input sign video X into a sequence of discrete sign tokens \hat{Z} , aligning the sign representations with text’s *discrete characteristics*. These sign tokens \hat{Z} are character-level sign tokens that are retrieved from a learned discrete character-level codebook. (2) The **CRA module** maps meaningful compositions of character-level sign tokens \hat{Z} into word-level sign tokens that form a sign sentence

W , further imparting a language-like *hierarchical structure* to the video sign representations. Moreover, we also align the sign token codebooks towards the text token space to improve semantic compatibility. (3) An **off-the-shelf LLM** takes the sign sentence W as input, along with an instructive text prompt that guides the LLM to generate the translation in the desired language. More details regarding the text prompt are in the Supplementary. Next, we present our VQ-Sign and CRA modules in detail.

3.2. Vector-Quantized Visual Sign Module

First, in order to produce a language-like representation, we would like to impart *discrete characteristics* to the input sign videos, aligning them more closely with spoken language representations that are inherently discrete and consist of distinct tokens in a vocabulary. However, achieving this is not straightforward because sign videos are a continuous signal in a high-dimensional spatio-temporal space which cannot be easily represented by a set of discrete tokens, and for which the vocabulary is not readily available. Hence, we introduce our Vector-Quantized Visual Sign (VQ-Sign) module to quantize the sign video X into a sequence of discrete sign tokens \hat{Z} via a sign codebook \mathbb{S}^c . As illustrated in Fig. 1, our VQ-Sign module involves a series of steps, which we present in detail next.

In the first step, we extract a compact feature Z from the high-dimensional input sign video $X \in \mathbb{R}^{N \times H \times W}$, where N is the number of video frames, while H and W are the height and width of the video frames respectively. To be precise, the sign video X is first organized into a sequence of short overlapping video clips, then each short video clip is fed into a visual encoder E_v to extract a compact feature representation of dimensionality d . Overall, this step transforms the original high-dimensional input sign video $X \in \mathbb{R}^{N \times H \times W}$ into a compact feature $Z \in \mathbb{R}^{\frac{N}{n} \times d}$, where n represents the number of frames between the start of neighboring clips. Notably, since Z is obtained by processing $\frac{N}{n}$ short clips, Z can also be seen as a sequence of $\frac{N}{n}$ clip-wise features, i.e., $\{z_t\}_{t=1}^{\frac{N}{n}}$, where each $z_t \in \mathbb{R}^d$ corresponds to the feature of the t -th short clip.

In the next step, we transform the feature $Z = \{z_t\}_{t=1}^{\frac{N}{n}}$ into a sequence of discrete tokens \hat{Z} using a codebook \mathbb{S}^c . Specifically, we discretize each clip’s feature z_t into a discrete token \hat{z}_t by finding the matching token s_j from the codebook $\mathbb{S}^c = \{s_i\}_{i=1}^M$, where the i -th token in the codebook is denoted as $s_i \in \mathbb{R}^d$ and M is the number of tokens in the codebook. The matching token s_j is the codebook’s closest element to the feature z_t in terms of Euclidean distance, i.e., $j = \arg \min_i (\|z_t - s_i\|_2^2)$. After the matching, each feature z_t is replaced by $\hat{z}_t = s_j$, which results in a discrete token sequence as shown in Fig. 1, e.g., $[s_2, s_3, s_4, s_M, s_5, s_1]$. Note that, we randomly initialize all

the tokens $\{s_i\}_{i=1}^M$ in the sign codebook \mathbb{S}^c at the start and optimize them during training, as introduced next.

However, we face a challenge in learning the discrete codebook \mathbb{S}^c . In particular, although autoencoding [59] has been a popular method for producing a codebook of discrete units, the high complexity of sign videos makes autoencoding (i.e., self-reconstruction of sign videos) challenging and costly. Therefore, inspired from predictive coding [4], a widely-used method in text and speech representation learning [6, 44, 46], we propose to learn discrete representations of sign videos through a context prediction task. Context prediction [46] is a self-supervised task that focuses on recognizing the future content in latent space based on the current information, which can learn discrete representations while eliminating the need for reconstructing the high-dimensional input video data. Furthermore, previous works show that training with context prediction effectively captures the temporal dependencies and relationships between elements in a sequence [5, 27], and the learned representations are often transferable to downstream tasks [6, 46].

Specifically, we employ a context prediction task where we try to distinguish a future sample $z_{\tau+k}$ based on the current context representation c_τ at various time steps τ . To facilitate this task, after we obtain the discrete token sequence \hat{Z} , we further produce a context latent representation c_τ using an auto-regressive model g that summarizes all the discrete tokens before a certain time step τ (i.e., $\{\hat{z}_t\}_{t \leq \tau}$) to produce context latent representations $c_\tau = g(\{\hat{z}_t\}_{t \leq \tau})$. Then, we optimize our module by minimizing the following context prediction contrastive loss $\mathcal{L}_k^{\text{cp}}$:

$$\mathcal{L}_k^{\text{cp}} = - \sum_{\tau=1}^{\frac{N}{n}-k} (\log \sigma(z_{\tau+k}^\top h_\tau) + \lambda \mathbb{E}_{\tilde{z} \sim p_n} [\log \sigma(-\tilde{z}^\top h_\tau)]) \quad (1)$$

where h_τ is obtained by applying a trainable linear layer to c_τ , $\sigma(z_{\tau+k}^\top h_\tau)$ is the probability of $z_{\tau+k}$ being the true sample among the negatives, \tilde{z} are negative samples drawn from a minibatch p_n , and λ is a hyperparameter. We sum $\mathcal{L}_k^{\text{cp}}$ over different step sizes k to obtain the context prediction loss $\mathcal{L}^{\text{cp}} = \sum_{k=1}^K \mathcal{L}_k^{\text{cp}}$, where K is the maximum number of future clips that we are interested in.

Following [18], in order to optimize the matching between \hat{z}_t and z_t , we further add two losses to optimize the matching distance between \hat{z}_t and z_t , such that the overall loss \mathcal{L}^{VQ} to optimize our VQ-Sign module is as follows:

$$\mathcal{L}^{\text{VQ}} = \mathcal{L}^{\text{cp}} + \sum_{t=1}^{N/n} \|\text{sg}(z_t) - \hat{z}_t\|^2 + \gamma \sum_{t=1}^{N/n} \|z_t - \text{sg}(\hat{z}_t)\|^2, \quad (2)$$

where $\text{sg}(z) \equiv z$ is the stop-gradient operator and γ is a hyperparameter. By optimizing \mathcal{L}^{VQ} , we can train our VQ-Sign and the discrete codebook without the need for reconstructing high-dimensional video clips, which makes the codebook construction a viable and relatively cheap option.

In summary, our VQ-Sign transforms sign videos into sequences of discrete sign tokens \hat{Z} , which are friendlier and more understandable to LLMs. Notably, the *produced discrete tokens \hat{Z} can be likened to character-level tokens*, in the sense that each discrete token \hat{z}_t corresponds to a short clip and may not contain much semantic meaning on its own (similar to linguistic characters), but they can be combined into a sequence to convey a clear semantic meaning (akin to forming a word or sentence). Thus, inspired by this, we call VQ-Sign’s codebook \mathbb{S}^c the character-level codebook that contains character-level sign tokens.

3.3. Codebook Reconstruction and Alignment

In the previous section, we quantize sign videos into discrete character-level sign tokens, which aligns them closer to language representations. In this section, our goal is to impart a *hierarchical structure* to our sign video representations, which makes them align even closer to language representations. Specifically, we aim to compose our character-level sign tokens into word-level sign tokens to mirror the observed hierarchical structure in spoken language, which makes them even more compatible with LLMs.

Intuitively, considering a spoken language sentence, we can represent it as a sequence of words, with each word formed by one or multiple characters. For example, a sentence ‘I love AI’ can be decomposed into a word sequence [‘I’, ‘love’, ‘AI’], where the word ‘love’ is in turn formed by the character sequence [‘l’, ‘o’, ‘v’, ‘e’]. We observe that, although each individual character may not contain much semantic meaning on its own, they can be composed to form words with clearer semantic meaning. In a similar fashion, we also want to impart such hierarchical structure to our character-level sign tokens by composing them to form meaningful word-level sign tokens.

Hence, we aim to find an optimal transformation from character-level sign tokens to word-level sign tokens for enhanced readability and compatibility with LLMs. To this end, we introduce the Codebook Reconstruction and Alignment (CRA) module to transform the character-level codebook \mathbb{S}^c from VQ-Sign into a word-level codebook \mathbb{S}^w whose tokens convey richer and clearer semantic meaning. Inspired by optimal transport methods [16, 61, 65], we observe that the above transformation can be formulated as an optimal transport problem of transporting characters into words, thus we introduce a *codebook reconstruction algorithm* with an optimal transport formulation to find an optimal transformation. Additionally, to further reduce the distribution gap between the sign tokens and the text tokens, our CRA module also performs *sign-text alignment*, enhancing the semantic compatibility of the sign tokens with LLMs. We introduce the details below.

To begin, the objective of our **Codebook Reconstruction Algorithm** is to create a word-level codebook \mathbb{S}^w

based on VQ-Sign’s character-level codebook \mathbb{S}^c . The challenge lies in determining which character-level sign tokens should be assembled together to form word-level sign tokens, which is a complex problem. To address this complexity, we adopt an approach based on two fundamental principles. Firstly, in order to maximize the overall predictability of the word-level tokens and enhance the distinctiveness of each token, we seek to minimize the *entropy* of each word-level token within the vocabulary [42]. We remark that, several approaches for establishing language-based subword vocabularies [8, 55] can be seen as entropy-minimizing approaches, employing different heuristics to establish the vocabulary with the goal of minimizing entropy [23]. On the other hand, considering the limited availability of sign video data, we incorporate *codebook size* as another key factor in our word-level codebook construction, since studies on languages with limited data [26, 53] have also identified vocabulary size as a crucial aspect. Specifically, too small a vocabulary can result in sub-optimal entropy values, while an excessively large vocabulary size can lead to issues such as parameter explosion and token sparsity which hinder understanding [2], and finding the right balance between these effects becomes even more sensitive for languages with limited data [19, 53].

Based on these principles, our objective is to determine an optimal codebook size that maximizes the entropy reduction while taking into account the increase in codebook size. In other words, we would like to find an optimal codebook size that maximizes the gradient of the entropy reduction with respect to the codebook size increase. To simplify the optimal size searching problem, we define a fixed size increment m and search through codebooks of various sizes (where the difference between each codebook size is m tokens). Specifically, we define the r -th codebook (\mathbb{S}_r^w) as the codebook with $r \times m$ tokens. Then, we seek to identify the optimal set of word-level tokens, where each word-level token is composed of character-level tokens. We approach this by formulating the character compositions as an optimal transport problem, where characters are transported to words.

However, it can be challenging to identify specific character combinations that convey precise semantic information, due to the temporal complexity of sign videos, which often makes the character-level token sequences \hat{Z} quite messy. For instance, some signers may execute the signing motions at a slower speed, which can lead to consecutive short video clips being highly similar, resulting in consecutively-repeating character-level discrete tokens. Thus, the character-level sequences between different signers can differ significantly (e.g., $[s_1, s_2]$ vs $[s_1, s_1, s_1, s_2, s_2]$) due to such duplication of character-level tokens, even though they may contain the same semantic information. At the same time, simply filtering out re-

peated character-level tokens straightforwardly (e.g., setting all $[s_1, s_1, s_1, s_2, s_2]$ to $[s_1, s_2]$) is sub-optimal, since the speed of some signs can also convey some information [31, 60], e.g., if a signer signs “ugly” quickly, it conveys “very ugly” in American Sign Language.

Pre-processing of Repeated Characters. Therefore, to alleviate the impact of the signer’s speed while keeping the information regarding each sign’s speed, we *first pre-process the character-level sequences* as follows: First, we find all the repeated tokens in the character-level sequence and compute the average number of repeated tokens (α) in each sequence. Then, for each repeated sequence (e.g., $[s_1, s_1, s_1]$), we keep the first character and remove the tailing repetitive ones (e.g., $[s_1, s_1, s_1]$ to $[s_1]$). At the same time, if the character-level tokens repeat more than α times, we insert a single character-level token s_0 as a “slowing down” sign, e.g., $[s_1, s_1, s_1]$ to $[s_1, s_0]$ if $\alpha < 3$. Crucially, this allows us to reduce redundancy, while still representing “quick” or “slow” signs that account for differences in the signer’s speed. Overall, this pre-processing and reducing of repeating characters makes the character-level sequence less messy, facilitating the search for specific meaningful character combinations.

Optimal Transport Formulation. Then, using the pre-processed character-level sequences, we search for an optimal word-level codebook via an optimal transport formulation [16, 61, 65]. Following our discussion above, we aim to find an optimal word-level codebook (i.e., \mathbb{S}_r^w) with low entropy and compact size. Specifically, to measure codebook entropy, we follow [25, 42, 45] to define the entropy of the r -th codebook \mathbb{S}_r^w as:

$$\mathcal{H}_{\mathbb{S}_r^w} = - \sum_{w_j \in \mathbb{S}_r^w} P(w_j) \log P(w_j), \quad (3)$$

where $P(w_j)$ is the relative frequency of the j^{th} token w_j from the word-level codebook \mathbb{S}_r^w . Then, based on VQ-Sign’s character-level codebook \mathbb{S}^c , the entropy of the word-level codebook \mathbb{S}_r^w can be computed through the following (with proof in Supplementary):

$$\begin{aligned} \mathcal{H}_{\mathbb{S}_r^w} = & - \sum_{w_j \in \mathbb{S}_r^w} \sum_{s_i \in \mathbb{S}^c} P(w_j, s_i) \log P(w_j, s_i) \\ & - \sum_{w_j \in \mathbb{S}_r^w} \sum_{s_i \in \mathbb{S}^c} P(w_j, s_i) (-\log P(s_i|w_j)), \end{aligned} \quad (4)$$

where $P(s_i|w_j)$ is the probability of the character-level token s_i appearing in the word-level token w_j .

Then, to formally define our objective function, we follow existing works [16, 48, 65] to define a transport matrix \mathbf{P} that represents the assignments of characters to words, and a distance matrix \mathbf{D} that represents the cost of transportation. We define the transport matrix $\mathbf{P} \in \mathbb{R}^{m \times (r \cdot m)}$ with the (j, i) -th element as $P(w_j, s_i)$, and define the distance matrix $\mathbf{D} \in \mathbb{R}^{m \times (r \cdot m)}$ as a matrix whose (j, i) -th element is $\log P(s_i|w_j)$. Note that, if w_j contains s_i , we use

$\frac{1}{\text{length}(w_j)}$ to estimate $P(s_i|w_j)$, and if w_j does not contain s_i , then it is deemed an infeasible assignment, and we set the distance $\log P(s_i|w_j) = \infty$. We further define $H(\mathbf{P})$ as $-\sum_{w_j \in \mathbb{S}_r^w} \sum_{s_i \in \mathbb{S}^c} P(w_j, s_i) \log P(w_j, s_i)$, which is simply the entropy of the probability distribution $P(w_j, s_i)$.

Hence, based on Eq. 4, the objective function for minimizing the entropy of \mathbb{S}_r^w can be formulated as:

$$\arg \min_{\mathbf{P} \in \mathbb{R}^{m \times (r \cdot m)}} H(\mathbf{P}) + \sum_j \sum_i \mathbf{P}(j, i) \mathbf{D}(j, i). \quad (5)$$

Specifically, following previous works [16, 48, 65], we impose two constraints on the transport matrix \mathbf{P} : the sum of each row in \mathbf{P} should equal to the probability of character token s_i and the sum of each column in \mathbf{P} should equal to the probability of word token w_j . Formally, we constrain the transport matrix \mathbf{P} with: $|\sum_i \mathbf{P}(i, j) - P(w_j)| \leq \epsilon$ and $|\sum_j \mathbf{P}(i, j) - P(s_i)| \leq \epsilon$, where ϵ is a small positive fixed hyperparameter.

Intuitively, this optimization process can be regarded as an optimal transport problem to find the best way to transport mass from \mathbb{S}^c to \mathbb{S}_r^w . To handle this optimal transport problem, we leverage the Sinkhorn algorithm [16, 48, 65], allowing us to effectively construct the candidate word-level codebooks \mathbb{S}_r^w with minimal entropy. Since our increment m between each candidate codebook is fixed, we can find the optimal codebook size that maximizes the gradient of entropy reduction by simply computing and finding the maximum the entropy difference between \mathbb{S}_r^w and \mathbb{S}_{r-1}^w .

After finding the optimal word-level codebook, we construct the word-level sign tokens by composing all features of the character-level tokens into the word-level tokens via our autoregressive model g . Refer to Supplementary for more details. Overall, with our codebook reconstruction algorithm, we can construct an optimal word-level sign codebook with low entropy yet also with relatively small size.

Sign-Text Alignment. Next, we further align the sign tokens with the text tokens used in LLMs in order to further improve semantic compatibility between them. To achieve this, we measure the distribution gap between sign tokens and text tokens via Maximum Mean Discrepancy (MMD) [57] and then optimize the sign tokens’ embeddings by minimizing MMD, which narrows down the distribution gap. Specifically, we compute the gap between the sign embedding space \mathbb{F}^s and text embedding space \mathbb{F}^t via MMD as:

$$\begin{aligned} \mathcal{L}^{MMD}(\mathbb{F}^s, \mathbb{F}^t) = & \sum_{i=1}^{n_s} \sum_{j=1}^{n_s} \frac{k(f(p_i), f(p_j))}{n_s^2} \\ & + \sum_{i=1}^{n_t} \sum_{j=1}^{n_t} \frac{k(q_i, q_j)}{n_t^2} - \sum_{i=1}^{n_s} \sum_{j=1}^{n_t} \frac{2 \cdot k(f(p_i), q_j)}{n_s n_t}, \end{aligned} \quad (6)$$

where p and q are the tokens in \mathbb{F}^s and \mathbb{F}^t , f is a small projection module that projects \mathbb{F}^s to \mathbb{F}^t , n_s and n_t are the numbers of tokens in \mathbb{F}^s and \mathbb{F}^t , and $k(\cdot)$ represents the radial

kernel [57] that measures the distance between two samples. We apply MMD loss to both the word-level and the character-level tokens to narrow the overall sign-text gap.

3.4. Training and Inference

Inference. Given a sign video X , we first extract compact features Z via the visual encoder E_v , and quantize Z to \hat{Z} via VQ-Sign’s learned character-level codebook. Then, we transform the sequence of discrete character-level tokens \hat{Z} into word-level tokens via our CRA, which produces a sign sentence W . Lastly, we project the sign sentence W to LLM embedding space via the small projection module f , and then feed the sign sentence W into the LLM along with a text prompt to instruct the LLM to perform the SLT task and generate text in the desired language.

Training. Our SignLLM is optimized in two stages: pre-training and fine-tuning. Specifically, the pre-training stage, which does not require explicit SLT supervision, includes two sub-stages: (i) We first pre-train VQ-Sign via the context prediction task with \mathcal{L}^{VQ} in Eq. 2. (ii) Then, based on the VQ-Sign’s learned character-level codebook, we construct the word-level codebook with the codebook reconstruction algorithm (Sec. 3.3) and apply the MMD loss \mathcal{L}^{MMD} in Eq. 6 to align the sign codebooks (\mathbb{S}^c and \mathbb{S}^w) and the text vocabulary of the desired language.

After the pre-training, we fine-tune SignLLM. To aid LLMs in understanding the sign sentences as texts, we additionally maximize the similarity between the text tokens Y generated by LLM and the ground truth tokens \bar{Y} as: \mathcal{L}^{sim} by minimizing the cross-entropy between them. We fine-tune our SignLLM (with frozen LLM) via the loss \mathcal{L}^{ft} as follows: $\mathcal{L}^{ft} = \mathcal{L}^{VQ} + \lambda_1 \mathcal{L}^{MMD} + \lambda_2 \mathcal{L}^{sim}$, where λ_1 and λ_2 are hyperparameters. Note that we follow previous gloss-free works [67, 72] to train our SignLLM framework, eliminating the need for additional gloss data.

4. Experiments

4.1. Implementation Details

Our visual encoder E_v is constructed by appending two Conv3D layers with a kernel size of (5, 3, 3) and a stride of (2, 1, 1) to a ResNet18 [28] pre-trained on ImageNet [17]. Each clip consists of 13 frames and the gap between the neighboring clips (n) is 4. Besides, the auto-regressive model g is implemented as a Convolutional Gated Recurrent Layer with a kernel size of (1, 1). We set the total number of discrete vectors M at 256 and each vector’s dimension d at 1024 for our character-level codebook. In VQ-Sign’s pre-training phase, we set $\gamma = 0.25$, initialize the learning rate at 0.01 and use the Adam algorithm, training the model to predict the future three clips ($K = 3$) for 200 epochs. During codebook reconstruction, we set the increment m to 32. We employ the frozen LLaMA-7B-16bit [58]

as our LLM and project the codebook space to LLaMA’s embedding space via f , which consists of two fc-layers with ReLU. We set $\lambda_1 = 0.5$, $\lambda_2 = 1$, and initialize the learning rate at 0.001 to fine-tune our SignLLM over 20 epochs. Please see Supplementary for more implementation details.

4.2. Datasets and Evaluation Metrics

Datasets. We follow previous works [13, 14, 34, 72–74] to run experiments on the Phoenix-2014T [10] and CSL-Daily [73] datasets for SLT, and evaluate on their dev and test sets. **Phoenix-2014T** [10] is a German sign language dataset with a vocabulary size of 2887 German words. The training, dev, and test sets contain 7096, 519, and 642 samples. **CSL-Daily** [73] is a Chinese sign language dataset with a vocabulary size of 2343 Chinese words. The training, dev, and test sets contain 18401, 1077, and 1176 samples.

Evaluation Metrics. Following previous works [13, 14, 37, 67, 72–74], we adopt BLEU [47] and ROUGE-L [38] as the evaluation metrics for SLT. BLEU-n evaluates the average translation precision up to n-grams, and we follow previous works [13, 14, 72–74] to report results for BLEU-1 to BLEU-4 (i.e., B1, B2, B3 and B4). ROUGE-L (or ROUGE) computes the F1 score based on the longest common subsequence between the predicted and ground truth texts.

4.3. Main Results

Results on PHOENIX2014T dataset. Tab. 1 presents a comparison of our approach with state-of-the-art gloss-based and gloss-free methods for SLT. Our method consistently improves upon all reported metrics as compared to other gloss-free approaches.

Results on CSL-Daily dataset. We compare our method with state-of-the-art approaches on the CSLDaily dataset in Tab. 2. We outperform previous gloss-free works on all metrics, showing the efficacy of our approach.

4.4. Ablation Study

To further investigate the proposed method, we follow previous works [13, 33, 67, 69, 72] to conduct extensive ablation experiments on the Phoenix-2014T dev and test sets. Refer to Supplementary for more experiment results.

Impact of LLM. First, we evaluate the impact of leveraging LLMs for the SLT task. Specifically, we establish three baselines: 1) **Ours (w/o LLM)** where we replace the LLM with a trainable lightweight text generator (mBART [40]). 2) **Ours (w/ T5)** where we replace our LLM with a smaller LLM (T5 [50]). As shown in Tab. 3, our approach (w/ LLaMA [58]) achieves a much better performance than Ours (w/o LLM), showing the efficacy of leveraging a powerful LLM. We also find that using a smaller and less powerful LLM (T5) leads to worse performance, as expected.

Impact of SignLLM. Next, we explore the impact of SignLLM by comparing against the following baselines: 1) **En-**

Table 1. Results on Phoenix-2014T dataset [10].

Setting	Method	Dev					Test				
		B1	B2	B3	B4	ROUGE	B1	B2	B3	B4	ROUGE
Gloss-based	SLRT [11]	47.26	34.40	27.05	22.38	-	46.61	33.73	26.19	21.32	-
	ConSLT [22]	-	-	-	21.11	47.74	-	-	-	21.59	47.69
	STN-SLT [62]	49.12	36.29	28.34	23.23	-	48.61	35.97	28.37	23.65	-
	STMC-T [74]	47.60	36.43	29.18	24.09	48.24	46.98	36.09	28.70	23.65	46.65
	BN-TIN-Transf.+SignBT [73]	51.11	37.90	29.80	24.45	50.29	50.80	37.75	29.72	24.32	49.54
	PET [33]	-	-	-	-	-	49.54	37.19	29.30	24.02	49.97
	MMTLB [13]	53.95	41.12	33.14	27.61	53.10	53.97	41.75	33.84	28.39	52.65
	TS-SLT [14]	54.32	41.99	34.15	28.66	54.08	54.90	42.43	34.46	28.95	53.48
SLTUNET [69]	-	-	-	27.87	52.23	52.92	41.76	33.99	28.47	52.11	
Gloss-free	NSLT [10]	28.10	16.81	11.82	9.12	31.00	27.10	15.61	10.82	8.35	29.70
	NSLT+Bahdanau [7, 10]	31.87	19.11	13.16	9.94	31.80	32.24	19.03	12.83	9.58	31.80
	NSLT+Luong [10, 41]	31.58	18.98	13.22	10.00	32.60	29.86	17.52	11.96	9.00	30.70
	TSPNet [37]	-	-	-	-	-	36.10	23.12	16.88	13.41	34.96
	CSGCR [70]	35.85	24.77	18.65	15.08	38.96	36.71	25.40	18.86	15.18	38.85
	GASLT [67]	-	-	-	-	-	39.07	26.74	21.86	15.74	39.86
	GFSLT-VLP [72]	44.08	33.56	26.74	22.12	43.72	43.71	33.18	26.11	21.44	42.49
	Ours	46.88	36.59	29.91	25.25	47.23	45.21	34.78	28.05	23.40	44.49

Table 2. Results on CSL-Daily dataset [73]. * means that the result was reproduced by [72]

Setting	Method	Dev					Test				
		B1	B2	B3	B4	ROUGE	B1	B2	B3	B4	ROUGE
Gloss-based	SLRT [11]	37.47	24.67	16.86	11.88	37.96	37.38	24.36	16.55	11.79	36.74
	ConSLT [22]	-	-	-	14.80	41.46	-	-	-	14.53	40.98
	BN-TIN-Transf.+SignBT [73]	51.46	37.23	27.51	20.80	49.49	51.42	37.26	27.76	21.34	49.31
	MMTLB [13]	53.81	40.84	31.29	24.42	53.38	53.31	40.41	30.87	23.92	53.25
	TS-SLT [14]	55.21	42.31	32.71	25.76	55.10	55.44	42.59	32.87	25.79	55.72
	SLTUNET [69]	-	-	-	23.99	53.58	54.98	41.44	31.84	25.01	54.08
Gloss-free	SLRT* [11]	21.03	9.97	5.96	4.04	20.51	20.00	9.11	4.93	3.03	19.67
	NSLT+Luong [10, 41]	34.22	19.72	12.24	7.96	34.28	34.16	19.57	11.84	7.56	34.54
	GASLT [67]	-	-	-	-	-	19.90	9.94	5.98	4.07	20.35
	GFSLT-VLP [72]	39.20	25.02	16.35	11.07	36.70	39.37	24.93	16.26	11.00	36.44
	Ours	42.45	26.88	17.90	12.23	39.18	39.55	28.13	20.07	15.75	39.91

Table 3. Ablation study for impact of LLM

Method	Dev				Test			
	B1	B2	B3	B4	B1	B2	B3	B4
Ours (w/o LLM)	29.75	20.04	14.96	11.95	27.20	18.29	13.32	10.36
Ours (w/ T5)	44.20	34.55	27.15	22.90	44.03	34.12	27.23	22.51
Ours	46.88	36.59	29.91	25.25	45.21	34.78	28.05	23.40

Table 4. Ablation study for impact of SignLLM.

Method	Dev				Test			
	B1	B2	B3	B4	B1	B2	B3	B4
Encoder Only	26.95	17.04	12.11	9.41	25.63	16.10	11.20	8.42
Encoder + FT	39.29	29.29	22.55	18.18	39.92	29.14	22.54	18.17
Ours	46.88	36.59	29.91	25.25	45.21	34.78	28.05	23.40

Table 5. Ablation study for main components of SignLLM.

Method	Dev				Test			
	B1	B2	B3	B4	B1	B2	B3	B4
Ours (w/o VQ-Sign)	35.78	24.65	18.08	14.15	33.14	22.80	16.74	13.13
Ours (w/o Codebook Reconstruction)	40.45	30.40	24.07	19.79	40.25	30.05	23.63	19.47
Ours (w/o Sign-text Alignment)	29.05	19.33	13.72	10.40	28.67	19.22	13.73	10.63
Ours	46.88	36.59	29.91	25.25	45.21	34.78	28.05	23.40

coder Only where we directly feed the outputs of the visual encoder E_v into the LLM, and keep the LLM frozen. 2) **Encoder + FT** where we directly feed the outputs of the visual encoder E_v into the LLM, and follow LLaVA [39] to fine-tune them to translate sign videos via LoRA [32]. We report the results in Tab. 4, where we significantly outperform the baselines. This shows that SignLLM is effective in harnessing off-the-shelf LLMs for the SLT task.

Impact of Main Components of SignLLM. We also verify the impact of the key components of SignLLM by comparing against the following baselines: 1) **Ours (w/o VQ-Sign)** where we directly quantize the visual encoder E_v 's output feature using k -means algorithm. The visual encoder E_v is pre-trained via the similarity loss \mathcal{L}^{sim} instead of our proposed VQ-Sign's loss \mathcal{L}^{VQ} . 2) **Ours (w/o Codebook Reconstruction)** where we feed the character-level sign tokens \hat{Z} from VQ-Sign directly into the LLM, while applying the sign-text alignment loss \mathcal{L}^{MMD} to the character-level sign tokens. 3) **Ours (w/o Sign-text Alignment)** where we do not apply the sign-text alignment loss \mathcal{L}^{MMD} . As shown in Tab. 5, removing any of our main designs leads to a significant performance drop, showing their efficacy.

Ground Truth:	jetzt wünsche ich ihnen noch einen schönen abend (Now I wish you a nice evening)
GFSLT-VLP:	liebe zuschauer guten abend (Dear viewers, good evening)
Ours:	ich wünsche ihnen noch einen schönen abend (I wish you a nice evening)
Ground Truth:	am tag ist es überwiegend stark bewölkt oder neblig trüb (During the day it is mostly cloudy or foggy)
GFSLT-VLP:	am tag, teils stark bewölkt oder neblig trüb teils freundlich regional kann sich die sonne auch längere zeit zeigen (During the day, partly very cloudy or foggy, partly friendly. In regions, the sun can also appear for a long time)
Ours:	am tag ist es meist stark bewölkt oder neblig (During the day it is usually cloudy or foggy)

Figure 2. Visualization of translation results. Correct translations are in blue while the wrong translations are in red.

cant performance drop, showing their efficacy.

Qualitative Results. We present two sample translations generated by our SignLLM and the current state-of-the-art (GFSLT-VLP [72]) in Fig. 2 for qualitative analysis. In the first sample (top), our model produces a highly accurate translation, whereas [72] inaccurately represents the semantic information. In the second sample (bottom), our model successfully preserves sentence semantics, while [72] introduces a translation error, resulting in redundant and erroneous information. These examples qualitatively show our SignLLM's efficacy in producing accurate translations.

5. Conclusion

We present SignLLM, a novel framework to harness off-the-shelf and frozen LLMs for SLT. SignLLM imparts language-like characteristics to sign video representations through the VQ-Sign and CRA modules, and a sign-text alignment loss improves semantic compatibility. We empirically observe that applying our SignLLM leads to state-of-the-art gloss-free results on two popular SLT benchmarks.

Acknowledgements. This project is supported by the Ministry of Education, Singapore, under the AcRF Tier 2 Projects (MOE-T2EP20222-0009 and MOE-T2EP20123-0014), National Research Foundation Singapore under its AI Singapore Programme (AISG-100E-2023-121).

References

- [1] Samuel Albanie, Gül Varol, Liliane Momeni, Hannah Bull, Triantafyllos Afouras, Himel Chowdhury, Neil Fox, Bencie Woll, Rob Cooper, Andrew McParland, et al. Bbc-oxford british sign language dataset. *arXiv preprint arXiv:2111.03635*, 2021. [2](#)
- [2] Ben Allison, David Guthrie, and Louise Guthrie. Another look at the data sparsity problem. In *Text, Speech and Dialogue: 9th International Conference, TSD 2006, Brno, Czech Republic, September 11-15, 2006. Proceedings 9*, pages 327–334. Springer, 2006. [5](#)
- [3] Ulrich Ammon. World languages: Trends and futures. *The handbook of language and globalization*, pages 101–122, 2010. [1](#)
- [4] Bishnu S Atal and Manfred R Schroeder. Adaptive predictive coding of speech signals. *Bell System Technical Journal*, 49(8):1973–1986, 1970. [4](#)
- [5] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33:12449–12460, 2020. [4](#)
- [6] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. Data2vec: A general framework for self-supervised learning in speech, vision and language. In *International Conference on Machine Learning*, pages 1298–1312. PMLR, 2022. [4](#)
- [7] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 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*, 2015. [8](#)
- [8] Kaj Bostrom and Greg Durrett. Byte pair encoding is sub-optimal for language model pretraining. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4617–4624, 2020. [2](#), [3](#), [5](#)
- [9] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. [1](#), [2](#), [3](#)
- [10] Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, Hermann Ney, and Richard Bowden. Neural sign language translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7784–7793, 2018. [1](#), [2](#), [7](#), [8](#)
- [11] Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. 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*, pages 10023–10033, 2020. [1](#), [2](#), [8](#)
- [12] Sanyuan Chen, Yutai Hou, Yiming Cui, Wanxiang Che, Ting Liu, and Xiangzhan Yu. Recall and learn: Fine-tuning deep pretrained language models with less forgetting. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7870–7881, 2020. [1](#)
- [13] Yutong Chen, Fangyun Wei, Xiao Sun, Zhirong Wu, and Stephen Lin. A simple multi-modality transfer learning baseline for sign language translation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 5110–5120. IEEE, 2022. [1](#), [2](#), [7](#), [8](#)
- [14] Yutong Chen, Ronglai Zuo, Fangyun Wei, Yu Wu, Shujie Liu, and Brian Mak. Two-stream network for sign language recognition and translation. In *NeurIPS*, 2022. [1](#), [2](#), [7](#), [8](#)
- [15] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022. [1](#), [2](#), [3](#)
- [16] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013. [5](#), [6](#)
- [17] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [7](#)
- [18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019. [4](#)
- [19] Shuoyang Ding, Adithya Renduchintala, and Kevin Duh. A call for prudent choice of subword merge operations in neural machine translation. In *Proceedings of Machine Translation Summit XVII: Research Track*, pages 204–213, 2019. [5](#)
- [20] Amanda Duarte, Shruti Palaskar, Lucas Ventura, Deepti Ghadiyaram, Kenneth DeHaan, Florian Metze, Jordi Torres, and Xavier Giro-i Nieto. How2sign: a large-scale multi-modal dataset for continuous american sign language. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2735–2744, 2021. [2](#)
- [21] Lin Geng Foo, Hossein Rahmani, and Jun Liu. Ai-generated content (aigc) for various data modalities: A survey. *arXiv preprint arXiv:2308.14177*, 2, 2023. [2](#)
- [22] Biao Fu, Peigen Ye, Liang Zhang, Pei Yu, Cong Hu, Xiaodong Shi, and Yidong Chen. A token-level contrastive framework for sign language translation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023. [1](#), [2](#), [8](#)
- [23] Philip Gage. A new algorithm for data compression. *C Users Journal*, 12(2):23–38, 1994. [5](#)
- [24] Tanmay Gupta and Aniruddha Kembhavi. Visual programming: Compositional visual reasoning without training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14953–14962, 2023. [1](#)
- [25] Ximena Gutierrez-Vasques, Christian Bentz, Olga Sozinova, and Tanja Samardzic. From characters to words: the turning point of bpe merges. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3454–3468, 2021. [6](#)

- [26] Barry Haddow, Rachel Bawden, Antonio Valerio Miceli Barone, Jindřich Helcl, and Alexandra Birch. Survey of low-resource machine translation. *Computational Linguistics*, 48(3):673–732, 2022. 5
- [27] Awni Hannun. Sequence modeling with ctc. *Distill*, 2(11):e8, 2017. 4
- [28] Kaiying He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 7
- [29] Ruidan He, Linlin Liu, Hai Ye, Qingyu Tan, Bosheng Ding, Liying Cheng, Jiawei Low, Lidong Bing, and Luo Si. On the effectiveness of adapter-based tuning for pretrained language model adaptation. 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 2208–2222, 2021. 1
- [30] Valentin Hofmann, Janet Pierrehumbert, and Hinrich Schütze. Dagobert: Generating derivational morphology with a pretrained language model. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3848–3861, 2020. 2
- [31] Jiahui Hou, Xiang-Yang Li, Peide Zhu, Zefan Wang, Yu Wang, Jianwei Qian, and Panlong Yang. Signspeaker: A real-time, high-precision smartwatch-based sign language translator. In *The 25th Annual International Conference on Mobile Computing and Networking*, pages 1–15, 2019. 6
- [32] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. 8
- [33] Tao Jin, Zhou Zhao, Meng Zhang, and Xingshan Zeng. Prior knowledge and memory enriched transformer for sign language translation. In *Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 3766–3775. Association for Computational Linguistics, 2022. 7, 8
- [34] Jichao Kan, Kun Hu, Markus Hagenbuchner, Ah Chung Tsoi, Mohammed Bennamoun, and Zhiyong Wang. Sign language translation with hierarchical spatio-temporal graph neural network. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3367–3376, 2022. 2, 7
- [35] Sang-Ki Ko, Chang Jo Kim, Hyedong Jung, and Choongsang Cho. Neural sign language translation based on human key-point estimation. *Applied sciences*, 9(13):2683, 2019. 2
- [36] W Levelt. Producing spoken language. *The neurocognition of language*, pages 83–122, 1999. 2, 3
- [37] Dongxu Li, Chenchen Xu, Xin Yu, Kaihao Zhang, Benjamin Swift, Hanna Suominen, and Hongdong Li. Tspnet: Hierarchical feature learning via temporal semantic pyramid for sign language translation. *Advances in Neural Information Processing Systems*, 33:12034–12045, 2020. 1, 2, 7, 8
- [38] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004. 7
- [39] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024. 8
- [40] Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742, 2020. 7
- [41] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, 2015. 8
- [42] Nathaniel FG Martin and James W England. *Mathematical theory of entropy*. Number 12. Cambridge university press, 2011. 5, 6
- [43] Igor Melnyk, Vijil Chenthamarakshan, Pin-Yu Chen, Payel Das, Amit Dhurandhar, Inkit Padhi, and Devleena Das. Re-programming pretrained language models for antibody sequence infilling. In *Proceedings of the 2023 International Conference on Machine Learning (ICML)*, 2023. 1
- [44] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013. 4
- [45] Arijit Nag, Bidisha Samanta, Animesh Mukherjee, Niloy Ganguly, and Soumen Chakrabarti. Entropy-guided vocabulary augmentation of multilingual language models for low-resource tasks. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8619–8629, 2023. 6
- [46] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 4
- [47] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002. 7
- [48] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019. 6
- [49] Haoxuan Qu, Yujun Cai, and Jun Liu. Llms are good action recognizers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2
- [50] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020. 7
- [51] Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*, 2023. 2
- [52] Stephanie Seneff. The use of linguistic hierarchies in speech understanding. In *in Proc. ICSLP*, 1998. 2, 3

- [53] Rico Sennrich and Biao Zhang. Revisiting low-resource neural machine translation: A case study. In *57th Annual Meeting of the Association for Computational Linguistics*, pages 211–221. Association for Computational Linguistics (ACL), 2019. 5
- [54] Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2998–3009, 2023. 1
- [55] Xinying Song, Alex Salcianu, Yang Song, Dave Dopson, and Denny Zhou. Fast wordpiece tokenization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2089–2103, 2021. 5
- [56] Shengeng Tang, Dan Guo, Richang Hong, and Meng Wang. Graph-based multimodal sequential embedding for sign language translation. *IEEE Transactions on Multimedia*, 24:4433–4445, 2021. 2
- [57] Ilya O Tolstikhin, Bharath K Sriperumbudur, and Bernhard Schölkopf. Minimax estimation of maximum mean discrepancy with radial kernels. *Advances in Neural Information Processing Systems*, 29:1930–1938, 2016. 6, 7
- [58] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 1, 7
- [59] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 2, 3, 4
- [60] William Vicars. American sign language: “slow”. <https://www.lifefprint.com/asl1101/pages-signs/s/slow.htm>, 2017. 6
- [61] Cédric Villani et al. *Optimal transport: old and new*. Springer, 2009. 5, 6
- [62] Andreas Voskou, Konstantinos P Panousis, Dimitrios Kosmopoulos, Dimitris N Metaxas, and Sotirios Chatzis. Stochastic transformer networks with linear competing units: Application to end-to-end sl translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11946–11955, 2021. 2, 8
- [63] Ivan Vulić, Edoardo Maria Ponti, Robert Litschko, Goran Glavaš, and Anna Korhonen. Probing pretrained language models for lexical semantics. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7222–7240, 2020. 2
- [64] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3349–3364, 2020. 1
- [65] Jingjing Xu, Hao Zhou, Chun Gan, Zaixiang Zheng, and Lei Li. Vocabulary learning via optimal transport 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 (Volume 1: Long Papers)*, pages 7361–7373, 2021. 5, 6
- [66] Wen Yang, Chong Li, Jiajun Zhang, and Chengqing Zong. Bigtrans: Augmenting large language models with multilingual translation capability over 100 languages. *arXiv preprint arXiv:2305.18098*, 2023. 1, 3
- [67] Aoxiong Yin, Tianyun Zhong, Li Tang, Weike Jin, Tao Jin, and Zhou Zhao. Gloss attention for gloss-free sign language translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2551–2562, 2023. 2, 7, 8
- [68] Kayo Yin and Jesse Read. Better sign language translation with stmc-transformer. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5975–5989, 2020. 1, 2
- [69] Biao Zhang, Mathias Müller, and Rico Sennrich. SLTUNET: A simple unified model for sign language translation. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. 1, 2, 7, 8
- [70] Jian Zhao, Weizhen Qi, Wengang Zhou, Nan Duan, Ming Zhou, and Houqiang Li. Conditional sentence generation and cross-modal reranking for sign language translation. *IEEE Transactions on Multimedia*, 24:2662–2672, 2021. 8
- [71] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023. 2
- [72] Benjia Zhou, Zhigang Chen, Albert Clapés, Jun Wan, Yanyan Liang, Sergio Escalera, Zhen Lei, and Du Zhang. Gloss-free sign language translation: Improving from visual-language pretraining, 2023. 1, 2, 7, 8
- [73] Hao Zhou, Wengang Zhou, Weizhen Qi, Junfu Pu, and Houqiang Li. Improving sign language translation with monolingual data by sign back-translation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 1316–1325. Computer Vision Foundation / IEEE, 2021. 1, 2, 7, 8
- [74] Hao Zhou, Wengang Zhou, Yun Zhou, and Houqiang Li. Spatial-temporal multi-cue network for sign language recognition and translation. *IEEE Transactions on Multimedia*, 24:768–779, 2022. 1, 2, 7, 8
- [75] Fengbin Zhu, Wenqiang Lei, Chao Wang, Jianming Zheng, Soujanya Poria, and Tat-Seng Chua. Retrieving and reading: A comprehensive survey on open-domain question answering. *arXiv preprint arXiv:2101.00774*, 2021. 2
- [76] Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu, Lingpeng Kong, Jiajun Chen, Lei Li, and Shujian Huang. Multilingual machine translation with large language models: Empirical results and analysis. *arXiv preprint arXiv:2304.04675*, 2023. 1, 3