

BIGVIDEO: A Large-scale Video Subtitle Translation Dataset for Multimodal Machine Translation

Liyan Kang^{1,3*} Luyang Huang^{2*} Ningxin Peng² Peihao Zhu² Zewei Sun²

Shanbo Cheng² Mingxuan Wang² Degen Huang⁴ Jinsong Su^{1,3†}

¹School of Informatics, Xiamen University ²Bytedance

³Key Laboratory of Digital Protection and Intelligent Processing of Intangible Cultural Heritage of Fujian and Taiwan (Xiamen University), Ministry of Culture and Tourism, China

⁴Dalian University of Technology

liyankang@stu.xmu.edu.cn huangluyang@bytedance.com jssu@xmu.edu.cn

Abstract

We present a large-scale video subtitle translation dataset, BIGVIDEO, to facilitate the study of multi-modality machine translation. Compared with the widely used HOW2 and VATEX datasets, BIGVIDEO is more than 10 times larger, consisting of 4.5 million sentence pairs and 9,981 hours of videos. We also introduce two deliberately designed test sets to verify the necessity of visual information: AMBIGUOUS with the presence of ambiguous words, and UNAMBIGUOUS in which the text context is self-contained for translation. To better model the common semantics shared across texts and videos, we introduce a contrastive learning method in the cross-modal encoder. Extensive experiments on the BIGVIDEO show that: *a*) Visual information consistently improves the NMT model in terms of BLEU, BLEURT, and COMET on both AMBIGUOUS and UNAMBIGUOUS test sets. *b*) Visual information helps disambiguation, compared to the strong text baseline on terminology-targeted scores and human evaluation. Dataset and our implementations are available at <https://github.com/DeepLearnXMU/BigVideo-VMT>.

1 Introduction

Humans are able to integrate both language and visual context to understand the world. From the perspective of NMT, it is also much needed to make use of such information to approach human-level translation abilities. To facilitate Multimodal Machine Translation (MMT) research, a number of datasets have been proposed including image-guided translation datasets (Elliott et al., 2016; Gella et al., 2019; Wang et al., 2022) and video-guided translation datasets (Sanabria et al., 2018; Wang et al., 2019; Li et al., 2022b).

*Equal contribution. Work was done while Liyan Kang was interning at ByteDance.

†Corresponding author.



Source Subtitle: **Clear shot** could also get you used to the air flow of the **courts**.

Ground Truth: **高远球**也可以让你习惯**场馆**的空气阻力。

System w/o Video: **清晰的镜头**也可以让你习惯**法庭**的气流。

System w/ Video: **清晰的击球**也会让你习惯**球场**的气流。

Figure 1: An example with semantic ambiguity in BIGVIDEO. The phrases with semantic ambiguity are highlighted in **red**. The wrong translations are in **blue** and the correct translations are in **yellow**.

However, the conclusion about the effects of visual information is still unclear for MMT research (Caglayan et al., 2019). Previous work has suggested that visual information is only marginally beneficial for machine translation (Li et al., 2021; Caglayan et al., 2021), especially when the text context is not complete. The most possible reason is that existing datasets focus on captions describing images or videos, which are not large and diverse enough. The text inputs are often simple and sufficient for translation tasks (Wu et al., 2021). Take the widely used Multi30K as an example. Multi30K consists of only 30K image captions, while typical text translation systems are often trained with several million sentence pairs.

We argue that studying the effects of visual contexts in machine translation requires a large-scale and diverse data set for training and a real-world and complex benchmark for testing. To this end, we propose BIGVIDEO, a large-scale video subtitle translation dataset. We collect human-written subtitles from two famous online video platforms, Xigua

and YouTube. BIGVIDEO consists of 155 thousand videos and 4.5 million high-quality parallel sentences in English and Chinese. We highlight the key features of BIGVIDEO as follows: *a)* The size of BIGVIDEO bypasses the largest available video machine translation dataset HOW2 and VATEX by one order of magnitude. *b)* To investigate the need for visual information, two test sets are annotated by language experts, referred as AMBIGUOUS and UNAMBIGUOUS. In AMBIGUOUS, the source input is not sufficient enough and requires videos to disambiguate for translation. The experts also labelled unambiguous words to help evaluate whether the improvement comes from visual contexts. In UNAMBIGUOUS, actions or visual scenes in the videos are mentioned in the subtitles but source sentences are self-contained for translation.

To make the most of visual information for MMT, we propose a unified encoder-decoder framework for MMT. The model has a cross-modal encoder that takes both videos and texts as inputs. Motivated by recent work on cross-modal learning (Li et al., 2020; Qi et al., 2020; Xia et al., 2021), we also introduce a contrastive learning objective to further bridge the representation gap between the text and video and project them in a shared space. As such, the visual information can potentially contribute more to the translation model.

We conduct extensive experiments on the proposed benchmark BIGVIDEO and report the results on BLEU (Papineni et al., 2002), BLEURT (Selam et al., 2020), COMET (Rei et al., 2020), terminology-targeted metrics and human evaluation. We also introduce the large scale WMT19 training data, which contains 20.4M parallel sentences to build the strong baseline model. The experiments show that visual contexts consistently improve the performance of both the AMBIGUOUS and UNAMBIGUOUS test set over the strong text-only model. The finding is slightly different with previous studies and address the importance of a large scale and high-quality video translation data. Further, the contrastive learning method can further boost the translation performance over other visual-guided models, which shows the benefits of closing the representation gap of texts and videos.

2 Related Work

Video-guided Machine Translation. The VATEX dataset has been introduced for video-guided machine translation task (Wang et al., 2019). It con-

tains 129K bilingual captions paired with video clips. However, as pointed out by Yang et al. (2022), captions in VATEX have sufficient information for translation, and models trained on VATEX tend to ignore video information. Beyond captions, Sanabria et al. (2018) considers video subtitles to construct the HOW2 dataset. HOW2 collects instructional videos from YouTube and obtains 186K bilingual subtitle sentences. To construct a challenging VMT dataset, Li et al. (2022b) collect 39K ambiguous subtitles from movies or TV episodes to build VISA. However, both HOW2 and VISA are limited on scale and diversity, given the training needs of large models. In contrast, we release a larger video subtitle translation dataset, with millions of bilingual ambiguous subtitles, covering all categories on YouTube and Xigua platforms.

To leverage video inputs in machine translation models, Hirasawa et al. (2020) use pretrained models such as ResNet (He et al., 2016), FasterRCNN (Ren et al., 2015) and I3D (Carreira and Zisserman, 2017). An additional attention module is designed in the RNN decoder to fuse visual features. To better learn temporal information in videos, (Gu et al., 2021) propose a hierarchical attention network to model video-level features. Different from previous work, we use a unified encoder to learn both video and text features. Specifically, a contrastive learning objective is adopted to learn cross-modal interaction.

Image-guided Machine Translation. Images as additional inputs have long been used for machine translation (Hitschler et al., 2016). For neural models, several attempts have been focused on enhancing the sequence-to-sequence model with strong image features (Elliott and Kádár, 2017; Yao and Wan, 2020; Lin et al., 2020; Yin et al., 2020; Su et al., 2021; Li et al., 2022a; Lin et al., 2020; Zhu et al., 2022; Lan et al., 2023). However, Li et al. (2021) and Wu et al. (2021) point out that images in Multi30K provide little information for translation. In this work, we focus on videos as additional visual inputs for subtitle translation. Videos illustrate objects, actions, and scenes, which contain more information compared to images. Subtitles are often in spoken language, which contains inherent ambiguities due to multiple potential interpretations (Mehrabi et al., 2022). Hence, our dataset can be a complement to existing MMT datasets.

Dataset	Text		Video		
	# Sent	Len.	# Video	# Clip	Sec.
VISA	35K	7.0	2K	35K	10.0
VATEX	129K	15.2	25K	25K	10.0
HOW2	186K	20.6	13K	186K	5.8
BIGVIDEO	4.5M	22.8	156K	4.5M	8.0

Table 1: Statistics of the BIGVIDEO and existing video-guided machine translation datasets. The size of the BIGVIDEO bypasses the size of largest available datasets by one order of magnitude. For the text length (Len.), we report average length of source sentences (English) for fair comparison. For the video length (Sec.), we report average seconds of videos for fair comparison.

3 Dataset

We present BIGVIDEO, consisting of 150 thousand unique videos (9,981 hours in total) with both English and Chinese subtitles. The videos are collected from two popular online video platforms, YouTube and Xigua. All subtitles are human-written. Table 1 lists statistics of our dataset and existing video-guided translation datasets. Among existing datasets, our dataset is significantly larger, with more videos and parallel sentences.

3.1 BIGVIDEO Dataset

To obtain high-quality video-subtitle pairs, we collect videos with both English and Chinese subtitles from *YouTube*¹ and *Xigua*². Both two platforms provide three types of subtitles: 1) *creator* which is uploaded by the creator, 2) *auto-generate* which is generated by the automatic speech recognition model and 3) *auto-translate* which is produced by machine translation model. We only consider videos with both English and Chinese subtitles uploaded by creators in order to obtain high-quality parallel subtitles. These videos and subtitles are often created by native or fluent English and Chinese speakers. In total, we collect 13K videos (6K hours in total) from YouTube and 2K videos from Xigua (3.9K hours in total).

Preprocessing. We first re-segment English subtitles into full sentences. To ensure the quality of parallel subtitles, we use quality estimation scores (e.g., the COMET score) to filter out low-quality pairs. More details are provided in Appendix B.1. Ultimately, 3.3M sentences paired with video clips are kept for YouTube, and 1.2M for Xigua. The average lengths of English and Chinese sentences

¹<https://www.youtube.com>

²<https://www.ixigua.com>

Source	Fluency \uparrow	Translation Quality \uparrow
YOUTUBE	4.81	4.11
XIGUA	4.60	4.20

Table 2: Human evaluation results on BIGVIDEO. The **fluency** score (1-5) measures whether source sentences (English) are fluent and error-free and the **translation quality** (1-5) measures whether sentence pairs are equivalent in meaning. Inter-annotator agreement with Krippendorff’s α for all columns: 0.76, 0.82.




Figure 2: The numbers of unique n -grams and POS tags in BIGVIDEO exceed existing VMT datasets by one order. Our dataset is lexically richer than VATEX and HOW2. For all datasets, we report statistics of source sentences (English) for fair comparison.

are 17.6 and 15.4 words for YouTube, 37.7 and 32.4 words for Xigua.

3.2 Dataset Analysis

Quality Evaluation. To assess the quality of text pairs, we randomly select 200 videos from each source and recruit seven annotators to rate the quality of subtitle pairs. For each video, we randomly select at most 20 clips for evaluation. All annotators are fluent in English and Chinese. After watching video clips and subtitles, human annotators are asked to rate subtitle pairs from 1 (worst) to 5 (best) on **fluency**—whether the source sentence (English) is fluent and grammatically correct, and **translation quality**—whether the Chinese subtitle is semantically equivalent to the English subtitle. Detailed guidelines are provided in Appendix F.

From Table 2, English sentences from both YouTube and Xigua have an average of 4.8 and 4.6 fluency scores, which shows that English subtitles are fluent and rarely have errors. In terms of translation quality, we find more than 96 percent of the pairs are equivalent or mostly-equivalent, with only minor differences (e.g., style).

Diversity Evaluation. In addition to the size and

Test set	Number	Length	# phrases
AMBIGUOUS	877	28.61	745
UNAMBIGUOUS	1,517	27.22	—

Table 3: Statistics of our test sets. We report number of samples, average length and number of ambiguous terms in two test sets.

quality, diversity is also critical for modeling alignments between parallel texts (Tiedemann, 2012). Prior work calculates unique n -grams and part-of-speech (POS) tags to evaluate linguistic complexity (Wang et al., 2019). Besides word-level metrics, we use video category distribution to assess video-level diversity.

Since the source text of our dataset, VATEX and HOW2 are in English, we compare unique n -grams and POS tags on the source texts. For unique POS tags, we compare four most common types: verb, noun, adjective and adverb. In Figure 2, our data from both XIGUA and YOUTUBE have substantially more unique n -grams and POS tags than VATEX and HOW2. Evidently, our dataset covers a wider range of actions, objects and visual scenes.

To evaluate video-level diversity, we compare category distributions among three datasets. The YouTube platform classifies videos into 15 categories. Since videos collected from the Xigua platform do not have category labels, we train a classifier on the YouTube data to label them. Details of the classifier are in Appendix B.2. Figure 3 depicts the distributions of three datasets. While both VATEX and HOW2 have a long-tail distribution on several categories (e.g., “Nonprofits & Activism” and “News & Politic”), BIGVIDEO has at least 1,000 videos in each category, which forms a more diverse training set.

3.3 Test Set Annotation Procedure

Subtitles often contain semantic ambiguities (Gu et al., 2021), which can be potentially solved by watching videos. In order to study “*How visual contexts benefit machine translation*”, we create two test sets: AMBIGUOUS contains ambiguous subtitles that videos provide strong disambiguation signal, while UNAMBIGUOUS consists of self-contained subtitles that videos are related but subtitles themselves contain enough context for translation. Statistics of two test sets are listed in Figure 3.

We randomly sample 200 videos from each of Xigua and YouTube and hire four professional





Figure 3: Category distribution on BIGVIDEO, VATEX and HOW2. BIGVIDEO covers a wide range of domains.



English Subtitle:That 's a good **chip** .
Chinese Subtitle:那款**玉米片**很好吃。
Ambiguous Annotation: chip → 玉米片

Figure 4: An example from our AMBIGUOUS test set. The ambiguous term “chip” is in red.

speakers in both English and Chinese to annotate the test set. Annotators are first asked to remove sentences which are not related to videos. In this step, we filter out about twenty percent of sentences. Annotators are then asked to re-write the Chinese subtitle if it is not perfectly equivalent to the English subtitle. Next, we ask the annotators to distinguish whether the source sentence contains semantic ambiguity. Specifically, annotators are instructed to identify ambiguous words or phrases in both English and Chinese sentences, as illustrated in Figure 4. We finally obtain 2394 samples in our test set. 36.6% of the sentences are in the AMBIGUOUS and 63.4% of the sentences are in the UNAMBIGUOUS. In the AMBIGUOUS, we annotate 745 ambiguous terms. The statistics indicate that videos play important roles in our dataset. Annotation instructions and detailed procedures are provided in Appendix F.

4 Method

4.1 Model

To better leverage videos to help translation, we present our video-guided machine translation model, as displayed in Figure 5. Our model can be seamlessly plugged into the pretrained NMT model, which can benefit from large-scale parallel training data. Importantly, we design a contrastive learning objective to further drive the translation model to learn shared semantics between videos and text.

Cross-modal Encoder. Our model takes both videos and text as inputs. Text inputs are first represented as a sequence of tokens \mathbf{x} and then converted to word embeddings through the embedding layer. Video inputs are represented as a sequence of continuous frames \mathbf{v} . We use a pretrained encoder to extract frame-level features, which is frozen for all experiments. Concretely, we apply the linear projection to obtain video features with the same dimension as text embeddings. To further model temporal information, we add positional embeddings to video features, followed by the layer normalization. Video features \mathbf{v}^{emb} and text embeddings \mathbf{x}^{emb} are then concatenated and fed into the Transformer encoder.

Text Decoder. Our decoder is the original Transformer decoder, which generates tokens autoregressively conditioned on the encoder outputs. We consider the cross entropy loss as a training objective:

$$\mathcal{L}_{CE} = - \sum_i^N \log P(\mathbf{y}_i | \mathbf{v}_i, \mathbf{x}_i), \quad (1)$$

where \mathbf{y}_i denotes the text sequence in the target language for the i -th sample in a batch of N samples.

4.2 Contrastive Learning Objective

In order to learn shared semantics between videos and text, we introduce a cross-modal contrastive learning (CTR)-based objective. The idea of the CTR objective is to bring the representations of video-text pairs closer and push irrelevant ones further.

Formally, given a positive text-video pair $(\mathbf{x}_i, \mathbf{v}_i)$, we use remaining $N - 1$ irrelevant text-video pairs $(\mathbf{x}_i, \mathbf{v}_j)$ in the batch as negative samples.

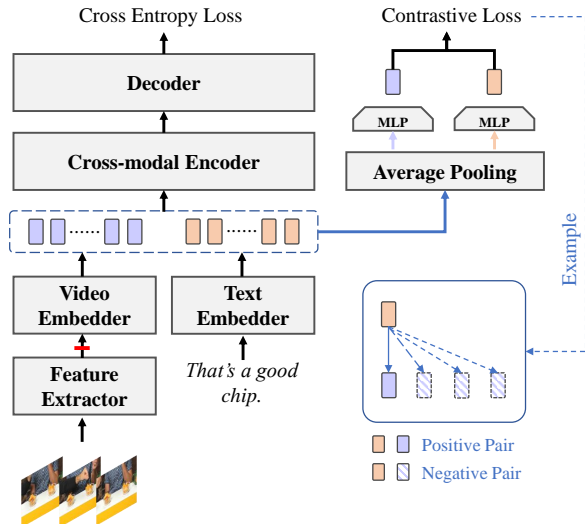


Figure 5: An illustration of our machine translation system. An example of our contrastive learning is presented in the blue box.

The contrastive learning objective (Sohn, 2016) is:

$$\mathcal{L}_{CTR} = - \sum_{i=1}^N \log \frac{\exp(\text{sim}(\mathbf{x}_i^p, \mathbf{v}_i^p)/\tau)}{\sum_{j=1, j \neq i}^N \exp(\text{sim}(\mathbf{x}_i^p, \mathbf{v}_j^p)/\tau)}, \quad (2)$$

where \mathbf{x}_i^p and \mathbf{v}_i^p are representations for the text and video, $\text{sim}(\cdot)$ is the cosine similarity function and the temperature τ is used to control the strength of penalties on hard negative samples (Wang and Liu, 2021).

Text and Video Representations. Importantly, since videos and subtitles are weakly aligned on the temporal dimension (Miech et al., 2019), we first average video embeddings and text embeddings in terms of the time dimension. Concretely, we apply two *projection heads* ("MLP" in Figure 5) to map representations to the same semantic space (Chen et al., 2020).

In the end, we sum up the two losses to obtain the final loss:

$$\mathcal{L} = \mathcal{L}_{CE} + \alpha \mathcal{L}_{CTR}, \quad (3)$$

where α is a hyper-parameter to balance the two loss items.

5 Experiments

5.1 Experimental Setup

Implementation Details. We evaluate our method on three video translation datasets: VATEX, HOW2

System	BLEU			COMET			BLEURT		
	All	Amb.	Unamb.	All	Amb.	Unamb.	All	Amb.	Unamb.
<i>w/o pretraining</i>									
TEXT-ONLY	43.97	43.59	44.19	37.31	33.06	39.77	61.11	59.68	61.93
GATED FUSION	44.33	44.12	44.45	37.88	34.47	39.85	61.25	60.19	61.86
SELECTIVE ATTN	44.39	44.20	44.51	38.48	34.84	40.59	61.37	60.30	62.07
Ours									
+ ViT	44.26	44.10	44.37	38.13	34.75	40.08	61.24	60.29	61.78
+ SLOWFAST	44.21	44.12	44.26	37.81	34.99	39.44	61.22	60.28	61.77
+ ViT + CTR	44.45	44.48	44.40	38.15	35.76	39.54	61.36	60.72	61.73
+ SLOWFAST + CTR	44.44	44.20	44.58	38.37	35.18	40.22	61.31	60.41	61.82
<i>w/ pretraining</i>									
TEXT-ONLY	44.45	43.89	44.79	38.36	33.40	41.23	61.41	59.85	62.31
+ ViT + CTR	44.83	44.62	44.96	39.44	36.42	41.19	61.76	60.75	62.34
+ SLOWFAST + CTR	44.77	44.43	44.97	39.26	36.03	41.12	61.71	60.52	62.40

Table 4: sacreBLEU(%), COMET(%) and BLEURT(%) scores on BIGVIDEO testset. We report results on AMBIGUOUS (**Amb.**), UNAMBIGUOUS (**Unamb.**) and whole test set (**All**). “+ CTR” denotes our cross-model framework with the contrastive learning objective. All results are mean values of five different random seeds. Complete results with standard deviations can be seen in Appendix A. The Best result in each group is in **bold**. The Best result in each column is in **red**.

and our proposed dataset BIGVIDEO. More dataset details can be found in Appendix C.1.

Our code is based on the *fairseq* toolkit (Ott et al., 2019). The Transformer-base model follows (Vaswani et al., 2017). Both encoder and decoder have 6 layers, 8 attention heads, hidden size = 512, and FFN size = 2048. We utilize post-layer normalization for all models. On VATEX, we follow the Transformer-small setting from Wu et al. (2021) for better performance, 6 layers for encoder/decoder, hidden size = 512, FFN size = 1024 and attention heads = 4.

All experiments are done on 8 NVIDIA V100 GPUS with mixed-precision training (Das et al., 2018), where the batch assigned to each GPU contains 4,096 tokens. More training details can be found in Appendix C.2. We stop the training if the performance on the validation set does not improve for ten consecutive epochs. The running time is about 64 GPU hours for our system. During the inference, the beam size and the length penalty are set to 4 and 1.0. We apply byte pair encoding (BPE) with 32K merge operations to preprocess sentences of our dataset. During training and testing, we uniformly sample a maximum of 12 frames as the video input. The text length is limited to 256. For the contrastive learning loss, we set α to 1.0 and τ to 0.002. The choices of hyper-parameters are in Appendix D.

For video features, we extract 2D features and 3D features to compare their effects. Concretely, we experiment with two pretrained models to extract the video feature: *a*) The vision transformer (ViT) (Dosovitskiy et al., 2021) which extracts frame-level features. *b*) The SlowFast model (SLOWFAST) which extracts video-level features (Feichtenhofer et al., 2019). For 2D features, we first extract images at a fixed frame rate (3 frames per second). Then we utilize pretrained Vision Transformer³ (ViT) to extract 2D video features into 768-dimensional vectors. Here the representation of $[CLS]$ token is considered as the global information of one frame. For 3D features, we extract 2304-dimensional SlowFast⁴ features at 2/3 frames per second.

Baselines and Comparisons. For baselines, we consider the base version of the Transformer (TEXT-ONLY), which only takes texts as inputs. For comparisons, since most recent MMT studies focus on image-guided machine translation, we implement two recent image-based MMT models: *a*) The gated fusion model (GATED FUSION) which fuses visual representations and text representations with a gate mechanism (Wu et al., 2021). *b*) The selective attention model (SELECTIVE ATTN) which uses a single-head attention to connect text and im-

³The model architecture is vit_base_patch16_224.

⁴SLOWFAST_8x8_R50.

age representations (Li et al., 2022a). We extract image features using ViT and obtain the visual feature by averaging image features on the temporal dimension. The visual feature is then fused with the text representations which is the same as original GATED FUSION and SELECTIVE ATTN. For HOW2 and VATEX, we additionally include the baseline models provided by the original paper.

Evaluation Metrics. We evaluate our results with the following three metrics: detokenized sacreBLEU⁵, COMET⁶ (Rei et al., 2020) and BLEURT⁷ (Sellam et al., 2020). In order to evaluate whether videos are leveraged to disambiguate, we further consider three terminology-targeted metrics (Alam et al., 2021):

- **Exact Match:** the accuracy over the annotated ambiguous words. If the correct ambiguous words or phrases appear in the output, we count it as correct.
- **Window Overlap:** indicating whether the ambiguous terms are placed in the correct context. For each target ambiguous term, a window is set to contain its left and right words, ignoring stopwords. We calculate the percentage of words in the window that are correct. In practice, we set window sizes to 2 (Window Overlap-2) and 3 (Window Overlap-3).
- **Terminology-biased Translation Edit Rate (1-TERm):** modified translation edit rate (Snover et al., 2006) in which words in ambiguous terms are set to 2 edit cost and others are 1.

6 Results

6.1 Main Results

Videos Consistently Improve the NMT Model.

As displayed in Table 4, on BIGVIDEO, our models equipped with videos obtain higher automatic scores. This indicates the benefit of using videos as additional inputs. Notably, our model trained with the additional contrastive learning objective yields better scores compared to the variant trained only with the cross entropy loss. This signifies that our contrastive learning objective can guide better acquisition of video inputs. Furthermore, we

⁵<https://github.com/mjpost/sacrebleu>.

Signature: BLEU+case.mixed+numrefs.1+smooth.exp+tok.zh+version.1.5.1 (Post, 2018)

⁶<https://github.com/Unbabel/COMET>. We use the default wmt20-comet-da.

⁷<https://github.com/google-research/bleurt>. We use the BLEURT-20.

System	Exact Match	Window Overlap-2	Window Overlap-3	1-TERm
<i>w/o pre-training</i>				
TEXT-ONLY	23.03	14.22	14.28	49.50
GATED FUSION	23.68	14.60	14.76	49.68
SELECTIVE ATTN	23.66	14.95	15.08	49.77
Ours				
+ ViT	24.27	15.09	15.27	49.78
+ SLOWFAST	24.05	14.97	15.13	49.32
+ ViT + CTR	25.02	15.56	15.77	49.91
+ SLOWFAST + CTR	24.08	15.04	15.12	49.32
<i>w/ pre-training</i>				
TEXT-ONLY	22.71	14.35	14.37	49.73
+ ViT + CTR	24.30	15.17	15.39	50.28
+ SLOWFAST + CTR	23.62	14.76	14.90	50.04

Table 5: Terminology-targeted evaluation on AMBIGUOUS test set. Complete results with standard deviations can be seen in Appendix A.

find image-based pretrained model ViT and video-based pretrained model SLOWFAST yield comparable results, indicating that two vision features perform equally well on BIGVIDEO.

Noticeably, compared to the text-only baseline, our models trained with the CTR objective achieves larger gain on AMBIGUOUS than that on UNAMBIGUOUS. This demonstrates that it is more difficult to correctly translate sentences of AMBIGUOUS, while taking videos as additional inputs can help the model generate better translations.

To better study the role of videos in translation, we introduce additional training data to build a stronger NMT baseline. We introduce the WMT19 Zh-En dataset with 20.4M parallel sentences for pretraining. We aim to answer: *how will the model perform if more text data are included?*

As displayed in Table 4, *Model with video inputs outperforms the strong NMT baseline*. Pretraining on large corpus benefits models on BIGVIDEO. However, we find improvements mainly come from the UNAMBIGUOUS. This shows that videos play more crucial roles in AMBIGUOUS, which suggests that BIGVIDEO can serve as a valuable benchmark for studying the role of videos in MMT research.

Videos Help Disambiguation. We further evaluate the model ability of disambiguation. We present results on terminology-targeted metrics in Table 5. First, our systems with video features achieve consistent improvements both on exact match and window overlap metrics compared to the text-only variant, indicating that models augmented by video inputs correctly translate more ambiguous words

Systems	Score	Win \uparrow	Tie	Lose \downarrow
AMBIGUOUS				
TEXT-ONLY	3.48	—	—	—
+ ViT + CTR	3.53	19.3%	65.3%	15.3%
UNAMBIGUOUS				
TEXT-ONLY	3.71	—	—	—
+ ViT + CTR	3.72	24%	51.3%	21.7%

Table 6: Human evaluation results on randomly sampled set. “Win”/“Tie”/“Lose” stands for the percentage of translations where our system is better than, tied with, or worse than the text-only system. Inter-annotator agreement with Krippendorff’s α for the translation quality score is 0.53.

and place them in the proper contexts. It is also worth noticing that our system with pretraining achieves better scores compared to the strong text baseline, which further highlights the importance of video inputs. Moreover, we find it hard to correctly translate ambiguous words since the best exact match score is 25.02%, which suggests that our AMBIGUOUS set is challenging.

Video-augmented Model Improves Translation Quality. We further conduct human evaluation to analyze the translation quality. We randomly pick 100 sentences from the AMBIGUOUS and the UNAMBIGUOUS respectively and recruit three human judges for evaluation. For each sentence, the judges read the source sentence and two candidate translations, which are from TEXT-ONLY and our model + ViT + CTR. The judges are required to rate each candidate on a scale of 1 to 5 and pick the better one. Detailed guidelines are in Appendix G.

From Table 6, we can see *our system with video inputs are more frequently rated as better translation than the text-only model on both AMBIGUOUS and UNAMBIGUOUS test sets*. This echoes automatic evaluations and implies that taking videos as inputs improve translation quality. Moreover, overall scores on UNAMBIGUOUS are better than those on AMBIGUOUS, which demonstrates that AMBIGUOUS is more challenging.

6.2 Incongruent Decoding

In this section, we explore *whether visual inputs contribute to the translation model*. Following (Caglayan et al., 2019; Li et al., 2022a), we use incongruent decoding to probe the need for visual modality on BIGVIDEO. During the inference, we replace the original video with a mismatched video




Figure 6: The results of incongruent decoding of the + ViT + CTR model. “Congruent” denotes the original results and “Incongruent” denotes the results of incongruent decoding where videos are replaced by wrong ones. “Match” represents the exact match score and “Overlap” stands for the “Window Overlap-2”.

for each sentence. As shown in Figure 6, on AMBIGUOUS and UNAMBIGUOUS, we observe that all automatic metrics of our system drop significantly with incongruent decoding, suggesting the effectiveness of leveraging videos as inputs. Interestingly, we also find that the drop of the BLEU and COMET scores is larger on AMBIGUOUS than that on UNAMBIGUOUS, which further proves our point that videos are more crucial for disambiguation.

6.3 Results on Public Datasets

Next, we conduct experiments on public datasets, VATEX and HOW2. Results are displayed in Table 7. On HOW2, our best system achieves higher BLEU score compared to the text-only model. However, the text-only model achieves best COMET and BLEURT, compared to all systems that take videos as inputs. On VATEX, our model with SLOWFAST features also achieves the highest scores on three evaluation metrics, compared to the text-only model and comparisons. Notably, the model with SLOWFAST features is significantly better than models with ViT features, which is probably because VATEX focuses on human actions and the SLOWFAST model is trained on the action recognition dataset. However, the performance gap between the TEXT-ONLY and our model + SLOWFAST + CTR is marginal. After we introduce 20M external MT data, we observe that the TEXT-ONLY and our best system are comparable on automatic metrics. Since the cross-modal encoder often requires large-scale paired videos and text to train robust representations, our model does not achieve huge performance gain on VATEX and HOW2. We hope our BIGVIDEO dataset can serve as a comple-

System	BLEU	COMET	BLEURT
How2			
<i>w/o pretraining</i>			
TEXT-ONLY	57.57	65.95	72.52
Sanabria et al. (2018)	54.40	—	—
GATED FUSHION	57.65	65.12	72.26
SELECTIVE ATTN	57.51	65.77	72.35
Ours			
+ ViT + CTR	57.95	65.53	72.46
+ SLOWFAST + CTR	57.78	65.58	72.41
VATEX			
<i>w/o pretraining</i>			
TEXT-ONLY	35.01	15.32	56.99
Wang et al. (2019)	30.11	4.50	53.85
GATED FUSHION	33.79	13.55	55.66
SELECTIVE ATTN	34.25	13.55	56.80
Ours			
+ ViT + CTR	34.84	12.44	56.25
+ SLOWFAST + CTR	35.15	15.65	57.06
<i>w/ pretraining</i>			
TEXT-ONLY	37.57	25.22	59.33
+ ViT + CTR	37.34	24.07	58.87
+ SLOWFAST + CTR	37.58	25.05	59.20

Table 7: Experimental results on HOW2 and VATEX. Complete results with standard deviations can be seen in Appendix A.

ment to existing video-guided machine translation datasets.

7 Conclusion

In this paper, we present BIGVIDEO—a large-scale video subtitle Translation dataset for multi-modal machine translation. We collect 155 thousand videos accompanied by over 4.5 million bilingual subtitles. Specially, we annotate two test subsets: AMBIGUOUS where videos are required for disambiguation and UNAMBIGUOUS where text contents are self-contained for translation. We also propose a cross-modal encoder enhanced with a contrastive learning objective to build cross-modal interaction for machine translation. Experimental results prove that videos consistently improve the NMT model in terms of the translation evaluation metrics and terminology-targeted metrics. Moreover, human annotators prefer our system outputs, compared to the strong text-only baseline. We hope our BIGVIDEO dataset can facilitate the research of multi-modal machine translation.

Limitations

BIGVIDEO is collected from two video platforms Xigua and YouTube. All videos are publicly available. However, some videos may contain user information (e.g., portraits) or other sensitive information. Similar to VATEX and HOW2, we will release our test set annotation and the code to reproduce our dataset. For videos without copyright or sensitive issues, we will make them public but limit for research, and non-commercial use (We will require dataset users to apply for access). For videos with copyright or sensitive risks, we will provide ids, which can be used to download the video. This step will be done under the instruction of professional lawyers.

Though we show that our model with video inputs helps disambiguation, we find that our model could yield incorrect translation due to the lack of world knowledge. For example, model can not distinguish famous table tennis player Fan Zhengdong and give correct translation. We find this is due to video pretrained models are often trained on action dataset (e.g., Kinetics-600 (Long et al., 2020)) and hardly learn such world knowledge. In this work, we do not further study methods that leverage world knowledge.

Ethical Considerations

Collection of BIGVIDEO. We comply with the terms of use and copyright policies of all data sources during collection from the YouTube and Xigua platform. User and other sensitive information is not collected to ensure the privacy of video creators. The data sources are publicly available videos and our preprocessing procedure does not involve privacy issues. For all annotation or human evaluation mentioned in the paper, we hire seven full-time professional translators in total and pay them with market wage. All of our annotators are graduates.

Potential Risks of BIGVIDEO and our model. While BIGVIDEO consists of high-quality parallel subtitles, we recognize that our data may still contain incorrect samples. Our model may as well generate degraded or even improper contents. As our dataset is based on YouTube or Xigua videos, models trained on our dataset might be biased towards US or Chinese user perspective, which could yield outputs that are harmful to certain populations.

Acknowledgements

The project was supported by the National Key Research and Development Program of China (No. 2020AAA0108004), National Natural Science Foundation of China (No. 62276219) and Natural Science Foundation of Fujian Province of China (No. 2020J06001). We also thank the reviewers for their insightful comments.

References

- Md Mahfuz Ibn Alam, Antonios Anastasopoulos, Laurent Besacier, James Cross, Matthias Gallé, Philipp Koehn, and Vassilina Nikoulina. 2021. On the evaluation of machine translation for terminology consistency. *arXiv preprint arXiv:2106.11891*.
- Mikel Artetxe and Holger Schwenk. 2019. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *Transactions of the Association for Computational Linguistics*.
- Ozan Caglayan, Menekse Kuyu, Mustafa Sercan Amac, Pranava Madhyastha, Erkut Erdem, Aykut Erdem, and Lucia Specia. 2021. Cross-lingual visual pre-training for multimodal machine translation. *arXiv preprint arXiv:2101.10044*.
- Ozan Caglayan, Pranava Madhyastha, Lucia Specia, and Loïc Barrault. 2019. Probing the need for visual context in multimodal machine translation. In *Proc. of NAACL*.
- João Carreira and Andrew Zisserman. 2017. Quo vadis, action recognition? A new model and the kinetics dataset. In *Proc. of CVPR*.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020. A simple framework for contrastive learning of visual representations. In *Proc. of ICML*.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. Xnli: Evaluating cross-lingual sentence representations. In *Proc. of EMNLP*.
- Dipankar Das, Naveen Mellempudi, Dheevatsa Mudigere, Dhiraj D. Kalamkar, Sasikanth Avancha, Kunal Banerjee, Srinivas Sridharan, Karthik Vaidyanathan, Bharat Kaul, Evangelos Georganas, Alexander Heinecke, Pradeep Dubey, Jesús Corbal, Nikita Shustrov, Roman Dubtsov, Evarist Fomenko, and Vadim O. Pirogov. 2018. Mixed precision training of convolutional neural networks using integer operations. In *Proc. of ICLR*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proc. of ICLR*.
- Desmond Elliott, Stella Frank, Khalil Sima'an, and Lucia Specia. 2016. Multi30K: Multilingual English-German image descriptions. In *Proc. of the 5th Workshop on Vision and Language*.
- Desmond Elliott and Ákos Kádár. 2017. Imagination improves multimodal translation. In *Proc. of IJCNLP*.
- Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. 2019. Slowfast networks for video recognition. In *Proc. of ICCV*.
- Spandana Gella, Desmond Elliott, and Frank Keller. 2019. Cross-lingual visual verb sense disambiguation. In *Proc. of ACL*.
- Weiqi Gu, Haiyue Song, Chenhui Chu, and Sadao Kurohashi. 2021. Video-guided machine translation with spatial hierarchical attention network. In *Proc. of ACL: Student Research Workshop*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proc. of CVPR*.
- Tosho Hirasawa, Zhishen Yang, Mamoru Komachi, and Naoaki Okazaki. 2020. Keyframe segmentation and positional encoding for video-guided machine translation challenge 2020. *arXiv preprint arxiv:2006.12799*.
- Julian Hitschler, Shigehiko Schamoni, and Stefan Riezler. 2016. Multimodal pivots for image caption translation. In *Proc. of ACL*.
- Zhibin Lan, Jiawei Yu, Xiang Li, Jian Luan Wen Zhang, Bin Wang, Degen Huang, and Jin-song Su. 2023. Exploring better text image trans-

- lation with multimodal codebook. In *Proc. of ACL*.
- Bei Li, Chuanhao Lv, Zefan Zhou, Tao Zhou, Tong Xiao, Anxiang Ma, and JingBo Zhu. 2022a. On vision features in multimodal machine translation. In *Proc. of ACL*.
- Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. 2020. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *Proc. of AAAI*.
- Jiaoda Li, Duygu Ataman, and Rico Sennrich. 2021. Vision matters when it should: Sanity checking multimodal machine translation models. In *Proc. of EMNLP*.
- Yihang Li, Shuichiro Shimizu, Weiqi Gu, Chenhui Chu, and Sadao Kurohashi. 2022b. VISA: An ambiguous subtitles dataset for visual scene-aware machine translation. In *Proc. of LREC*.
- Huan Lin, Fandong Meng, Jinsong Su, Yongjing Yin, Zhengyuan Yang, Yubin Ge, Jie Zhou, and Jiebo Luo. 2020. Dynamic context-guided capsule network for multimodal machine translation. In *Proc. of ACMMM*.
- Fuchen Long, Ting Yao, Zhaofan Qiu, Xinmei Tian, Jiebo Luo, and Tao Mei. 2020. Learning to localize actions from moments. In *Proc. of ECCV*.
- Ninareh Mehrabi, Palash Goyal, Apurv Verma, Jwala Dhamala, Varun Kumar, Qian Hu, Kai-Wei Chang, Richard S. Zemel, Aram Galstyan, and Rahul Gupta. 2022. Is the elephant flying? resolving ambiguities in text-to-image generative models. *arXiv preprint arxiv:2211.12503*.
- Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. 2019. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *Proc. of ICCV*.
- Jihyung Moon, Hyunchang Cho, and Eunjeong L. Park. 2020. Revisiting round-trip translation for quality estimation. In *Proc. of EACL*.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proc. of NAACL*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proc. of ACL*.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. In *Proc. of CMT*.
- Di Qi, Lin Su, Jia Song, Edward Cui, Taroon Bharti, and Arun Sacheti. 2020. Imagebert: Cross-modal pre-training with large-scale weak-supervised image-text data. *arXiv preprint arXiv:2001.07966*.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. COMET: A neural framework for MT evaluation. In *Proc. of EMNLP*.
- Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. Faster R-CNN: towards real-time object detection with region proposal networks. In *Proc. of NIPS*.
- Ramon Sanabria, Ozan Caglayan, Shruti Palaskar, Desmond Elliott, Loïc Barrault, Lucia Specia, and Florian Metze. 2018. How2: A large-scale dataset for multimodal language understanding. *arXiv preprint arxiv:1811.00347*.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. BLEURT: Learning robust metrics for text generation. In *Proc. of ACL*.
- Matthew G. Snover, Bonnie J. Dorr, Richard M. Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proc. of AMTA*.
- Kihyuk Sohn. 2016. Improved deep metric learning with multi-class n-pair loss objective. In *Proc. of NIPS*.
- Jinsong Su, Jinchang Chen, Hui Jiang, Chulun Zhou, Huan Lin, Yubin Ge, Qingqiang Wu, and Yongxuan Lai. 2021. Multi-modal neural machine translation with deep semantic interactions. *Inf. Sci.*
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in OPUS. In *Proc. of LREC*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proc. of NIPS*.

- Feng Wang and Huaping Liu. 2021. Understanding the behaviour of contrastive loss. In *Proc. of CVPR*.
- Josiah Wang, Josiel Figueiredo, and Lucia Specia. 2022. MultiSubs: A large-scale multimodal and multilingual dataset. In *Proc. of LREC*.
- Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuanfang Wang, and William Yang Wang. 2019. VateX: A large-scale, high-quality multilingual dataset for video-and-language research. In *Proc. of ICCV*.
- Zhiyong Wu, Lingpeng Kong, Wei Bi, Xiang Li, and Ben Kao. 2021. Good for misconceived reasons: An empirical revisiting on the need for visual context in multimodal machine translation. In *Proc. of ACL*.
- Qiaolin Xia, Haoyang Huang, Nan Duan, Dongdong Zhang, Lei Ji, Zhifang Sui, Edward Cui, Taroon Bharti, and Ming Zhou. 2021. Xgpt: Cross-modal generative pre-training for image captioning. In *Proc. of NLPCC*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive 897 pretraining for language understanding. In *Proc. of NIPS*.
- Zhishen Yang, Toshio Hirasawa, Mamoru Komachi, and Naoaki Okazaki. 2022. Why videos do not guide translations in video-guided machine translation? an empirical evaluation of video-guided machine translation dataset. *J. Inf. Process.*
- Shaowei Yao and Xiaojun Wan. 2020. Multimodal transformer for multimodal machine translation. In *Proc. of ACL*.
- Yongjing Yin, Fandong Meng, Jinsong Su, Chulun Zhou, Zhengyuan Yang, Jie Zhou, and Jiebo Luo. 2020. A novel graph-based multi-modal fusion encoder for neural machine translation. In *Proc. of ACL*.
- Yaoming Zhu, Zewei Sun, Shanbo Cheng, Yuyang Huang, Liwei Wu, and Mingxuan Wang. 2022. Beyond triplet: Leveraging the most data for multimodal machine translation. *arXiv preprint arXiv:2212.10313*.

A Complete Results

The Complete results with standard deviations can be found in Table 8, Table 9 and Table 10.

B Data Collection

B.1 Preprocessing

Subtitles are organized as a list of text chunks. Each chunk contains both English and Chinese lines and a corresponding timestamp. To obtain complete sentences, we start processing subtitles by merging chunks. Since English subtitles are often with strong punctuation marks, we greedily merge continuous segments (The start time of the second segment and the end time of the first segment are within 0.5 seconds) until an end mark is met at the end of the segment. To preserve context, we keep merging continuous sentences until a maximum time limit of 15 seconds is reached. Finally, we pair each merged segment with the video clip from the time interval corresponding to the segment.

English sentences from both YouTube and Xigua have an average of 4.6 fluency score, which shows that English subtitles are fluent and rarely have errors. In terms of translation quality, subtitles collected from Xigua have an average of 4.2 translation quality score, which indicates most of the subtitle pairs are equivalent or near-equivalent. In YouTube data, we find about 20 percent of sentence pairs are not equivalent or have major errors such as mistranslation or omission.

To remove low-quality pairs, we try three commonly-used quality estimation scores: 1) the COMET score, 2) the Euclidean distance based on the multilingual sentence embedding (Artetxe and Schwenk, 2019), and 3) the round-trip translation BLEU score (Moon et al., 2020). We filter out pairs if more than one score is lower than the threshold (set to 0.1, 4 and 20, respectively). On annotated samples, the average translation quality reaches 4.1 after cleaning.

B.2 Video Category Classifier Details

To construct a large-scale video-guided dataset, we collect videos from a variety of domains and categorized them into 15 classes based on their video categories in YouTube. We use the official youtube-dl⁸ toolkit to retrieve video categories and other metadata from YouTube. To ensure consistency

⁸<https://youtube-dl.org>

System	BLEU			COMET			BLEURT		
	ALL	AMBIGUOUS	UNAMBIGUOUS	ALL	AMBIGUOUS	UNAMBIGUOUS	ALL	AMBIGUOUS	UNAMBIGUOUS
<i>w/o pre-training</i>									
TEXT-ONLY	43.97 _{0.10}	43.59 _{0.29}	44.19 _{0.20}	37.31 _{0.34}	33.06 _{0.24}	39.77 _{0.56}	61.11 _{0.09}	59.68 _{0.21}	61.93 _{0.13}
GATED FUSHION + ViT	44.33 _{0.13}	44.12 _{0.33}	44.45 _{0.12}	37.88 _{0.32}	34.47 _{0.62}	39.85 _{0.35}	61.25 _{0.07}	60.19 _{0.08}	61.86 _{0.08}
SELECTIVE ATTN + ViT	44.39 _{0.13}	44.20 _{0.22}	44.51 _{0.10}	38.48_{0.19}	34.84 _{0.39}	40.59_{0.32}	61.37_{0.10}	60.30 _{0.24}	62.07_{0.13}
Ours									
+ ViT	44.26 _{0.20}	44.10 _{0.24}	44.37 _{0.21}	38.13 _{0.65}	34.75 _{0.79}	40.08 _{0.71}	61.24 _{0.17}	60.29 _{0.30}	61.78 _{0.20}
+ SLOWFAST	44.21 _{0.17}	44.12 _{0.27}	44.26 _{0.27}	37.81 _{0.38}	34.99 _{0.25}	39.44 _{0.53}	61.22 _{0.09}	60.28 _{0.08}	61.77 _{0.17}
+ ViT + CTR	44.45_{0.13}	44.48_{0.16}	44.40 _{0.11}	38.15 _{0.56}	35.76_{0.42}	39.54 _{0.75}	61.36 _{0.17}	60.72_{0.21}	61.73 _{0.16}
+ SLOWFAST + CTR	44.44 _{0.12}	44.20 _{0.12}	44.58_{0.17}	38.37 _{0.41}	35.18 _{0.71}	40.22 _{0.53}	61.31 _{0.10}	60.41 _{0.09}	61.82 _{0.14}
<i>w/ pre-training</i>									
TEXT-ONLY	44.45 _{0.11}	43.89 _{0.19}	44.79 _{0.16}	38.36 _{0.33}	33.40 _{0.29}	41.23_{0.54}	61.41 _{0.13}	59.85 _{0.16}	62.31 _{0.19}
+ ViT + CTR	44.83_{0.09}	44.62_{0.10}	44.96 _{0.11}	39.44_{0.51}	36.42_{0.41}	41.19 _{0.89}	61.76_{0.08}	60.75_{0.12}	62.34 _{0.17}
+ SLOWFAST + CTR	44.77 _{0.31}	44.43 _{0.15}	44.97_{0.45}	39.26 _{0.31}	36.03 _{0.52}	41.12 _{0.76}	61.71 _{0.17}	60.52 _{0.11}	62.40_{0.25}

Table 8: Complete sacreBLEU(%), COMET(%) and BLEURT(%) scores on BIGVIDEO testset. “+ CTR” denotes our cross-model framework with contrastive learning loss. “ALL” represents the results on the whole test set. All results are mean values of five different random seeds. The best result in each group is in **bold**.

System	Exact Match	Window Overlap-2	Window Overlap-3	1-TERm
<i>w/o pre-training</i>				
TEXT-ONLY	23.03 _{0.67}	14.22 _{0.41}	14.28 _{0.42}	49.50 _{0.13}
GATED FUSHION	23.68 _{0.52}	14.60 _{0.56}	14.76 _{0.47}	49.68 _{0.19}
SELECTIVE ATTN	23.66 _{0.47}	14.95 _{0.76}	15.08 _{0.65}	49.77 _{0.22}
Ours				
+ ViT	24.27 _{0.37}	15.09 _{0.47}	15.27 _{0.50}	49.78 _{0.41}
+ SLOWFAST	24.05 _{0.58}	14.97 _{0.77}	15.13 _{0.85}	49.32 _{0.44}
+ ViT + CTR	25.02_{0.74}	15.56_{0.65}	15.77_{0.55}	49.91_{0.20}
+ SLOWFAST + CTR	24.08 _{0.31}	15.04 _{0.14}	15.12 _{0.17}	49.32 _{0.18}
<i>w/ pre-training</i>				
TEXT-ONLY	22.71 _{0.72}	14.35 _{0.54}	14.37 _{0.61}	49.73 _{0.21}
+ ViT + CTR	24.30_{0.67}	15.17_{0.58}	15.39_{0.52}	50.28_{0.39}
+ SLOWFAST + CTR	23.62 _{0.58}	14.76 _{0.48}	14.90 _{0.36}	50.04 _{0.13}

Table 9: Complete terminology-targeted results on BIGVIDEO test set. All results are mean values of five different random seeds with standard deviations as subscripts. The Best result in each group is in **bold**.

between YouTube and Xigua videos, we train a category classifier to classify the video category tags of Xigua videos and those YouTube videos whose category information are missing. We train the category classifier using the English subtitles and category information of pre-labeled videos and use it to predict the category tags for the rest of the unlabeled videos. Specifically, we first group consecutive subtitles in a video by 5 and then concatenate them as input for the classifier. During inference, we predicted the category tags of groups in each video and obtain the video’s label by voting. The category classifier model was fine-tuned based on the pre-trained XLNet-large-cased model, which performed well on other classification tasks such as XNLI (Yang et al., 2019; Conneau et al., 2018). A statistical summary of the video categories of the train set can be found in Figure 3.

In addition, we also count the category tags of the videos in the test set, as listed in Table 11. Sim-

System	BLEU	COMET	BLEURT
<i>HOW2</i>			
<i>w/o pretraining</i>			
TEXT-ONLY	57.57 _{0.26}	65.95_{0.71}	72.52_{0.24}
Sanabria et al. (2018)	54.40	—	—
GATED FUSHION	57.65 _{0.35}	65.12 _{0.43}	72.26 _{0.27}
SELECTIVE ATTN	57.51 _{0.20}	65.77 _{0.93}	72.35 _{0.23}
Ours			
+ ViT + CTR	57.95_{0.24}	65.53 _{0.68}	72.46 _{0.25}
+ SLOWFAST + CTR	57.78 _{0.09}	65.58 _{0.71}	72.41 _{0.15}
<i>VATEX</i>			
<i>w/o pretraining</i>			
TEXT-ONLY	35.01 _{0.14}	15.32 _{0.45}	56.99 _{0.307}
Wang et al. (2019)	30.11 _{0.72}	4.50 _{0.81}	53.85 _{0.55}
GATED FUSHION	33.79 _{0.14}	13.55 _{0.42}	55.66 _{0.11}
SELECTIVE ATTN	34.25 _{0.30}	13.55 _{0.10}	56.80 _{0.11}
Ours			
+ ViT + CTR	34.84 _{0.25}	12.44 _{0.64}	56.25 _{0.25}
+ SLOWFAST + CTR	35.15_{0.24}	15.65_{0.35}	57.06_{0.06}
<i>w/ pretraining</i>			
TEXT-ONLY	37.57 _{0.35}	25.22 _{0.66}	59.33 _{0.09}
+ ViT + CTR	37.34 _{0.14}	25.07 _{1.21}	58.87 _{0.33}
+ SLOWFAST + CTR	37.58 _{0.15}	25.05 _{0.58}	59.20 _{0.13}

Table 10: Complete experimental results on HOW2 and VATEX. All results are mean values of five different random seeds with standard deviations as subscripts. The Best result in each group is in **bold**.

ilar to the train set, we directly obtain the category tags for most YouTube videos directly from YouTube and predict the tags of the remaining videos using the category classifier mentioned earlier. The statistics for video categories show that the videos in our dataset are diverse in terms of domain, both in the training and test sets. These statistics for the video categories provide a more comprehensive view of BIGVIDEO.

Category	Xigua	YouTube	All
People & Blogs	15	29	44
Entertainment	48	14	62
Howto & Sytle	11	37	48
Education	33	3	36
Travel & Events	14	4	18
Music	0	0	0
Gaming	13	6	19
Sports	5	1	6
Science & Tech	30	2	32
Comedy	0	3	3
Film & Animation	3	1	4
Pets & Animals	12	6	18
News & Politics	1	1	2
Autos & Vehicles	9	2	11
Activism	0	2	2
All	194	111	305

Table 11: Video category tags of our test set. For most of the YouTube videos, we obtain the category tags from YouTube official. For the rest of the YouTube videos and all of the Xigua videos, we train a classifier and predict the category tags according to the subtitles.

Hyperparameters	BIGVIDEO	HOW2	VATEX
GPUs	8	2	1
Batch Size	4,096	4,096	4,096
Dropout	0.1	0.3	0.3
Weight Decay	0.1	0.1	0.1
Learning Rate	7e-4	5e-4	1e-3
Warmup Steps	4000	4000	2000
Layer Normalization	PostNorm	PostNorm	PostNorm

Table 12: Training hyperparameters details.

C Experimental Detatils

C.1 Dataset

We additionally conduct experiments on two public video-guided translation datasets, VATEX (Wang et al., 2019) and HOW2 (Sanabria et al., 2018). The HOW2 dataset is a collection of instructional videos from Youtube. The corpus contains 184,948 English-Portuguese pairs for training, each associated with a video clip. We utilize val (2,022) as the validation set and dev5 (2,305) as the testing set. The VATEX dataset is a video-and-language dataset containing over 41,250 unique videos. The released version of the bilingual collection includes 129,955 sentence pairs for training, 15,000 sentence pairs for validation, and 30,000 for testing. Since the testing set is not publicly available, we split the original validation set into two halves for validation and testing. Some video clips of VATEX are no longer available on the Youtube. So after removal, the used corpus contains 115,480




Figure 7: Bleu scores on BIGVIDEO validation, test and AMBIGUOUS sets. The x-axis is the choice of the different temperatures for the contrastive learning objective.




Figure 8: Bleu scores on BIGVIDEO validation, test and AMBIGUOUS sets. The x-axis is the choice of different weights for the contrastive learning objective.

sentence pairs for training, 6,645 sentence pairs for validation, and 6,645 sentence pairs for testing, each associated with a video clip.

C.2 Training and Implementation Details

More training details can be found in Table 12. For the pretraining on WMT19 Zh-En dataset, we utilize the same training parameters as that on BIGVIDEO and train the model for 300k steps.

D The Choice of Hyper-parameters

Temperature for contrastive learning objective.

Performances of different temperature are presented in Figure 7. Here we fix the weight for contrastive learning objective to 1. On the validation set, there exists no significant difference in BLEU scores among choices of temperature. For better translation performance, a small temperature is more suitable.

Weight for contrastive learning objective. We fix the $\tau = 0.002$ and adjust the weight from 0.5 to 1.5. We can observe that contrastive learning objective with varying weights benefits the model to different degrees. 1.0 is the most suitable weight for our system.

Length of Video Frames. To investigate how the




Figure 9: Bleu scores on BIGVIDEO validation, test and AMBIGUOUS sets. The x-axis is the length of video frames for our system.



Source Subtitle:Number one **drive shot** requires smaller swing but more focus.

Target Subtitle:第一、**抽球**。挥杆幅度要小，但是要集中力量。

System w/o Video:第一个**驾驶镜头**需要较小的挥杆，但更多的焦点。


System w/ Video:第一，**抽球**需要更小的挥杆动作，但要集中注意力

Figure 10: A case. The phrases with semantic ambiguity are highlighted in **red**. The wrong translations are in **blue** and the correct translations are in **yellow**.

length of video frames affects translation, we adjust the number of sampled video frames in [1,12,36]. Figure 9 depicts their performances. Here the video features we use are 2D features extracted by ViT. We can observe that when only one video frame is sampled, the video degrades into one image and its positive impact on the system is reduced. A maximum of 12 video frames achieves the best performance.

E Case Study

We additionally present two cases in the appendix. In figure 10, the phrase “drive shot” is better translated by our system by understanding the meaning of “shot”. In Figure 11, we can find both the text-only baseline and our system fail to correctly translate the source title. The objects in the video are *cards of Duel Monsters*, which need world knowledge to understand. So the source title is complicated for text-only and our system.



Source Subtitle: I spent And there we go , a **Blackfeather Darkrage Dragon** , which is like **Black Winged Dragon**.

Ground Truth: 好了，**玄翼龙黑羽**，就像**黑翼龙**。

System w/o Video: 好了，**黑羽毛darkragedragon**就像**黑风龙**。

System w/ Video: 这是**黑羽毛的龙龙**，就像**黑窗龙**。

Figure 11: A case. The phrases with semantic ambiguity are highlighted in **red**. The wrong translations are in **blue** and the correct translations are in **yellow**.

F Annotation Guidelines

We hire seven full-time annotators who are fluent in both Chinese and English. They are recruited to annotate translation data or conduct human evaluations. The annotators are shown one English and corresponding Chinese subtitle of the given video clip. After watching videos and reading subtitles, they are required to decide whether videos are related to subtitles. If not, the sample will be discarded. Then the annotators are required to rate on three aspects:

- **Fluency Score (1-5, 1 is the worst and 5 is the best):** If the audio is in English, the annotators will need to check whether the English subtitle is the transcript of the audio. If the audio is not in English, the annotators will need to rate if the sentence is grammatically correct.
- **Translation Quality (1-5, 1 is the worst and 5 is the best):** Whether the Chinese subtitle is equivalent in meaning to the English Subtitle.
- **Ambiguous (0/1):** The annotators need to decide whether the video information is required to complete the translation. "1" means "the video information is required" and otherwise "0".

G Human Evaluation Guidelines

We hire three annotators to conduct the human evaluation. Each annotator is required to rate 100 samples from AMBIGUOUS and 100 samples from UNAMBIGUOUS on **translation quality** and rank

two systems. The definition of the **translation quality** is the same as that in annotation guidelines.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
We discuss the limitations of our work in Section Limitation.
- A2. Did you discuss any potential risks of your work?
We discuss the limitations of our work in Section Ethical Considerations.
- A3. Do the abstract and introduction summarize the paper’s main claims?
It can be seen in the abstract and the introduction of the paper.
- A4. Have you used AI writing assistants when working on this paper?
We use Grammarly to check our writing.

B Did you use or create scientific artifacts?

We use some publicly available scientific artifacts, which can be seen in Section Experiments and Appendix.

- B1. Did you cite the creators of artifacts you used?
We cite the creators of the datasets and code we use, which can be seen in Section Experiments and Appendix.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
We discuss this problem in Section Ethical Considerations.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
We have discussed these problems in the Section Ethical Considerations and Limitations. All the datasets and code repos we use are for research purposes and conform to the intention of the original creators. For the dataset we create, we specify intended use in the Limitations.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
We avoid collecting user information or other sensitive information. However, in our dataset, we recognize that there might exist videos that contain sensitive information, such as portraits. Upon publication, we plan to release video features and publicly available url for videos that may contain sensitive information.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
We provide information about the artifacts in Section Dataset and Appendix.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
We report relevant statistics in Section Dataset and Appendix.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

C **Did you run computational experiments?**

We provide information about computational experiments in Section Experiments.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

We provide the information in Section Experiments and Appendix.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

We provide the information in Section Experiments and Appendix.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

We provide the information in Section Results and Appendix. We report results which are the mean values of five runs and their standard deviations.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

We provide the information in Section Experiments.

D **Did you use human annotators (e.g., crowdworkers) or research with human participants?**

We provide the information in Appendix.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

We provide the information in Appendix.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

We provide the information in Section Limitations.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

We provide the information in Section Dataset and Appendix.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Our data collection protocol is followed under the instruction of professional attorneys.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

We report in Appendix.