

MTVR: Multilingual Moment Retrieval in Videos

Jie Lei Tamara L. Berg Mohit Bansal

Department of Computer Science

University of North Carolina at Chapel Hill

{jielei, tlberg, mbansal}@cs.unc.edu

Abstract

We introduce MTVR, a large-scale multilingual video moment retrieval dataset, containing 218K English and Chinese queries from 21.8K TV show video clips. The dataset is collected by extending the popular TVR dataset (in English) with paired Chinese queries and subtitles. Compared to existing moment retrieval datasets, MTVR is multilingual, larger, and comes with diverse annotations. We further propose mXML, a multilingual moment retrieval model that learns and operates on data from both languages, via encoder parameter sharing and language neighborhood constraints. We demonstrate the effectiveness of mXML on the newly collected MTVR dataset, where mXML outperforms strong monolingual baselines while using fewer parameters. In addition, we also provide detailed dataset analyses and model ablations. Data and code are publicly available at <https://github.com/jayleicn/mTVRretrieval>

1 Introduction

The number of videos available online is growing at an unprecedented speed. Recent work (Escorcia et al., 2019; Lei et al., 2020) introduced the Video Corpus Moment Retrieval (VCMR) task: given a natural language query, a system needs to retrieve a short moment from a large video corpus. Figure 1 shows a VCMR example. Compared to the standard text-to-video retrieval task (Xu et al., 2016; Yu et al., 2018), it allows more fine-grained moment-level retrieval, as it requires the system to not only retrieve the most relevant videos, but also localize the most relevant moments inside these videos. Various datasets (Krishna et al., 2017; Hendricks et al., 2017; Gao et al., 2017; Lei et al., 2020) have been proposed or adapted for the task. However, they are all created for a single language (English), though the application could be useful for users speaking other languages as well. Besides, it is also unclear

Video Corpus:



00:00,327 → 00:04,320 Whitney: This is my fiancé...
惠特尼: 这是我的未婚夫... 00:32,192 → 00:34,626 House: Nine months later, a miracle...
豪斯: 9个月之后, 一个奇迹...



00:07,786 → 00:13,156 Monica: Who wasn't invited...
莫妮卡: 还没有被邀请到... 00:44,223 → 00:52,929 Rachel: Daddy, I can't marry him...
瑞秋: 爸爸, 我不能嫁给他...



00:03,897 → 00:07,731 Ross: Somebody seems to be...
罗斯: 有人在... 00:36,497 → 00:38,761 Rachel: Call me when you get this...
瑞秋: 听到留言请回电。

Query:

Rachel explains to her dad on the phone why she can't marry her fiancé.
瑞秋在电话里向她父亲解释了她不能和其未婚夫结婚的原因。

Query Type: video + subtitle

Figure 1: A MTVR example in the Video Corpus Moment Retrieval (VCMR) task. Ground truth moment is shown in green box. Colors in the query text indicate whether the words are more related to video (orchid) or subtitle (salmon) or both (orange). The query and the subtitle text are presented in both English and Chinese. The video corpus typically contains thousands of videos, for brevity, we only show 3 videos here.

whether the progress and findings in one language generalizes to another language (Bender, 2009). While there are multiple existing multilingual image datasets (Gao et al., 2015; Elliott et al., 2016; Shimizu et al., 2018; Pappas et al., 2016; Lan et al., 2017; Li et al., 2019), the availability of multilingual video datasets (Wang et al., 2019a; Chen and Dolan, 2011) is still limited.

Therefore, we introduce MTVR, a large-scale, multilingual moment retrieval dataset, with 218K human-annotated natural language queries in two

languages, English and Chinese. MTVR extends the TVR (Lei et al., 2020) dataset by collecting paired Chinese queries and Chinese subtitle text (see Figure 1). We choose TVR over other moment retrieval datasets (Krishna et al., 2017; Hendricks et al., 2017; Gao et al., 2017) because TVR is the largest moment retrieval dataset, and also has the advantage of having dialogues (in the form of subtitle text) as additional context for retrieval, in contrast to pure video context in the other datasets. We further propose mXML, a compact, multilingual model that learns jointly from both English and Chinese data for moment retrieval. Specifically, on top of the state-of-the-art monolingual moment retrieval model XML (Lei et al., 2020), we enforce encoder parameter sharing (Sachan and Neubig, 2018; Dong et al., 2015) where the queries and subtitles from the two languages are encoded using shared encoders. We also incorporate a language neighborhood constraint (Wang et al., 2018; Kim et al., 2020) to the output query and subtitle embeddings. It encourages sentences of the same meaning in different languages to lie close to each other in the embedding space. Compared to separately trained monolingual models, mXML substantially reduces the total model size while improving retrieval performance (over monolingual models) as we show in Section 4. Detailed dataset analyses and model ablations are provided.

2 Dataset

The TVR (Lei et al., 2020) dataset contains 108,965 high-quality English queries from 21,793 videos from 6 long-running TV shows (provided by TVQA (Lei et al., 2018)). The videos are associated with English dialogues in the form of subtitle text. MTVR extends this dataset with translated dialogues and queries in Chinese to support multilingual multimodal research.

2.1 Data Collection

Dialogue Subtitles. We crawl fan translated Chinese subtitles from subtitle sites.¹ All subtitles are manually checked by the authors to ensure they are of good quality and are aligned with the videos. The original English subtitles come with speaker names from transcripts that we map to the Chinese subtitles, to ensure that the Chinese subtitles have the same amount of information as the English version.

¹<https://subhd.tv>, <http://zimuku.la>

QType (%)	Query Examples (in English and Chinese)
video-only (74.2)	Howard places his plate onto the coffee table. 霍华德将盘子放在咖啡桌子上。
sub-only (9.1)	Alexis and Castle talk about the timeline of the murder. 亚历克西斯和卡塞尔谈论谋杀的时间顺序。
video+sub (16.6)	Joey waives his hand when he asks for his food. 乔伊催餐时摆了摆手。

Table 1: MTVR English and Chinese query examples in different query types. The percentage of the queries in each query type is shown in brackets.

Query. To obtain Chinese queries, we hire human translators from Amazon Mechanical Turk (AMT). Each AMT worker is asked to write a Chinese translation of a given English query. Languages are ambiguous, hence we also present the original videos to the workers at the time of translation to help clarify query meaning via spatio-temporal visual grounding. The Chinese translations are required to have the exact same meaning as the original English queries and the translation should be made based on the aligned video content. To facilitate the translation process, we provide machine translated Chinese queries from Google Cloud Translation² as references, similar to (Wang et al., 2019b). To find qualified bilingual workers in AMT, we created a qualification test with 5 multiple-choice questions designed to evaluate workers’ Chinese language proficiency and their ability to perform our translation task. We only allow workers that correctly answer all 5 questions to participate our annotation task. In total, 99 workers finished the test and 44 passed, earning our qualification. To further ensure data quality, we also manually inspect the submitted results during the annotation process and disqualify workers with poor annotations. We pay workers \$0.24 every three sentences, this results in an average hourly pay of \$8.70. The whole annotation process took about 3 months and cost approximately \$12,000.00.

2.2 Data Analysis

In Table 2, we compare the average sentence lengths and the number of unique words under different part-of-speech (POS) tags, between the two languages, English and Chinese, and between query and subtitle text. For both languages, dialogue subtitles are linguistically more diverse than queries, i.e., they have more unique words in all

²<https://cloud.google.com/translate>

Data	Avg Len	#unique words by POS tags				
		all	verb	noun	adj.	adv.
English						
Q	13.45	15,201	3,015	7,143	2,290	763
Sub	10.78	49,325	6,441	19,223	7,504	1,740
Q+Sub	11.27	52,545	7,151	20,689	8,021	1,976
Chinese						
Q	12.55	34,752	12,773	18,706	1,415	1,669
Sub	9.04	101,018	36,810	53,736	4,958	5,568
Q+Sub	9.67	117,448	42,284	62,611	5,505	6,185

Table 2: Comparison of English and Chinese data in MTRV. We show average sentence length, and number of unique tokens by POS tags, for Query (Q) and or Subtitle (Sub).

categories. This is potentially because the language used in subtitles are unconstrained human dialogues while the queries are collected as declarative sentences referring to specific moments in videos (Lei et al., 2020). Comparing the two languages, the Chinese data is typically more diverse than the English data.³ In Table 1, we show English and their translated Chinese query examples in Table 1, by query type. In the appendix, we compare MTRV with existing video and language datasets.

3 Method

Our multilingual moment retrieval model mXML is built on top of the Cross-model Moment Localization (XML) (Lei et al., 2020) model, which performs efficient video-level retrieval at its shallow layers and accurate moment-level retrieval at its deep layers. To adapt the monolingual XML model into the multilingual setting in MTRV and improve its efficiency and effectiveness, we apply encoder parameter sharing and neighborhood constraints (Wang et al., 2018; Kim et al., 2020) which encourages the model to better utilize multilingual data to improve monolingual task performance while maintaining smaller model size.

Query and Context Representations. We represent videos using ResNet-152 (He et al., 2016) and I3D (Carreira and Zisserman, 2017) features extracted every 1.5 seconds. We extract language features using pre-trained, then finetuned (on our queries and subtitles) RoBERTa-base (Liu et al., 2019), for English (Liu et al., 2019) and Chinese (Cui et al., 2020), respectively. For queries, we use token-level features. For subtitles, we max-

³ The differences might be due to the different morphemes in the languages. E.g., the Chinese word 长发 (‘long hair’) is labeled as a single noun, but as an adjective (‘long’) and a noun (‘hair’) in English (Wang et al., 2019b).




Figure 2: Illustration of mXML’s encoding process. Compared to monolingual models, mXML learns from the two languages simultaneously, and allows them to benefit each other via *encoder parameter sharing* and *neighborhood constraints*. We show the detailed encoding process of the model in the appendix (Figure 3).

pool the token-level features every 1.5 seconds to align with the video features. We then project the extracted features into a low-dimensional space via a linear layer, and add learned positional encoding (Devlin et al., 2018) after the projection. We denote the resulting video features as $E^v \in \mathbb{R}^{l \times d}$, subtitle features as $E_{en}^s \in \mathbb{R}^{l \times d}$, $E_{zh}^s \in \mathbb{R}^{l \times d}$, and query features as $E_{en}^q \in \mathbb{R}^{l_q \times d}$, $E_{zh}^q \in \mathbb{R}^{l_q \times d}$. l is video length, l_q is query length, and d is hidden size. The subscripts en and zh denote English and Chinese text features, respectively.

Encoders and Parameter Sharing. We follow Lei et al. (2020) to use *Self-Encoder* as our main component for query and context encoding. A Self-Encoder consists of a self-attention (Vaswani et al., 2017) layer, a linear layer, and a residual (He et al., 2016) connection followed by layer normalization (Ba et al., 2016). We use a Self-Encoder followed by a modular attention (Lei et al., 2020) to encode each query into two modularized query vectors $q_{lang}^v, q_{lang}^s \in \mathbb{R}^d$ ($lang \in \{en, zh\}$) for video and subtitle retrieval, respectively. For videos, we apply two Self-Encoders instead of a Self-Encoder and a Cross-Encoder as in XML, because we found this modification simplifies the implementation while maintains the performance. We use the outputs from the first and the second Self-Encoder $H_{vr,lang}^v, H_{mr,lang}^v \in \mathbb{R}^{l \times d}$ for video retrieval and moment retrieval. Similarly, we have $H_{vr,lang}^s, H_{mr,lang}^s \in \mathbb{R}^{l \times d}$ for subtitles. All the Self-Encoders are shared across languages, e.g., we use the same Self-Encoder to encode both English and Chinese queries, as illustrated in Figure 2. This parameter sharing strategy greatly reduces the

model size while maintaining or even improving model performance, as we show in Section 4.

Language Neighborhood Constraint. To facilitate stronger multilingual learning, we add neighborhood constraints (Wang et al., 2018; Kim et al., 2020; Burns et al., 2020) to the model. This encourages sentences that express the same or similar meanings to lie close to each other in the embedding space, via a triplet loss. Given paired sentence embeddings $e_{en}^i \in \mathbb{R}^d$ and $e_{zh}^i \in \mathbb{R}^d$, we sample negative sentence embeddings $e_{en}^j \in \mathbb{R}^d$ and $e_{zh}^k \in \mathbb{R}^d$ from the same mini-batch, where $i \neq j, i \neq k$. We use cosine similarity function \mathcal{S} to measure the similarity between embeddings. Our language neighborhood constraint can be formulated as:

$$\mathcal{L}_{nc} = \frac{1}{n} \sum_i [\max(0, \mathcal{S}(e_{en}^i, e_{zh}^k) - \mathcal{S}(e_{en}^i, e_{zh}^i) + \Delta) + \max(0, \mathcal{S}(e_{en}^j, e_{zh}^i) - \mathcal{S}(e_{en}^j, e_{zh}^j) + \Delta)], \quad (1)$$

where $\Delta=0.2$ is the margin. We apply this constraint on both query and subtitle embeddings, across the two languages, as illustrated in Figure 2. For queries, we directly apply it on the query vectors q_{lang}^v, q_{lang}^s . For the subtitle embeddings, we apply it on the embeddings $H_{vr,lang}^s, H_{mr,lang}^s$, after max-pooling them in the temporal dimension.

Training and Inference. During training, we optimize video retrieval scores with a triplet loss, and moment scores with a cross-entropy loss. At inference, these two scores are aggregated together as the final score for video corpus moment retrieval. See appendix for details.

4 Experiments and Results

We evaluate our proposed mXML model on the newly collected MTRV dataset, and compare it with several existing monolingual baselines. We also provide ablation studies evaluating our model design and the importance of each input modality (videos and subtitles).

Data Splits and Evaluation Metrics. We follow TVR (Lei et al., 2020) to split the data into 80% train, 10% val, 5% test-public and 5% test-private. We report average recall (R@1) on the Video Corpus Moment Retrieval (VCMR) task. A predicted moment is correct if it has high Intersection-over-Union (IoU) with the ground-truth.

Baseline Comparison. In Table 3, we compare mXML with multiple baseline approaches. Given

Method	#param	English R@1		Chinese R@1	
		IoU=0.5	IoU=0.7	IoU=0.5	IoU=0.7
Chance	-	0.00	0.00	0.00	0.00
Proposal based					
MCN	6.4M	0.02	0.00	0.13	0.02
CAL	6.4M	0.09	0.04	0.11	0.04
Retrieval + Re-ranking					
MEE+MCN	10.4M	0.92	0.42	1.43	0.64
MEE+CAL	10.4M	0.97	0.39	1.51	0.62
MEE+ExCL	10.0M	0.92	0.33	1.43	0.72
XML	6.4M	7.25	3.25	5.91	2.57
mXML	4.5M	8.30	3.82	6.76	3.20

Table 3: Baseline comparison on MTRV *test-public* split. mXML achieves better retrieval performance on both languages while using fewer parameters.

a natural language query, the goal of video corpus moment retrieval is to retrieve relevant moments from a large video corpus. The methods for this task can be grouped into two categories, (i) proposal based approaches (MCN (Hendricks et al., 2017) and CAL (Escorcia et al., 2019)), where they perform video retrieval on the pre-segmented moments from the videos; (ii) retrieval+re-ranking methods (MEE (Miech et al., 2018)+MCN, MEE+CAL, MEE+ExCL (Ghosh et al., 2019) and XML (Lei et al., 2020)), where one approach is first used to retrieve a set of videos, then another approach is used to re-rank the moments inside these retrieved videos to get the final moments. Our method mXML also belongs to the retrieval+re-ranking category. Across all metrics and both languages, we notice retrieval+re-ranking approaches achieve better performance than proposal based approaches, indicating that retrieval+re-ranking is potentially better suited for the VCMR task. Meanwhile, our mXML outperforms the strong baseline XML significantly⁴ while using few parameters. XML is a monolingual model, where a separate model is trained for each language. In contrast, mXML is multilingual, trained on both languages simultaneously, with parameter sharing and language neighborhood constraints to encourage multilingual learning. mXML prediction examples are provided in the appendix.

Ablations on Model Design. In Table 4, we present our ablation study on mXML. We use ‘Baseline’ to denote the mXML model without parameter sharing and neighborhood constraint. Shar-

⁴Statistically significant with $p < 0.01$. We use bootstrap test (Efron and Tibshirani, 1994; Noreen, 1989).

Method	#param	English R@1		Chinese R@1	
		IoU=0.5	IoU=0.7	IoU=0.5	IoU=0.7
Baseline	6.4M	5.77	2.63	4.7	2.38
+ Share Enc.	4.5M	6.09	2.85	4.72	2.25
+ NC (mXML)	4.5M	6.22	2.96	5.17	2.41

Table 4: mXML ablation study on MTRV *val* split. *Share Enc.* = encoder parameter sharing, *NC* = Neighborhood Constraint. Each row adds an extra component to the row above it.

Model Type	English R@1		Chinese R@1	
	IoU=0.5	IoU=0.7	IoU=0.5	IoU=0.7
Query type: video				
Baseline	5.46	2.53	4.78	2.47
mXML	5.77	2.67	5.14	2.32
Query type: subtitle				
Baseline	4.15	1.97	3.11	1.14
mXML	6.12	3.32	4.05	1.87
Query type: video+subtitle				
baseline	8.02	3.38	5.18	2.62
mXML	8.29	4.09	5.89	3.11

Table 5: Comparison of mXML and the baseline on MTRV *val* set, with breakdown on query types. Both models are trained with video and subtitle as inputs.

ing encoder parameter across languages greatly reduces #parameters while maintaining (Chinese) or even improving (English) model performance. Adding neighborhood constraint does not introduce any new parameters but brings a notable ($p < 0.06$) performance gain to both languages. We hypothesize that this is because the learned information in the embeddings of the two languages are complementary (though the sentences in the two languages express the same meaning, their language encoders (Liu et al., 2019; Cui et al., 2020)) are pre-trained differently, which may lead to different meanings at the embedding level. In Table 5, we show a detailed comparison between mXML and its baseline version, by query types. Overall, we notice the mXML perform similarly with *Baseline* in ‘*video*’ queries, but shows a significant performance gain in ‘*subtitle*’ queries, suggesting the parameter sharing and neighborhood constraint are more useful for queries that need more language understanding.

Ablations on Input Modalities. In Table 6, we compare mXML variants with different context inputs, i.e., video or subtitle or both. We report their performance under the three annotated query types,

QType (percentage)	English R@1		Chinese R@1	
	IoU=0.5	IoU=0.7	IoU=0.5	IoU=0.7
Model input: video				
video (74.32%)	4.12	1.89	3.73	1.86
sub (8.85%)	1.97	1.24	1.35	1.04
video+sub (16.83%)	2.67	1.2	2.45	1.15
Model input: subtitle				
video	1.35	0.62	1.11	0.51
sub	6.33	2.9	4.15	1.97
video+sub	6.22	2.62	4.2	2.13
Model input: video+subtitle				
video	5.77	2.67	5.14	2.32
sub	6.12	3.32	4.05	1.87
video+sub	8.29	4.09	5.89	3.11

Table 6: mXML performance breakdown on MTRV *val* set by query types, with different inputs.

video, *sub* and *video+sub*. Overall, the model with both video and subtitle as inputs perform the best. The video model performs much better on the *video* queries than on the *sub* queries, while the subtitle model achieves higher scores on the *sub* queries than the *video* queries.

In the appendix, we also present results on ‘*generalization to unseen TV shows*’ setup.

5 Conclusion

In this work, we collect MTRV, a new large-scale, multilingual moment retrieval dataset. It contains 218K queries in English and in Chinese from 21.8K video clips from 6 TV shows. We also propose a multilingual moment retrieval model mXML as a strong baseline for the MTRV dataset. We show in experiments that mXML outperforms monolingual models while using fewer parameters.

Acknowledgements

We thank the reviewers for their helpful feedback. This research is supported by NSF Award #1562098, DARPA KAIROS Grant #FA8750-19-2-1004, and ARO-YIP Award #W911NF-18-1-0336. The views contained in this article are those of the authors and not of the funding agency.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Emily M Bender. 2009. Linguistically naïve!= language independent: Why nlp needs linguistic typology. In *Proceedings of the EACL 2009 Workshop*

- on the Interaction between Linguistics and Computational Linguistics: Virtuous, Vicious or Vacuous?*
- Andrea Burns, Donghyun Kim, Derry Wijaya, Kate Saenko, and Bryan A Plummer. 2020. Learning to scale multilingual representations for vision-language tasks. In *ECCV*.
- Joao Carreira and Andrew Zisserman. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*.
- David Chen and William B Dolan. 2011. Collecting highly parallel data for paraphrase evaluation. In *ACL*.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Shijin Wang, and Guoping Hu. 2020. Revisiting pre-trained models for chinese natural language processing. In *EMNLP*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. Multi-task learning for multiple language translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*.
- Bradley Efron and Robert J Tibshirani. 1994. *An introduction to the bootstrap*. CRC press.
- Desmond Elliott, Stella Frank, Khalil Sima'an, and Lucia Specia. 2016. **Multi30K: Multilingual English-German image descriptions**. In *Proceedings of the 5th Workshop on Vision and Language*, pages 70–74, Berlin, Germany. Association for Computational Linguistics.
- Victor Escorcia, Mattia Soldan, Josef Sivic, Bernard Ghanem, and Bryan Russell. 2019. Temporal localization of moments in video collections with natural language. *arXiv preprint arXiv:1907.12763*.
- Haoyuan Gao, Junhua Mao, Jie Zhou, Zhiheng Huang, Lei Wang, and Wei Xu. 2015. Are you talking to a machine? dataset and methods for multilingual image question. In *NIPS*.
- Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. 2017. Tall: Temporal activity localization via language query. In *ICCV*.
- Soham Ghosh, Anuva Agarwal, Zarana Parekh, and Alexander Hauptmann. 2019. Excl: Extractive clip localization using natural language descriptions. In *NAACL*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*.
- Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. 2017. Localizing moments in video with natural language. In *ICCV*.
- Donghyun Kim, Kuniaki Saito, Kate Saenko, Stan Sclaroff, and Bryan A Plummer. 2020. Mule: Multimodal universal language embedding. In *AAAI*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. Dense-captioning events in videos. In *ICCV*.
- Weiyu Lan, Xirong Li, and Jianfeng Dong. 2017. Fluency-guided cross-lingual image captioning. In *ACM MM*.
- Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L. Berg. 2018. Tvqa: Localized, compositional video question answering. In *EMNLP*.
- Jie Lei, Licheng Yu, Tamara L Berg, and Mohit Bansal. 2020. Tvr: A large-scale dataset for video-subtitle moment retrieval. In *ECCV*.
- Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. 2020. Hero: Hierarchical encoder for video+ language omni-representation pre-training. In *EMNLP*.
- Xirong Li, Chaoxi Xu, Xiaoxu Wang, Weiyu Lan, Zhengxiong Jia, Gang Yang, and Jieping Xu. 2019. Coco-cn for cross-lingual image tagging, captioning, and retrieval. *TMM*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Antoine Miech, Ivan Laptev, and Josef Sivic. 2018. Learning a text-video embedding from incomplete and heterogeneous data. *arXiv preprint arXiv:1804.02516*.
- Eric W Noreen. 1989. *Computer-intensive methods for testing hypotheses*. Wiley New York.
- Nikolaos Pappas, Miriam Redi, Mercan Topkara, Brendan Jou, Hongyi Liu, Tao Chen, and Shih-Fu Chang. 2016. Multilingual visual sentiment concept matching. In *ICMR*.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. In *NeurIPS Autodiff Workshop*.
- Michaela Regneri, Marcus Rohrbach, Dominikus Wetzel, Stefan Thater, Bernt Schiele, and Manfred Pinkal. 2013. Grounding action descriptions in videos. *TACL*.

Devendra Sachan and Graham Neubig. 2018. Parameter sharing methods for multilingual self-attentional translation models. In *Proceedings of the Third Conference on Machine Translation: Research Papers*.

Nobuyuki Shimizu, Na Rong, and Takashi Miyazaki. 2018. Visual question answering dataset for bilingual image understanding: A study of cross-lingual transfer using attention maps. In *COLING*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*.

Liwei Wang, Yin Li, Jing Huang, and Svetlana Lazebnik. 2018. Learning two-branch neural networks for image-text matching tasks. *TPAMI*.

Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuanfang Wang, and William Yang Wang. 2019a. Vatex: A large-scale, high-quality multilingual dataset for video-and-language research. In *ICCV*.

Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuanfang Wang, and William Yang Wang. 2019b. Vatex: A large-scale, high-quality multilingual dataset for video-and-language research. In *ICCV*.

Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*.

Youngjae Yu, Jongseok Kim, and Gunhee Kim. 2018. A joint sequence fusion model for video question answering and retrieval. In *ECCV*.

A Appendix

Data Analysis. In Table 8 we show a comparison of MTVR with existing moment retrieval datasets and related video and language datasets. Compared to other moment retrieval datasets, MTVR is significantly larger in scale, and comes with query type annotations that allows in-depth analyses for the models trained on it. Besides, it is also the only moment retrieval dataset with multilingual annotations, which is vital in studying the moment retrieval problem under the multilingual context. Compared to the existing multilingual video and language datasets, MTVR is unique as it has a more diverse set of context and annotations, i.e., dialogue, query type, and timestamps.

Training and Inference Details. In Figure 3 we show an overview of the mXML model. We compute video retrieval score as:

$$s^{vr} = \frac{1}{2} \sum_{m \in \{v, s\}} \max\left(\frac{H_{vr}^m}{\|H_{vr}^m\|} \frac{\mathbf{q}^m}{\|\mathbf{q}^m\|}\right). \quad (2)$$

Setting	English R@1		Chinese R@1	
	IoU=0.5	IoU=0.7	IoU=0.5	IoU=0.7
unseen	1.68	0.79	1	0.54
seen	4.82	2.79	4.18	2.32

Table 7: mXML performance on the MTVR val split *Friends* examples, in both *unseen* and *seen* settings.

The subscript $lang \in \{en, zh\}$ is omitted for simplicity. It is optimized using a triplet loss similar to main text Equation (1). For moment retrieval, we first compute the query-clip similarity scores $S^{q,c} \in \mathbb{R}^l$ as:

$$S^{q,c} = \frac{1}{2}(H_{mr}^s \mathbf{q}^s + H_{mr}^v \mathbf{q}^v). \quad (3)$$

Next, we apply Convolutional Start-End Detector (ConvSE module) (Lei et al., 2020) to obtain start, end probabilities $P_{st}, P_{ed} \in \mathbb{R}^l$. These scores are optimized using a cross-entropy loss. The single video moment retrieval score for moment $[t_{st}, t_{ed}]$ is computed as:

$$s^{mr}(t_{st}, t_{ed}) = P_{st}(t_{st})P_{ed}(t_{ed}), t_{st} \leq t_{ed}. \quad (4)$$

Given a query q_i , the retrieval score for moment $[t_{st}:t_{ed}]$ in video v_j is computed following the aggregation function as in (Lei et al., 2020):

$$s^{vcmr}(v_j, t_{st}, t_{ed}|q_i) = s^{mr}(t_{st}, t_{ed})\exp(\alpha s^{vr}(v_j|q_i)), \quad (5)$$

where $\alpha=20$ is used to assign higher weight to the video retrieval scores. The overall loss is a simple summation of video and moment retrieval loss across the two languages, and the language neighborhood constraint loss.

Implementation Details. mXML is implemented in PyTorch (Paszke et al., 2017). We use Adam (Kingma and Ba, 2014) with initial learning rate $1e-4$, $\beta_1=0.9$, $\beta_2=0.999$, L2 weight decay 0.01, learning rate warm-up over the first 5 epochs. We train mXML for at most 100 epochs at batch size 128, with early stop based on the sum of R@1 (IoU=0.7) scores for English and Chinese. The experiments are conducted on a NVIDIA RTX 2080Ti GPU. Each run takes around 7 hours.

Generalization to Unseen TV shows. To investigate whether the learned model can be transferred to other TV shows, we conduct an experiment by using the TV show ‘*Friends*’ as an ‘*unseen*’ TV

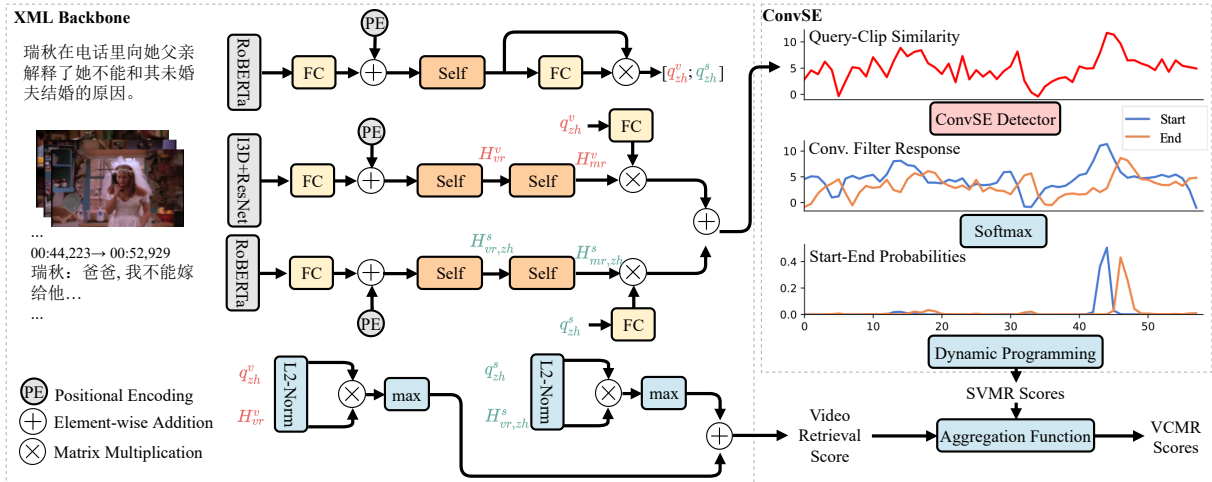


Figure 3: mXML overview. For brevity, we only show the modeling process for a single language (Chinese). The cross-language modifications, i.e., parameter sharing and neighborhood constraint are illustrated in Figure 2. This figure is edited from the Figure 4 in (Lei et al., 2020).

Dataset	Domain	#Q/#videos	Multilingual	Dialogue	QType	Timestamp
QA datasets with temporal annotation						
TVQA (Lei et al., 2018)	TV show	152.5K/21.8K	-	✓	-	✓
How2QA (Li et al., 2020)	Instructional	44K/22K	-	✓	-	✓
Multilingual video description datasets						
MSVD (Chen and Dolan, 2011)	Open	70K/2K	✓	-	-	-
VATEX (Wang et al., 2019b)	Activity	826K/41.3K	✓	-	-	-
Moment retrieval datasets						
TACoS (Regneri et al., 2013)	Cooking	16.2K/0.1K	-	-	-	✓
DiDeMo (Hendricks et al., 2017)	Flickr	41.2K/10.6K	-	-	-	✓
ActivityNet Captions (Krishna et al., 2017)	Activity	72K/15K	-	-	-	✓
CharadesSTA (Gao et al., 2017)	Activity	16.1K/6.7K	-	-	-	✓
How2R (Li et al., 2020)	Instructional	51K/24K	-	✓	-	✓
TVR (Lei et al., 2020)	TV show	109K/21.8K	-	✓	✓	✓
MTVR	TV show	218K/21.8K	✓	✓	✓	✓

Table 8: Comparison of MTVR with related video and language datasets.

show for testing, and train the model on all the other 5 TV shows. For comparison, we also include a model trained on ‘seen’ setting, where we use all the 6 TV shows including *Friends* for training. To ensure the models on these two settings are trained on the same number of examples, we downsample the examples in the *seen* setting to match the *unseen* setting. The results are shown in Table 7. We notice our mXML achieves a reasonable performance even though it does see a single example from the TV show *Friends*. Meanwhile, the gap between *unseen* and *seen* settings are still large, we encourage future work to further explore this direction.

Prediction Examples We show mXML prediction examples in Figure 4. We show both Chinese (*top*) and English (*bottom*) prediction examples, and correct (*left*) and incorrect (*right*) examples.



Figure 4: Qualitative examples of mXML. *Top*: examples in Chinese. *Bottom*: examples in English. *Left*: correct predictions. *Right*: incorrect predictions. We show top-3 retrieved moments for each query. salmon bar shows the predictions, green box indicates the ground truth.