

## Graphical Abstract (Optional)

### The MSR-Video to Text Dataset with Clean Annotations

Haoran Chen, Jianmin Li, Simone Frintrop, Xiaolin Hu



# ELSEVIER

Video captioning automatically generates short descriptions of the video content, usually in form of a single sentence. Many methods have been proposed for solving this task. A large dataset called MSR Video to Text (MSR-VTT) is often used as the benchmark dataset for testing the performance of the methods. However, we found that the human annotations, i.e., the descriptions of video contents in the dataset are quite noisy, e.g., there are many duplicate captions and many captions contain grammatical problems. These problems may pose difficulties to video captioning models for learning underlying patterns. We cleaned the MSR-VTT annotations by removing these problems, then tested several typical video captioning models on the cleaned dataset. Experimental results showed that data cleaning boosted the performances of the models measured by popular quantitative metrics. We recruited subjects to evaluate the results of a model trained on the original and cleaned datasets. The human behavior experiment demonstrated that trained on the cleaned dataset, the model generated captions that were more coherent and more relevant to the contents of the video clips.

### **Research Highlights (Required)**

To create your highlights, please type the highlights against each `\item` command.

It should be short collection of bullet points that convey the core findings of the article. It should include 3 to 5 bullet points (maximum 85 characters, including spaces, per bullet point.)

- Identify duplicate captions and grammatical problems in the MSR-VTT dataset.
- Clean the dataset and compare model performance before and after data cleaning.
- Inspect the impact of each step in data cleaning.
- Human evaluation shows the positive impact of data cleaning.



## The MSR-Video to Text Dataset with Clean Annotations

Haoran Chen<sup>a,b,c</sup>, Jianmin Li<sup>a,b,c</sup>, Simone Frintrop<sup>e</sup>, Xiaolin Hu<sup>a,b,c,d,\*\*</sup>

<sup>a</sup>State Key Laboratory of Intelligent Technology and Systems, Department of Computer Science and Technology, Tsinghua University, Beijing, 100084, China

<sup>b</sup>Beijing National Research Center for Information Science and Technology, Tsinghua University, Beijing, 100084, China

<sup>c</sup>Institute for Artificial Intelligence, Tsinghua University, Beijing, 100084, China

<sup>d</sup>Chinese Institute for Brain Research (CIBR), Beijing, 100010, China

<sup>e</sup>Department of Informatics, University of Hamburg, Hamburg, 20148, Germany

### ABSTRACT

Video captioning automatically generates short descriptions of the video content, usually in form of a single sentence. Many methods have been proposed for solving this task. A large dataset called MSR Video to Text (MSR-VTT) is often used as the benchmark dataset for testing the performance of the methods. However, we found that the human annotations, i.e., the descriptions of video contents in the dataset are quite noisy, e.g., there are many duplicate captions and many captions contain grammatical problems. These problems may pose difficulties to video captioning models for learning underlying patterns. We cleaned the MSR-VTT annotations by removing these problems, then tested several typical video captioning models on the cleaned dataset. Experimental results showed that data cleaning boosted the performances of the models measured by popular quantitative metrics. We recruited subjects to evaluate the results of a model trained on the original and cleaned datasets. The human behavior experiment demonstrated that trained on the cleaned dataset, the model generated captions that were more coherent and more relevant to the contents of the video clips.

© 2022 Elsevier Ltd. All rights reserved.

### 1. Introduction

The goal of video captioning is to summarize the content of a video clip by a single sentence, which is an extension of image captioning (Cho et al., 2014; Rennie et al., 2017; Yu et al., 2018; Anderson et al., 2018). To accomplish it, one must use both computer vision (CV) techniques and natural language processing (NLP) techniques. A benchmark dataset, called MSR-Video to Text 1.0 (MSR-VTT v1) (Xu et al., 2016b), was released in 2016. It contains 10,000 video clips and each clip is described by 20 captions, which are supposed to be different, given by human annotators. The dataset has become popular in the field of both video captioning and retrieval. Until March 31st, 2022, that work (Xu et al., 2016b) has been cited by 793 times according to Google scholar.

However, with a quick look, one can find many duplicate annotations, spelling mistakes and syntax errors in the annotations (Figs. 1, 2). It is unknown how many mistakes there are exactly

in the dataset and whether/how these mistakes would influence the performance of the video captioning models.

We quantitatively analyzed the annotations in the MSR-VTT dataset, and identified four main types of problems. First, thousands of annotations have duplicates for some of the video clips in the dataset. Second, thousands of special characters, such as "+", "-", ".", "/", ":", exist in the annotations. Third, thousands of spelling mistakes exist in the annotations. Fourth, hundreds of sentences are redundant or incomplete. We developed some techniques for cleaning the annotations to solve these problems. Our experiments demonstrated that existing models of video captioning, trained on the cleaned training set, had better performances compared to the results obtained by the models trained on the original training set. A human evaluation study also showed that a state-of-the-art model trained on the cleaned training set generated better captions than trained on the original training set in terms of semantic relevance and sentence coherence.

The cleaned dataset is available to the public<sup>1</sup>.

<sup>\*\*</sup>Corresponding author:

*e-mail:* [x1hu@tsinghua.edu.cn](mailto:x1hu@tsinghua.edu.cn) (Xiaolin Hu)

<sup>1</sup>GitHub repository: <https://github.com/WingsBrokenAngel/>



1. A man is throwing a football at a target.  $\times 2$
2. A man throws an American football at an aiming board.  $\times 3$
3. Kids throws football at target.  $\times 4$
4. Man throwing football to target in slow motion.  $\times 2$
5. People are playing sports.  $\times 2$
6. Someone is throwing a football at a target.  $\times 2$

**Figure 1.** An example video clip (No. 4290, starting from 0) with duplicate annotations.  $\times t$  denotes repeating  $t$  times

## 2. Related Work

### 2.1. Datasets

Three datasets MSVD (also called YouTube2Text), MSR-VTT and VATEX, unlimited to a specific domain, are widely used in recent studies of video captioning as benchmarks and video retrieval as well.

MSVD was published in 2013 (Guadarrama et al., 2013). It contains 1970 video clips and roughly 80,000 captions. Each video clip pairs with 40 captions.

MSR-VTT v1 was published in 2016 (Xu et al., 2016b). It contains 10,000 video clips and 200,000 captions. Each video clip pairs with 20 captions. The MSR-VTT v2 dataset was proposed in the second Video Captioning Competition<sup>2</sup> using the MSR-VTT v1 dataset as the training and validation sets and additional 3000 video clips as the test set. However, the annotations of the test set are not open to the public.

In 2019, a new large-scale video description dataset, named VATEX was presented, which is multilingual, linguistically complex and diverse in terms of video and annotations. It contains over 41,000 videos, reused from Kinetics-600, with 10 English text sentences for each of them (Wang et al., 2019c).

### 2.2. Recent Advances in Video Captioning

Kinds of models or methods or algorithms have been proposed for video captioning. With semantic concepts detected from the video, the probability distribution of each tag is integrated into the parameters of a recurrent unit in SCN (Gan et al., 2017). Video captioning is improved by sharing knowledge with two related tasks on the encoder and the decoder of a sequence-to-sequence model (Pasunuru and Bansal, 2017a). Reinforced learning is enhanced for video captioning with the mixed-loss function and the CIDEr-entailment reward in

CIDEnt-RL (Pasunuru and Bansal, 2017b). Multiple modalities are fused by hierarchical attention, which helps to improve the model performance, in HATT (Wu et al., 2018). The video feature produced by Efficient Convolutional Network is fed into a video captioning model, which boosts the quality of the generated caption, in the model named ECO (Zolfaghari et al., 2018). In the GRU-EVE, the Short Fourier Transform is applied to video features and high level semantics is derived from the object detector in order to generate captions rich in semantics (Wang et al., 2019a). A memory structure is used to capture the comprehensive visual information across the whole training set for a word in the MARN (Pei et al., 2019). The encoder employs a sibling (dual-branch) architecture to encode video clips in the SibNet (Liu et al., 2018). HACA fuses both global and local temporal dynamics existing in a video clip and generates an accurate description with knowledge from different modalities (Wang et al., 2018). Different expert modules are trained to provide knowledge for describing out-of-domain video clips in the TAMoE (Wang et al., 2019d). The model called SAM-SS is trained under the self-teaching manner to reduce the gap between the training and the test phase with meaningful semantic features (Chen et al., 2020b). Different types of representations are encoded and fused by the cross-gating block and captions are generated with Part-of-Speech information in the POS-RL (Wang et al., 2019b). In the VNS-GRU, “absolute equalitarianism” in the training process is alleviated by professional learning while a comprehensive selection method is used to choose the best checkpoint for the final test (Chen et al., 2020a). A new paradigm, named Open-book Video Captioning (Zhang et al., 2021), is adopted to generate natural language under the prompts of video-content-relevant sentences, unlimited to the video itself.

## 3. Analysis and Cleaning of the MSR-VTT dataset

Since MSR-VTT v2 uses MSR-VTT v1 for training and validation, and the annotations of the test set of MSR-VTT v2 are not open to the public, we performed analysis on MSR-VTT v1.

The MSR-VTT v1 dataset contains 10,000 video clips. Its training set has 6,513 video clips, the validation set has 497 video clips and the test set has 2,990 video clips. All clips are categorized into 20 classes with diverse contents and sceneries. A total of 0.2 million human annotations were collected to describe those video clips. The training/validation/test sets have 130,260/9,940/59,800 annotations, respectively. The vocabulary sizes of the training/validation/test set are 23,666/5,993/16,001, respectively.

### 3.1. Special Characters

There are 60 different characters in the dataset, including 0-9, a-z and 24 special characters in Table 1 (space is neglected). Generally speaking, those special characters are not used to train a model. We are intended to remove special characters while preserve information integrity in annotations.

We processed those special characters as follow:

MSR-VTT-DataCleaning

<sup>2</sup>Competition Website: <http://ms-multimedia-challenge.com/2017/challenge>



Figure 2. Three examples in the MSR-VTT dataset. The words in blue and red denote grammatical mistakes and spelling mistakes, respectively

1. Some special characters were removed from the sentences, include "#", "\*", "+", ".", ":", "=", ">", "[", "]", "(", ")", "\", where "[", "]", "(", and ")" were removed only when they were not in pairs.
2. The contents between bracket pairs "()" and "[]" were removed.
3. Special characters "-", "|", "'", "@", "\_", "''", "/" were replaced with spaces.
4. Characters from another language were replaced by the most similar English characters. For example, "é" was replaced by "e" in "error" and "B" by "b" in "beautiful".
5. "&" between two different words was substituted by "and".

In total, 7,248 out of 200,000 sentences and 4,190 out of 10,000 video clips were corrected (Table 3).

Table 1. Special characters in the MSR-VTT dataset

#	\$	%	&	(	)
*	+	-	.	—	:
=	>	@	[	\	]
/	‘		é	B	’

### 3.2. Spelling Mistakes

Massive spelling mistakes were found in the annotations during manual check. Tokenization is a process of demarcating a string of an input sentence into a list of words. After tokenization on each of the sentences, we used a popular spelling check software Hunspell<sup>3</sup> to check spelling errors.

Before Hunspell was used to do spelling checks, 784 new words were added to its vocabulary. These words were chosen manually by four criteria:

1. word abbreviations that are popular, eg. F1, WWF, RPG;
2. specific terms that are widely used, eg. Minecraft, Spongebob, Legos;
3. new words that are popular on the Internet, eg. gameplay, spiderman, talkshow;
4. names of persons, eg. Mariah, Fallon, Avril.

After that, spelling mistakes were found in 19,038 annotations out of 200,000 annotations. 21,826 words might have spelling mistakes suggested by Hunspell. Those candidates were corrected in the following steps:

1. Substituted British English spellings with the corresponding American English spellings. For instance, colour → color, travelling → traveling, programme → program, practising → practicing, theatre → theater. There were 61 such pairs.
2. Split unusual words that were created by concatenating two different words, e.g. rockclimbing → rock climbing, blowdrying → blow drying, swordfighting → sword fighting, screencaster → screen caster, rollercoaster → roller coaster. In total, 34 distinct words were found.
3. Corrected words that truly contain spelling mistakes, e.g., discusing → discussing, explaining → explaining, coversation → conversation, vedio → video, diffrent → different.

In total 35,668 words were substituted, split or corrected in these three steps for 27,954 sentences in 7,829 video clips as shown in Table 3.

### 3.3. Duplicate Annotations

Duplicate sentences were discovered in many annotations of video clips (Fig. 1). For each video clip, duplicates were removed. The similarity between two sentences was defined as follow

$$s_{a,b} = 0.5 \left( \frac{\mu(\mathbf{a}, \mathbf{b})}{\iota(\mathbf{a})} + \frac{\mu(\mathbf{a}, \mathbf{b})}{\iota(\mathbf{b})} \right), \quad (1)$$

<sup>3</sup>Available at <https://hunspell.github.io>

where  $\iota(\mathbf{x})$  denotes the word count in the sentence  $\mathbf{x} = \{x_1, x_2, \dots\}$  and  $\mu(\mathbf{a}, \mathbf{b})$  denotes the word count of the longest common subsequence in  $\mathbf{a}$  and  $\mathbf{b}$ .  $\mu(\mathbf{a}, \mathbf{b})$  is defined as follows,

$$\begin{aligned} \mu(\mathbf{a}, \mathbf{b}) &= \max(\iota(\mathbf{c})) \\ \text{s.t. } &\mathbf{c} \in \mathbf{a}, \mathbf{c} \in \mathbf{b}, \end{aligned} \quad (2)$$

where  $\mathbf{c} \in \mathbf{a}$  stands for that  $\mathbf{c}$  is a subsequence of  $\mathbf{a}$ . Word  $w_1$  and word  $w_2$  were regarded as the same if the Levenshtein distance (Levenshtein, 1966) between them was less than or equal to  $\bar{\epsilon}$ . Two sentences were regarded as duplicated if  $s_{a,b} > \bar{s}$ , where  $\bar{s}$  is the similarity threshold. With proper values of  $\bar{\epsilon}$  and  $\bar{s}$ , we could find duplicated sentences that had little difference. For example, considering the second pair of sentences in Table 2, the character ‘‘m’’ is missing in the word ‘‘woan’’ and the second sentence just has one more word ‘‘young’’ than the first sentence. These two sentences are almost the same in terms of meaning.

After duplicate removal, 183,856 video annotations remained in the dataset with 119,625 in the training set, 9,126 in the validation set and 55,105 in the test set with the hyper-parameters  $\hat{\epsilon} = 0$ ,  $\hat{s} = 0.85$ , tuned in Section 4. In another word, 17,733 sentences were removed in 7,129 video clips (Table 3). Each clip has 9 annotations at least, 20 at most and 18.4 on average.

### 3.4. Successive Sentences without Punctuations

In the task of video captioning, we expect each annotation contains one sentence. For many annotations in the dataset, each of them consists of multiple sentences. In Fig. 3, the first annotation can be split into three complete sentences: ‘‘A women in a dress talks about data scientist.’’, ‘‘She tells how they are problem solvers and well educated.’’, ‘‘She starts asking how you can stand out among other data scientist.’’ It causes two potential problems. First, the models trained on such annotations may output grammatically problematic sentences because these annotations are syntactically incorrect. Second, such annotations in the test set are no longer reliable ground truth so that the metrics, computed with them, are not reliable, neither.

However, many annotations in the dataset consist of multiple sentences. To solve it, one needs to manually separate the annotations into several complete sentences and merge them into a single sentence. For the sake of efficiency, the annotations, which consist of several sentences, were only divided and merged for the test set. For the training and validation sets, the sentences longer than  $l_a + 2\sigma$ , where  $l_a$  and  $\sigma$  denote the average sentence length and its standard deviation, were truncated respectively. After the process on it, 5,543 sentences were corrected in 2,758 video clips (Table 3).

## 4. Experiments

Experiments were conducted on the original and cleaned MSR-VTT datasets with several existing video captioning models, SCN (Gan et al., 2017), ECO (Zolfaghari et al., 2018), SAM-SS (Chen et al., 2020b) and VNS-GRU (Chen et al., 2020a). They were trained for 30, 30, 50, 80 epochs, respectively. They were evaluated on the validation set at the end



1. A women in a dress talks about data scientist she tells how they are problem solvers and well educated she starts asking how you can stand out among other data scientist.



2. A video game is displayed on the screen and in this game a man riding a motorcycle hits a car then we see a webpage with cars with a man speaking as a voice over.

**Figure 3. Redundancy samples in the MSR-VTT dataset. Caption 1 can be divided into three sentences. And Caption 2 can be divided into two or three sentences**

of each epoch. The first two models used the early stopping strategy with cross-entropy loss as the indicator. The last two models used the Comprehensive Selection Method to select a checkpoint for testing (Chen et al., 2020a). For the sake of fair comparison, the experiment settings were the same as the original papers. The two hyper-parameters  $\bar{\epsilon}$  and  $\bar{s}$  (see section 3.3) were set to 0 and 0.85 in our experiments, unless otherwise stated.

### 4.1. Evaluation Metrics

BLEU, CIDEr, METEOR and ROUGE-L are adopted as objective metrics for evaluating the results of the models. BLEU is a quick and easy-to-calculate metric, originally used for evaluating the performance of machine translation models (Papineni et al., 2002). CIDEr is a metric that captures human consensus (Vedantam et al., 2015). METEOR is a metric that involves precision, recall and order correlation, based on unigram matches (Banerjee and Lavie, 2005). ROUGE-L is a metric that determines the quality of a summary by finding the longest common subsequence (Lin, 2004). Besides these individual metrics, an overall score is presented to combine all of these metrics (Chen

**Table 2. Examples of duplicates and its similarity value**

Sentences	Similarity <sup>1</sup>
a woman is walking down the aisle in a wedding a woman is walking down the isle in a wedding dress	0.86, 0.96, 0.96
a man is talking to a woan a young man is talking to a woman	0.80, 0.94, 0.94
a woman is singing on a music video a young woman is singing in a music video	0.83, 0.94, 0.94

<sup>0</sup>Note: Each pair of sentences describe the same video clip.

<sup>1</sup>Those values are calculated by each pair of sentences with  $\bar{e} = 0, 1, 2$ .

**Table 3. Impact of each step in terms of sentences and videos**

Step	<i>Sentence</i> <sub>c</sub> <sup>1</sup>	<i>Video</i> <sub>c</sub> <sup>2</sup>
Special Characters (Section 3.1)	7,248	4,190
Spelling Mistakes (Section 3.2)	27,954	7,829
Duplicate Annotations (Section 3.3)	17,733	7,129
Successive Sentences (Section 3.4)	5,543	2,758

<sup>1</sup>The number of sentence that are corrected in a step

<sup>2</sup>The number of video that are corrected in a step

et al., 2020b):

$$O_i = \left( \frac{B4_i}{B4_b} + \frac{C_i}{C_b} + \frac{M_i}{M_b} + \frac{R_i}{R_b} \right) / 4, \quad (3)$$

where the subscript  $i$  denotes the model  $i$  and the subscript  $b$  denotes the best score of the metric  $b$  over a group of models for comparison. B4, C, M, R and O denote BLEU-4, CIDEr, METEOR, ROUGE-L and the overall score (3), respectively.

**Table 4. Influence of Edit Distance Threshold  $\bar{e}$  on the remaining annotation count and the performance of the model VNS-GRU**

$\bar{e}$	SC <sup>1</sup>	B4	C	M	R	O
0	184,078	47.6	52.6	30.4	64.1	0.9988
1	183,856	47.2	52.2	30.2	64.1	0.9931
2	183,545	47.2	52.4	30.5	64.2	0.9969

<sup>0</sup>Note:  $\bar{s} = 0.9$ . All metric values are presented in percentage.

<sup>1</sup>SC represents the number of remaining sentences in the dataset.

#### 4.2. Influence of Edit Distance Threshold and Similarity Threshold on Duplicates Removal

In the step of removing duplicated annotations, there are two hyper-parameters: the edit distance threshold  $\bar{e}$  and similarity threshold  $\bar{s}$ . The sensitivity of the hyper-parameters were investigated on the output of this step. As shown in Table 4, the threshold of edit distance  $\bar{e}$  was inversely proportionate to the

remained sentence count. The performance of the model VNS-GRU was the best when  $\bar{e} = 0$ . As shown in Table 5, the threshold of similarity was proportionate to the remained sentence count. The performance of the model VNS-GRU was the best when  $\bar{s} = 0.85$ . Table 2 shows that with the method described

**Table 5. Influence of Similarity Threshold  $\bar{s}$  on the remaining annotation count and the performance of the model VNS-GRU**

$\bar{s}$	SC <sup>1</sup>	B4	C	M	R	O
0.75	175,539	46.8	53.6	30.4	64.0	0.9850
0.80	179,169	47.6	54.2	30.4	64.3	0.9933
0.85	182,264	47.4	55.0	30.7	64.2	0.9982
0.90	183,705	47.4	53.7	30.4	64.0	0.9890
0.95	185,219	47.5	55.0	30.5	64.4	0.9978
1.00	185,330	47.6	53.4	30.2	63.9	0.9867

<sup>0</sup>Note:  $\bar{e} = 0$ .

<sup>1</sup>SC represents the number of remaining sentences in the dataset.

in the Section 3.3, we can find similar sentences, in terms of semantics, with one or two words different.

#### 4.3. Comparison between the Original/Cleaned MSR-VTT Datasets

In Table 6, a model name without any superscript indicates that the model was trained on the original training set and the metrics were calculated on the original test set. We had several observations. First, the models trained on the cleaned training set achieved higher scores of metrics than the models trained on the original training set, even though the metrics were calculated on the original test set. For instance, VNS-GRU<sup>a</sup> (Chen et al., 2020a) improves over VNS-GRU by 1.6% on BLEU-4, by 2.1% on CIDEr, by 0.9% on METEOR and by 1.1% on ROUGE-L. Second, the models trained on the cleaned training set and tested on the cleaned test set achieved higher scores of metrics than the models trained on the original training set and tested on the original test set. For instance, VNS-GRU<sup>b</sup> (Chen et al., 2020a) improves over VNS-GRU by 1.3% on BLEU-4, by 0.7% on METEOR and by 0.9% on ROUGE-L. Third, the scores of VNS-GRU<sup>b</sup> were slightly lower than the scores of

Table 6. Results on the original/cleaned MSR-VTT dataset

Model	B4	C	M	R	O
SCN (Gan et al., 2017)	42.1	48.3	28.7	61.6	0.9148
SCN <sup>1</sup>	44.3	51.5	29.7	63.0	0.9546
SCN <sup>2</sup>	44.4	50.4	29.7	63.0	0.9501
ECO (Zolfaghari et al., 2018)	43.0	49.8	28.9	62.1	0.9299
ECO <sup>1</sup>	44.4	51.6	29.5	63.1	0.9543
ECO <sup>2</sup>	44.5	50.6	29.6	63.1	0.9511
SAM-SS (Chen et al., 2020b)	43.8	51.4	28.9	62.4	0.9426
SAM-SS <sup>1</sup>	44.9	52.3	29.5	63.2	0.9605
SAM-SS <sup>2</sup>	45.0	51.1	29.5	63.2	0.9556
VNS-GRU (Chen et al., 2020a)	45.3	53.0	29.9	63.4	0.9699
VNS-GRU <sup>1</sup>	<b>46.9</b>	<b>55.1</b>	<b>30.8</b>	<b>64.5</b>	<b>0.9995</b>
VNS-GRU <sup>2</sup>	46.6	52.2	30.6	64.3	0.9828

<sup>1</sup>The model was trained on the cleaned training set and the metrics were calculated on the original test set.

<sup>2</sup>The model was trained on the cleaned training set and the metrics were calculated on the cleaned test set.

VNS-GRU<sup>a</sup>. We attributed this to the increase of annotation diversity in the cleaned test set.

The scores of BLEU-4, CIDEr, METEOR and ROUGE-L of popular video captioning models, proposed in recent years, were plotted in Fig. 4. One of the earliest models on the MSR-VTT dataset, VideoLAB, from ACM Multimedia MSR-VTT Challenge 2016 (Xu et al., 2016a), was used as the baseline, and all other models were compared with it. Then the relative changes of other models in percentage can be inferred on the right vertical axes in Fig. 4. By training on the cleaned training set, one of the state-of-the-art models, VNS-GRU was improved from 15.9% to 19.9% on BLEU-4, from 20.2% to 24.7% on CIDEr, from 7.9% to 10.8% on METEOR, from 4.6% to 6.1% on ROUGE-L, compared with the results obtained by the same model trained on the original training set. From the figure, it is seen that the relative improvements brought by annotation cleaning were non-negligible.

#### 4.4. Ablation Study

To analyze the utility of each step in data cleaning, we compared the performances of the model VNS-GRU (Chen et al., 2020a) on the original and cleaned test sets in Tables 7 and 8, trained on the training set cleaned by Step I (Section 3.1), Step II (Section 3.2), Step III (Section 3.3), Step IV (Section 3.4), accumulatively.

As shown in Tables 7 and 8, Step I brought improvements in all the metrics since it reduced the number of irregular words and phrases, which contain special characters. After Step II, the four metrics remained similar to those after Step I when measured on the original test set (Table 7), but the metrics were improved when measured on the cleaned test set (Table 8). After Step III, all metrics except METEOR increased in the both cases. The METEOR value slightly decreased when measured on the cleaned test set (Table 8). After the last step, almost all metrics were further improved, except BLEU-4. If we focus on the performance of the model measured on the cleaned test set

Table 7. Results on the origin test set

I <sup>1</sup>	II	III	IV	B4	C	M	R	O
×	×	×	×	45.3	53.0	29.9	63.4	0.9678
√	×	×	×	47.1	54.4	30.4	64.0	0.9901
√	√	×	×	47.3	53.9	30.2	64.1	0.9876
√	√	√	×	47.4	55.0	30.7	64.2	0.9975
√	√	√	√	46.9	55.1	30.8	64.5	0.9974

<sup>1</sup>The model was trained on the training set with data cleaning steps I, II, III and IV taken one by one.

(Table 8), we found that the overall score was improved after each step. These results suggest that all steps are necessary for cleaning the annotations.

Table 8. Results on the cleaned test set. The model was trained on the training set with data cleaning steps I, II, III and IV taken one by one

I	II	III	IV	B4	C	M	R	O
×	×	×	×	44.5	49.8	29.7	63.0	0.9598
√	×	×	×	46.9	50.9	30.2	63.8	0.9849
√	√	×	×	46.9	51.4	30.3	63.9	0.9885
√	√	√	×	47.6	51.7	30.2	64.1	0.9936
√	√	√	√	46.6	52.2	30.6	64.3	0.9947

## 5. Human Evaluation

It is well-known that the metrics including BLEU-4 (Papineni et al., 2002), CIDEr (Vedantam et al., 2015), METEOR (Banerjee and Lavie, 2005), ROUGE-L (Lin, 2004) do not fully reflect the quality of the video captioning results. We then






Figure 4. The performance of typical models on the MSR-VTT dataset during 2016 and 2020. The models include VideoLAB, Aalto, v2t\_navigator, MTVC (Pasunuru and Bansal, 2017a), CIDEnt-RL (Pasunuru and Bansal, 2017b), SibNet (Liu et al., 2018), HACA (Wang et al., 2018), TAMoE (Wang et al., 2019d), SAM-SS (Chen et al., 2020b) and POS\_RL (Wang et al., 2019b) and VNS-GRU (Chen et al., 2020a). The first three models are from ACM Multimedia MSR-VTT Challenge 2016 (Xu et al., 2016a). VideoLAB was used as the baseline (0% change).

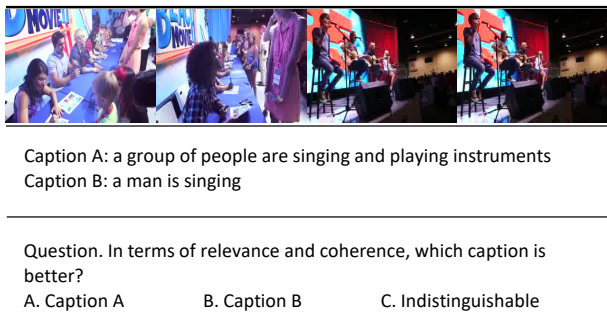


Figure 5. An example question in the human evaluation experiment. Captions A and B were generated by VNS-GRU or VNS-GRU\*

conducted a human evaluation study. We recruited 17 people (11 male and 6 female, ages between 20 and 35) with normal or corrected-to-normal vision to do this experiment. The subjects were mainly from Tsinghua University, Beijing, China. All subjects had at least college level English. This study was approved by the Department of Psychology Ethics Committee, Tsinghua University, Beijing, China.

The subjects watched video clips from the MSR-VTT dataset and compared the results of VNS-GRU trained on the original and cleaned annotations of the dataset (Figure 5). The subjects were instructed to compare the results based on two criteria:

1. relevance, the match between the contents of the video clip and the caption;
2. coherence, the language fluency and grammatical correctness in the caption.


For each video clip, there were three options: (A) Caption A is better; (B) Caption B is better; and (C) Indistinguishable. The two captions were generated by VNS-GRU or VNS-GRU\*, which were trained on the original and cleaned annotations of the dataset, respectively. The subjects needed to choose one and only one of three options. A total of 30 video clips were randomly sampled from the test set and presented to all subjects in an fixed order. Every subject completed the experiment within half an hour.

We noted down the number of votes for VNS-GRU, VNS-GRU\* and Indistinguishable for every subject and calculated the average over all subjects (Figure 6). On average, for 11.8 video clips the subjects voted for “VNS-GRU\* is better” and for 10.1 video clips the subjects voted for “VNS-GRU is better”. The one-sided student t-test indicated that VNS-GRU\* performed better than VNS-GRU ( $p = 0.02, n = 17$ ). On average, for 8.1 videos the subjects could not distinguish the quality of the results.

These results suggested that annotation cleaning could boost the quality of the generated captions by video captioning models from subjective evaluation of humans.

## 6. Conclusion

The MSR-VTT dataset is a widely used dataset in the areas of video captioning and video retrieval. Thousands of problems were found in its annotations, and many of them were obvious mistakes. We inspected the influence of these problems on the results of video captioning models. By four steps of data cleaning, we removed or corrected sentences to resolve these



**Figure 6. Human evaluation results.** “VNS-GRU\*”, “VNS-GRU” and “Indistinguishable” denote the numbers of videos which the subjects voted for “VNS-GRU\* is better than VNS-GRU”, “VNS-GRU is better than VNS-GRU” and “They are indistinguishable”, respectively. Error bars are standard deviations. The p-value between “VNS-GRU\*” and “VNS-GRU” is 0.02

problems, and compared the results of several popular video captioning models. The models trained on the cleaned dataset generated better captions than the models trained on the original dataset measured by both objective metrics and subjective evaluations. In particular, trained on the cleaned dataset, VNS-GRU achieved better results with improvement of at least 0.9% compared to the baseline. This cleaned dataset is recommended for developing new video captioning models in the future.

## Acknowledgments

The authors would like to thank Han Liu and Huiran Yu for insightful discussions. This work was supported by the National Natural Science Foundation of China under Grant 62061136001, Grant U19B2034 and Grant 61620106010, and by the deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – TRR 169/A6.

## References

- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L., 2018. Bottom-up and top-down attention for image captioning and visual question answering, in: Proceedings of the IEEE conference on computer vision and pattern recognition, CVPR, pp. 6077–6086.
- Awad, G., Fiscus, J., Joy, D., Michel, M., Smeaton, A., Kraaij, W., Eskevich, M., Aly, R., Ordelman, R., Ritter, M., et al., 2016. Trecvid 2016: Evaluating video search, video event detection, localization, and hyperlinking, in: TREC Video Retrieval Evaluation (TRECVID).
- Bain, M., Nagrani, A., Varol, G., Zisserman, A., 2021. Frozen in time: A joint video and image encoder for end-to-end retrieval, in: IEEE International Conference on Computer Vision.
- Banerjee, S., Lavie, A., 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments, in: Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization, pp. 65–72.
- Bertasius, G., Wang, H., Torresani, L., 2021. Is space-time attention all you need for video understanding. arXiv preprint arXiv:2102.05095 2, 4.
- Chen, H., Li, J., Hu, X., 2020a. Delving deeper into the decoder for video captioning, in: ECAI 2020 - 24th European Conference on Artificial Intelligence, pp. 1079–1086.
- Chen, H., Lin, K., Maye, A., Li, J., Hu, X., 2020b. A semantics-assisted video captioning model trained with scheduled sampling. *Frontiers in Robotics and AI* 7, 129.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y., 2014. Learning phrase representations using rnn encoder–decoder for statistical machine translation, in: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP, pp. 1724–1734.
- Dong, J., Li, X., Xu, C., Yang, X., Yang, G., Wang, X., Wang, M., 2021. Dual encoding for video retrieval by text. *IEEE Transactions on Pattern Analysis and Machine Intelligence* doi:10.1109/TPAMI.2021.3059295.
- Gabeur, V., Sun, C., Alahari, K., Schmid, C., 2020. Multi-modal transformer for video retrieval, in: European Conference on Computer Vision, Springer, pp. 214–229.
- Gan, Z., Gan, C., He, X., Pu, Y., Tran, K., Gao, J., Carin, L., Deng, L., 2017. Semantic compositional networks for visual captioning, in: Proceedings of the IEEE conference on computer vision and pattern recognition, CVPR, pp. 5630–5639.
- Guadarrama, S., Krishnamoorthy, N., Malkarnenkar, G., Venugopalan, S., Mooney, R., Darrell, T., Saenko, K., 2013. Youtube2text: Recognizing and describing arbitrary activities using semantic hierarchies and zero-shot recognition, in: Proceedings of the IEEE international conference on computer vision, ICCV, pp. 2712–2719.
- Levenshtein, V.I., 1966. Binary codes capable of correcting deletions, insertions, and reversals, in: *Soviet physics doklady*, pp. 707–710.
- Lin, C.Y., 2004. ROUGE: A package for automatic evaluation of summaries, in: *Text Summarization Branches Out*, pp. 74–81.
- Liu, S., Fan, H., Qian, S., Chen, Y., Ding, W., Wang, Z., 2021. Hit: Hierarchical transformer with momentum contrast for video-text retrieval, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 11915–11925.
- Liu, S., Ren, Z., Yuan, J., 2018. Sibnet: Sibling convolutional encoder for video captioning, in: Proceedings of the 26th ACM International Conference on Multimedia, p. 1425–1434.
- Papineni, K., Roukos, S., Ward, T., Zhu, W.J., 2002. BLEU: a method for automatic evaluation of machine translation, in: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, ACL, pp. 311–318.
- Pasunuru, R., Bansal, M., 2017a. Multi-task video captioning with video and entailment generation, in: Barzilay, R., Kan, M. (Eds.), Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL, pp. 1273–1283.
- Pasunuru, R., Bansal, M., 2017b. Reinforced video captioning with entailment rewards, in: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP, pp. 979–985.
- Pei, W., Zhang, J., Wang, X., Ke, L., Shen, X., Tai, Y.W., 2019. Memory-attended recurrent network for video captioning, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, pp. 8347–8356.
- Rennie, S.J., Marcheret, E., Mroueh, Y., Ross, J., Goel, V., 2017. Self-critical sequence training for image captioning, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, pp. 7008–7024.
- Vedantam, R., Zitnick, C.L., Parikh, D., 2015. Cider: Consensus-based image description evaluation, in: IEEE Conference on Computer Vision and Pattern Recognition, CVPR, pp. 4566–4575.
- Wang, B., Ma, L., Zhang, W., Jiang, W., Wang, J., Liu, W., 2019a. Controllable video captioning with pos sequence guidance based on gated fusion network, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, ICCV, pp. 2641–2650.
- Wang, B., Ma, L., Zhang, W., Jiang, W., Wang, J., Liu, W., 2019b. Controllable video captioning with pos sequence guidance based on gated fusion network, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, ICCV, pp. 2641–2650.
- Wang, X., Wang, Y., Wang, W.Y., 2018. Watch, listen, and describe: Globally and locally aligned cross-modal attentions for video captioning, in: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 795–801.
- Wang, X., Wu, J., Chen, J., Li, L., Wang, Y.F., Wang, W.Y., 2019c. Vatex:

- A large-scale, high-quality multilingual dataset for video-and-language research, in: 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 4580–4590. doi:10.1109/ICCV.2019.00468.
- Wang, X., Wu, J., Zhang, D., Su, Y., Wang, W.Y., 2019d. Learning to compose topic-aware mixture of experts for zero-shot video captioning, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 8965–8972.
- Wu, C., Wei, Y., Chu, X., Sun, W., Su, F., Wang, L., 2018. Hierarchical attention-based multimodal fusion for video captioning. *Neurocomputing* 315, 362–370.
- Xu, J., Mei, T., Yao, T., Rui, Y., 2016a. The 1st video to language challenge. URL: <http://ms-multimedia-challenge.com/2016/challenge>.
- Xu, J., Mei, T., Yao, T., Rui, Y., 2016b. MSR-VTT: A large video description dataset for bridging video and language, in: Proceedings of the IEEE conference on computer vision and pattern recognition, CVPR, pp. 5288–5296.
- Yu, N., Hu, X., Song, B., Yang, J., Zhang, J., 2018. Topic-oriented image captioning based on order-embedding. *IEEE Transactions on Image Processing* PP, 1–1.
- Zhang, Z., Qi, Z., Yuan, C., Shan, Y., Li, B., Deng, Y., Hu, W., 2021. Open-book video captioning with retrieve-copy-generate network, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9837–9846.
- Zolfaghari, M., Singh, K., Brox, T., 2018. ECO: efficient convolutional network for online video understanding, in: Proceedings of the European conference on computer vision ECCV, pp. 713–730.