

Multimodal Matching Transformer for Live Commenting

Chaoqun Duan¹ and Lei Cui² and Shuming Ma³ and Furu Wei⁴ and Conghui Zhu⁵ and Tiejun Zhao⁶

Abstract. Automatic live commenting aims to provide real-time comments on videos for viewers. It encourages users engagement on online video sites, and is also a good benchmark for video-to-text generation. Recent work on this task adopts encoder-decoder models to generate comments. However, these methods do not model the interaction between videos and comments explicitly, so they tend to generate popular comments that are often irrelevant to the videos. In this work, we aim to improve the relevance between live comments and videos by modeling the cross-modal interactions among different modalities. To this end, we propose a multimodal matching transformer to capture the relationships among comments, vision, and audio. The proposed model is based on the transformer framework and can iteratively learn the attention-aware representations for each modality. We evaluate the model on a publicly available live commenting dataset. Experiments show that the multimodal matching transformer model outperforms the state-of-the-art methods.

1 Introduction

Live commenting is an emerging feature of online video sites that allows real-time comments to fly across the screen or roll at the right side of the videos, so that viewers can see comments and videos at the same time. Automatic live commenting aims to provide some additional opinions of videos and respond to live comments from other viewers, which encourages users engagement on online video sites. Automatic live commenting is also a good testbed of a model’s ability of dealing with multi-modality information [16]. It requires the model to understand the vision, text, and audio, and organize the language to produce the comments of the videos. Therefore, it is an interesting and important task for human-AI interaction.

Although great progress has been made in multimodal learning [15, 24, 25], live commenting is still a challenging task. Recent work on live commenting implements an encoder-decoder model to generate the comments [16]. However, these methods do not model the interaction between the videos and the comments explicitly. Therefore, the generated comments are often general to any videos and irrelevant to the specific input videos. Figure 1 shows an example of the generated comments by an encoder-decoder model. It shows that the encoder-decoder model tends to output the popular sentences, such as “Oh my God !”, while the reference comment is much more informative and relevant to the video. The reason is that the encoder-decoder model cares more about the language model, rather than the interaction between the videos and the comments, so



Case 1	Oh My God !!!!!
Case 2	So am I.
Reference	The cat is afraid of spicy.

Figure 1: An example of the generated comments by the encoder-decoder model. Above is a frame extracted from a selected video. Below are two cases generated by the encoder-decoder model around the above frame, as well as a reference comment by human.

generating popular comments is a safe way for the model to reduce the empirical risk. As a result, the encoder-decoder model is more likely to generate a frequent sentence, rather than an informative and relevant comment.

Another problem with current state-of-the-art live commenting models is that they do not take the audio into consideration. Audio, as an important part of videos, carries information that may not appear in the vision or text. For example, when the video is about playing the piano, it is difficult to make a proper comment without the audio. The audio also includes dialogues or background music, which helps understand the story in videos. Therefore, the audio should not be neglected if the model needs to fully understand videos and make an informative comment.

In this work, we build a novel live commenting model to make more relevant comments. Based on existing observations, we propose a multimodal matching transformer to learn the cross-modal interaction between videos and comments explicitly. The proposed multimodal matching network can match the most relevant comments with the given videos from a candidate set, so it can encourage the produced comments to be more informative and less general. Our model is based on the transformer architecture, and it jointly learns the cross-modal representations of text, vision, and audio. We evaluate our model on a live commenting dataset [16]. Experiments show that the proposed multimodal matching transformer model is effective and significantly

¹ Harbin Institute of Technology, China, email: cqduan@stu.hit.edu.cn

² Microsoft Research Asia, China, email: lecu@microsoft.com

³ Microsoft Research Asia, China, email: Shuming.Ma@microsoft.com

⁴ Microsoft Research Asia, China, email: fuwei@microsoft.com

⁵ Harbin Institute of Technology, China, email: conghui@hit.edu.cn

⁶ Harbin Institute of Technology, China, email: tjzhao@hit.edu.cn

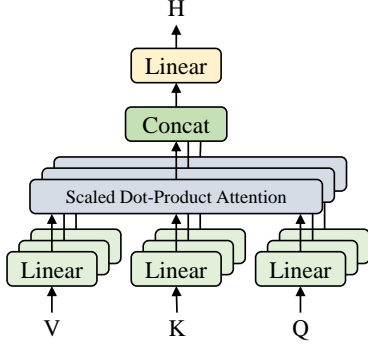


Figure 2: Architecture of the multi-head attention.

outperforms state-of-the-art methods.

The contributions of this paper can be summarized as follows:

- We propose using the audio information for the task of live commenting, which is neglected by previous work.
- We propose a novel multimodal matching network to capture the relationship among text, vision, and audio, based on the state-of-the-art transformer framework.
- Experiments show that the proposed multimodal matching model significantly outperforms the state-of-the-art methods.

2 Background

In this work, our proposed model is based on the Transformer network [20]. This section will give an introduction for core modules of the Transformer network.

2.1 Input Representation

In the Transformer network architecture, there is no recurrence and no convolution. To leverage the order of the sequence, it introduces positional embeddings and each positional embedding is computed based on the token's position in the sequence:

$$\begin{aligned} PE_{(pos,2i)} &= \sin(pos/10000^{2i/d_{model}}) \\ PE_{(pos,2i+1)} &= \cos(pos/10000^{2i/d_{model}}) \end{aligned} \quad (1)$$

where pos is the position in the sequence and i is the dimension.

Namely, the input representation of the Transformer network contains two parts: word embeddings $\{e_{w_1}, e_{w_2}, \dots, e_{w_n}\}$ and positional embeddings $\{e_{p_1}, e_{p_2}, \dots, e_{p_n}\}$, where n is the length of the input sentence. The two parts are fused by an addition operation.

2.2 Multi-Head Attention

After obtaining outputs of previous layers, the Transformer network uses a multi-head attention mechanism to learn the context-aware representation for the sequence.

Figure 2 shows the architecture of the multi-head attention. Q , K and V , three matrices, are inputs derived from previous layers and H is the output. The Multi-head attention can be denoted as:

$$\begin{aligned} H &= MultiHead(Q, K, V) \\ &= Concat(O_1, O_2, \dots, O_h)W^O \end{aligned} \quad (2)$$

where W^O is a trainable parameter. h is the number of parallel attention layers. O_i is computed by Eq. (3).

$$O_i = Attention(QW_i^Q, KW_i^K, VW_i^V) \quad (3)$$

where W_i^Q , W_i^K and W_i^V are trainable parameters and $Attention(;;)$ is the scaled dot-product attention which can be denoted as:

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (4)$$

3 Multimodal Matching Transformer

Automatic live commenting aims to make comments on a video clip. According to the analysis in [16], the live comments are relevant to not only the video clip, but also the surrounding comments from other viewers. In this work, we find it helpful to incorporate the audio information into the live commenting model. However, the surrounding comments, the vision part, and the audio part of the videos are from different modalities, and not trivial to model their relationships. To address these issues, we propose a multimodal matching model, which we denote as Matching Transformer. The model is based on the popular transformer architecture [20, 9]. Our model can jointly learn the cross-modal representations of textual context, visual context, audio context, and model the relationships among them.

3.1 Task Definition

We formulate the automatic live commenting as a ranking problem. Formally, given a video V and a time-stamp t , automatic live commenting aims to select a comment y^* from a candidate set Y , which is most relevant to the video clip near the time-stamp based on the surrounding comments C , the visual part F and the audio part A . Concretely, we extract N_c comments near the time-stamp t as $C = \{c_1, c_2, \dots, c_{N_c}\}$, where c_i is a comment. For the vision part F , we sample N_f video frames near the time-stamp t as $F = \{f_1, f_2, \dots, f_{N_f}\}$, where f_i is an image and the interval between two images is 1 second. We convert a 5-second audio clip surrounding the time-stamp t to a log magnitude mel-frequency spectrogram, and it can be denoted as $A = \{a_1, a_2, \dots, a_{N_a}\}$, where a_i is a vector and N_a is the length of the audio clip. The candidate comment can be denoted as $y = \{y_1, y_2, \dots, y_k\}$, where y_i is a word and k is the number of words.

In this way, the task can be formulated as searching the most relevant comment to the video clip in the multimodal semantic space:

$$y^* = \arg \max_{y \in Y} S_m(C, F, A, y) \quad (5)$$

where S_m is a model to produce the similarity between (C, F, A) and y .

3.2 Model Overview

Figure 3 shows the architecture of the Matching Transformer. The model consists of three components: (1) **an encoder layer** converts different modalities of a video clip (including surrounding comments, the vision part of the video, the audio part of the

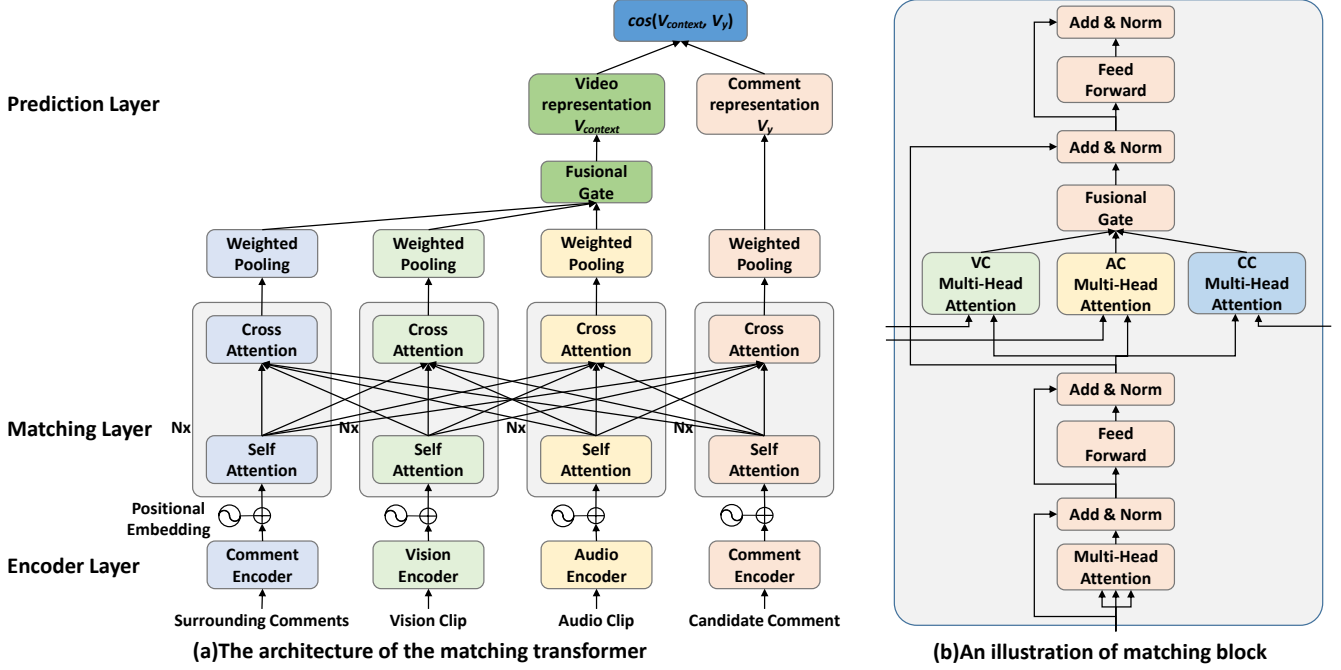


Figure 3: Architecture of the matching transformer. Part (a) illustrates the whole architecture. In this model, the matching layer consists of N matching blocks and part (b) illustrates the structure of a matching block.

video) and a candidate comment into vectors; (2) a **matching layer** iteratively learns the attention-aware representation for each modality; (3) a **prediction layer** outputs a score measuring the matching degree between a video clip and a comment.

Formally, given different contexts of a video clip (C, F, A) and a comment y , the model can be denoted as:

$$s = S_m(C, F, A, y) \quad (6)$$

Next, we introduce each layer in detail.

3.3 Encoder Layer

As shown in part (a) of Figure 3, our model contains three kinds of encoder: a comment encoder, a vision encoder and an audio encoder. These encoders convert a comment, a vision clip and an audio clip into vectors respectively. In our model, the existing surrounding comments C and the candidate comment y share the same comment encoder.

Comment Encoder In our model, N_c comments near the timestamp, $C = \{c_1, c_2, \dots, c_{N_c}\}$, are first concatenated into one comment $C = \{w_1, w_2, \dots, w_{l_c}\}$, where w_i is the i -th word in the comment and l_c is the total number of words. Then, the comment encoder converts words of the comment into vectors $\{e_{c_1}, e_{c_2}, \dots, e_{c_{l_c}}\}$ by looking up M , where $M \in \mathbf{R}^{d \times |V|}$ is the embedding table. d is the dimension of the embedding and $|V|$ is the size of the vocabulary. Similarly, the comment encoder also converts the candidate comment y into vectors: $\{e_{y_1}, e_{y_2}, \dots, e_{y_k}\}$.

Vision Encoder The vision encoder converts a vision clip $F = \{f_1, f_2, \dots, f_{N_f}\}$ into vectors $\{e_{f_1}, e_{f_2}, \dots, e_{f_{l_f}}\}$ by a pre-trained model, where l_f is equal to N_f . Similar to [16], we leverage a pre-trained 18-layer ResNet [13] to encode the frames within a vision

clip. It can be denoted as:

$$e_{f_i} = ResNet(f_i) \quad (7)$$

Audio Encoder For audio encoding, we first slice a 5-second audio clip $A = \{a_1, a_2, \dots, a_{N_a}\}$ into five audio frame sets, $\{\{a_1^t, a_2^t, \dots, a_{l_a}^t\}\}_{t=1}^5$, based on the timestamp. Then, we use a GRU [7] to encode each set. It can be denoted as:

$$h_i^t = GRU(a_i^t, h_{i-1}^t) \quad (8)$$

At last, we use the last hidden state of each set $\{h_{l_a}^1, h_{l_a}^2, \dots, h_{l_a}^5\}$ as the representation of the audio clip: $\{e_{a_1}, e_{a_2}, \dots, e_{a_{l_a}}\}$.

Positional Embedding To exploit the temporal information in each modality, following [20], we also use positional embedding (PE) by adding it to the output of each encoder.

3.4 Matching Layer

Inspired by the recent successful deep learning frameworks [13, 20], we adopt a matching layer which consists of N matching blocks to iteratively learn the attention-aware representation for each modality. The structure of a matching block is shown in part (b) of Figure 3. Each matching block is composed of four parts: a multi-head self-attention, a multi-head cross attention and two position-wise FNN. Compared to the basic block defined in [20], our matching block adds a multi-head cross attention and a position-wise FNN. We use these auxiliary mechanisms to learn attention-aware representation from other modalities.

For simplicity, we take the candidate comment as the example to illustrate the matching layer.

Formally, in the t -th block, given the output of previous matching block corresponding to the candidate comment: $H_y^{t-1} =$

$\{h_{y_1}^{t-1}, h_{y_2}^{t-1}, \dots, h_{y_k}^{t-1}\}$, we first utilize a multi-head self-attention and a position-wise FNN to learn the context of the candidate comment \hat{H}_y^t :

$$\tilde{H}_y^t = \text{MultiHead}(H_y^{t-1}, H_y^{t-1}, H_y^{t-1}) \quad (9)$$

$$\hat{H}_y^t = \text{MLP}(\text{ReLU}(\text{MLP}(\tilde{H}_y^t))) \quad (10)$$

Similar to Eq. (9) and Eq. (10), we also compute the context vectors of surrounding comment \hat{H}_c^t , visual clip \hat{H}_f^t , and audio clip \hat{H}_a^t . Then we employ a multi-head cross attention to learn the attention-aware representation from each modality:

$$\tilde{H}_{yc}^t = \text{MultiHead}(\hat{H}_y^t, \hat{H}_c^t, \hat{H}_c^t) \quad (11)$$

$$\tilde{H}_{yf}^t = \text{MultiHead}(\hat{H}_y^t, \hat{H}_f^t, \hat{H}_f^t) \quad (12)$$

$$\tilde{H}_{ya}^t = \text{MultiHead}(\hat{H}_y^t, \hat{H}_a^t, \hat{H}_a^t) \quad (13)$$

After getting these three attention-aware representations, we use MLP to build a fusional gate and combine them with the weighted sum:

$$g_c = \text{MLP}(\tilde{H}_{yc}^t, \tilde{H}_{yf}^t, \tilde{H}_{ya}^t) \quad (14)$$

$$\begin{aligned} \tilde{H}_y^t &= g_{c[1:d]} \odot \tilde{H}_{yc}^t + g_{c[d:2d]} \odot \tilde{H}_{yf}^t \\ &+ g_{c[2d:]} \odot \tilde{H}_{ya}^t \end{aligned} \quad (15)$$

where \odot means the element-wise dot and d is the dimension of \tilde{H}_{yc}^t , \tilde{H}_{yf}^t , and \tilde{H}_{ya}^t .

Finally, we feed \tilde{H}_y^t into a position-wise FNN to produce the output of the t -th matching block corresponding to the candidate comment:

$$H_y^t = \text{MLP}(\text{ReLU}(\text{MLP}(\tilde{H}_y^t))) \quad (16)$$

As described above, Eq. (9)-Eq. (16) illustrate how to compute the representation of candidate comment H_y^t . In implementation, we adopt the same way to compute the representations of surrounding comment H_c^t , vision clip H_f^t and audio clip H_a^t .

3.5 Prediction Layer

The prediction layer outputs a score measuring the matching degree between (C, F, A) and y . In this layer, we first employ a weighted pooling to convert the output of the last matching block to a fixed-length vector:

$$V_y = A_y^p H_y^N \quad (17)$$

$$A_y^p = \text{softmax}(\text{ReLU}(H_y^N \mathbf{W}_1^p + \mathbf{b}_1^p) \mathbf{W}_2^p + \mathbf{b}_2^p) \quad (18)$$

where \mathbf{W}_1^p , \mathbf{W}_2^p , \mathbf{b}_1^p and \mathbf{b}_2^p are trainable parameters.

Similarly, we get the vectors V_c , V_f and V_a for C , F and A respectively. Then, we adopt a fusional gate to combine V_c , V_f and V_a into $V_{context}$:

$$g_v = \text{MLP}(V_c, V_f, V_a) \quad (19)$$

$$\begin{aligned} V_{context} &= g_{v[1:d]} \odot V_c + g_{v[d:2d]} \odot V_f \\ &+ g_{v[2d:]} \odot V_a \end{aligned} \quad (20)$$

where d is the dimension of the V_c , V_f and V_a .

Finally, we use a cosine distance to measure the similarity between $V_{context}$ and V_y :

$$s = \cos(V_{context}, V_y) \quad (21)$$

3.6 Training

To learn the $S_m(\cdot, \cdot, \cdot, \cdot)$, we leverage the max-margin loss function, which can be formulated as:

$$\begin{aligned} L(\theta) &= \frac{1}{N} \sum_{i=1}^N \max(0, M \\ &+ S_m(\mathbf{F}^{(i)}, \mathbf{A}^{(i)}, \mathbf{C}^{(i)}, \mathbf{y}^{(i)-}; \theta) \\ &- S_m(\mathbf{F}^{(i)}, \mathbf{A}^{(i)}, \mathbf{C}^{(i)}, \mathbf{y}^{(i)+}; \theta)) \end{aligned} \quad (22)$$

where N is the number of instances in the training set, $(\mathbf{F}^{(i)}, \mathbf{A}^{(i)}, \mathbf{C}^{(i)}, \mathbf{y}^{(i)-})$ is the negative sample and $(\mathbf{F}^{(i)}, \mathbf{A}^{(i)}, \mathbf{C}^{(i)}, \mathbf{y}^{(i)+})$ is the positive sample. M is the margin that needs to be specified manually. θ denotes all the trainable parameters of our model. When training, we employ Adam [14] as the optimizer.

4 Experiments

4.1 Dataset

We evaluate our model on a live commenting dataset⁷ that is released by [16]. The live commenting dataset is a large-scale video-comment dataset. It contains 2,361 videos and 895,929 comments. The data is collected from a popular Chinese video streaming website, Bilibili. Therefore, it has strong authenticity and practicability. In our experiment, we use the same partition as in [16]. The detailed statistics of the dataset is shown in Table 1.

Table 1: Statistics of the Live Comment Dataset.

	Train	Dev	Test	Total
#Video	2,161	100	100	2,361
#Comment	820k	42k	34k	896k
#Word	4,419k	248k	193k	4,860k
Avg. Words	5.39	5.85	5.58	5.42
Hours	103.81	5.02	5.01	113.84

4.2 Evaluation Metric

Following the previous work [8, 16], we adopt **Recall@k**, **Mean Recall (MR)** and **Mean Reciprocal Rank (MRR)** for automatic evaluation, which are standard evaluation metrics of the ranking task. For testing, we construct a candidate comments set in which each video clip contains 100 comments, which is exactly the same as the previous work [16] for fair comparison. The comments in the candidate comment set are comprised of three parts: (1) the ground-truth comments; (2) top 20 popular comments; (3) random selected comments. We evaluate our model on the testing set.

In addition, we also test the performance of our model by human evaluation. Following [16], we use the metrics of **relevance (Rel)** and **correctness (Cor)** to evaluate our model. Relevance measures the relevance between produced comments and videos and correctness measures the confidence that produced comments are made by human in the context of the video. We do not evaluate the fluency of the produced comments, because our model just selects a proper comment from a candidate comment set, which is naturally fluent. For both relevance and correctness, we use a score $s \in \{1, 2, 3, 4, 5\}$

⁷ <https://github.com/lancopku/livebot>

Table 2: The performance comparison on the live commenting dataset (**Recall@k**, **MRR**: higher is better; **MR**: lower is better). Our matching transformer significantly outperforms the baselines in terms of all metrics. Meanwhile, our model achieves better performance than baselines by using the same two modalities.

Model	Text	Vision	Audio	Recall@1	Recall@5	Recall@10	MR	MRR
S2S	✓	✓		12.89	33.78	50.29	17.05	0.2454
Fusional RNN	✓	✓		17.25	37.96	56.10	16.14	0.2710
Unified Transformer	✓	✓		18.01	38.12	55.78	16.01	0.2753
Matching Transformer-C	✓			18.02	42.83	59.37	12.28	0.3087
Matching Transformer-CF	✓	✓		22.77	46.71	62.87	11.19	0.3519
Matching Transformer-CFA	✓	✓	✓	23.52	46.99	64.24	11.05	0.3596

Table 3: Effect of different modalities used in the Matching Transformer (**Recall@k**, **MRR**: higher is better; **MR**: lower is better). It shows that more modalities always lead to better performance, which indicates that the proposed model can capture the semantic information of different modalities to help the live commenting task.

	Text	Vision	Audio	Recall@1	Recall@5	Recall@10	MR	MRR
Single-Modal	✓			18.02	42.83	59.37	12.28	0.3087
		✓		18.55	38.38	50.98	16.33	0.2920
			✓	17.95	36.89	50.52	15.33	0.2861
Double-Modal	✓	✓		22.77	46.71	62.87	11.19	0.3519
	✓		✓	19.93	44.39	59.68	12.21	0.3276
		✓	✓	18.03	39.00	52.77	15.60	0.2933
Triple-Modal	✓	✓	✓	23.52	46.99	64.24	11.05	0.3596

Table 4: Human evaluation results of different models (**Rel** refers to the relevance score; **Cor** refers to the correctness score; **Human** means the natural comments in the dataset.). It shows that the produced comments of our model are more relevant than those of the baseline models. Besides, our model can produce more correct and proper comments.

Model	Rel	Cor
S2S	2.23	2.91
Fusional RNN	2.95	3.34
Unified Transformer	3.07	3.45
Matching Transformer	3.25	3.57
Human	3.31	4.11

to denote the degree, the higher the better. When testing, three human annotators are asked to give a score to evaluate the top one comment produced by our model and we use the average score as the final result.

4.3 Settings

In our experiments, the word embeddings and video frame vectors are in 512 dimensions while audio frame vectors are in 64 dimensions. The GRU for audio encoder is in 512 dimensions. For positional embedding, we use fixed sinusoidal positional embedding and set the dimension as 512. The word embeddings are randomly initialized and updated during training, while the video frame vectors and audio frame vectors are fixed. There are 6 matching blocks in the matching layer. In each matching block, the number of heads in the multi-head attention is 8 and the dimension of the position-wise FNN is 2,048. The margin M is set to 0.1 in our experiment. We employ the Adam [14] for training, whose default hyper-parameters β_1 and β_2 are set to 0.9 and 0.999 for optimization respectively. The initial learning rate of Adam is set to 0.00009. The learning rate is halved when the accuracy on the development set drops. We also employ a

dropout strategy [18] and layer normalization [4] to reduce the risk of over-fitting. The dropout rate is set to 0.2 and the batch size is 64.

For pre-processing, we use the Stanford tokenizer [17] to tokenize the comments and audio feature extractor⁸ released by [12] to process the audio clip. During training, we draw 1 negative sample for each video clip.

4.4 Baselines

- **S2S** [21] is a traditional sequence to sequence model without the attention mechanism. Specifically, the model uses two encoders to encode visual and textual information respectively. During decoding, the decoder uses the concatenation of the outputs from the two encoders as input.
- **Fusional RNN** [16] consists of three parts: a video encoder, a comment encoder and a comment decoder. The three parts are all RNN-based networks and they are related by an attention layer. This model uses the visual and textual context as input.
- **Unified Transformer** [16] is a transformer-based generative model. Similar to Fusional RNN, this model is comprised of three parts: a video encoder, a comment encoder and a comment decoder. The difference is that these three parts are all stacked attention-based transformer blocks.

4.5 Overall Results

Table 2 shows the automatic evaluation results of the baseline models and our proposed models. The baselines (S2S, Fusional RNN and Unified Transformer) only use the text and vision of the videos. Our matching transformer leverages three modalities (text, vision, and audio) and significantly outperforms the baselines in terms of all metrics. Moreover, we also report the results of our model

⁸ <https://github.com/tensorflow/models/tree/master/research/audioset>



(a) Frame 1.

(b) Frame 2.

(c) Frame 3.

Surrounding Comments	饿啦 (I'm hungry.) 我在中国都没有的吃, 你说说 (I didn't eat it in China. What a pity.) 这叫灌汤包 (This is soup dumpling.)
Target Comments	好久没吃小笼包啦!想吃想吃想吃 (I haven't eaten soup dumplings for a long time.) 香香的口味, 你尝过没 (Very delicious. Have you tried it yet?) 小笼包也有汤汁的 (There's soup in the soup dumpling.)
S2S	你的鸡蛋怎么和我家的不一样呢 (Why your eggs are different from mine?)
Fusional RNN	跳的是司马的舞 (Dancing)
Unified Transformer	掉了的包子都去喂流浪狗了 (A soup dumpling that fell on the ground are feed the stray dogs.)
Matching Transformer	好久没吃小笼包啦!想吃想吃想吃 (I haven't eaten soup dumplings for a long time.)

Figure 4: An example of the produced comments of different models on a video. Above are three selected frames in the videos. Below are the existing comments in the video and the produced comments of different models.

with only one modality (less than baselines) and two modalities (equal to baselines). It shows that the matching transformer achieves comparable performance with the baselines using only one modality. Meanwhile, our model achieves better performance than baselines by using the same two modalities, which verifies the efficiency of the proposed model. Finally, the triple-modality model significantly outperforms the baselines, achieving +5.51 points on Recall@1, +8.87 points on Recall@5 and +8.46 points on Recall@10.

4.6 Effect of Different Modalities

We also would like to know how different kinds of modalities contribute to our proposed model. Therefore, we conduct ablation experiments by removing different modalities from our model. Table 3 summarizes the results of the ablation experiments. Under the single-modality setting, it shows that the model with the modality of text achieves better performance over the other two modalities. Among the possible alternatives of double-modality, the combination of text and vision obtains the best performance. Finally, the model with triple-modalities get the highest scores in terms of all the automatic metrics. Besides, it is observed that more modalities always lead to better performance, which indicates that the proposed model can capture the semantic information of different modalities to help the live commenting task.

4.7 Human Evaluation

We randomly sample 100 video clips from the test set to evaluate our model in terms of the relevance and the correctness. For both metrics, we use a score range from 1 to 5 to denote the degree, where the higher the better. We have three human annotators to give a score that evaluates the top one comment produced by our model, and we use the average score as the reported result.

The result is shown in Table 4. It shows that the produced comments of our model are more relevant than those of the baseline models. Besides, our model can produce more correct and proper comments. The scores of both relevance and correctness degrees are also closer to that of the comments made by human. This result indicates that our model is able to produce the relevant comments to the videos by modeling the relevance in different modalities.

4.8 Case Study

To further compare our model with the baselines, we provide an example for case study. This example is talking about a Chinese food called soup dumplings. As illustrated in Figure 4, this example consists of three frames, three surrounding comments and three target comments. Since the audio is not visible, we do not provide the audio part of the video. The surrounding comments are in the first row of the table below the three frames. The second row contains three target comments which are naturally made by human viewers and correspond to a specific time-stamp. We compare the produced

comments of different models with the target comments. It shows that when we select the top one output as the produced comment, both unified transformer and matching transformer can produce a comment relevant to the target comments. However, the produced comments of both fusional RNN and S2S are of low relevance to the video. The output of S2S is talking about eggs and the output of fusional RNN is about dancing, both of which are far away from the video clip. Furthermore, we compare the produced comments between matching transformer and unified transformer. According to the case in Figure 4, it is obvious that the comment from the matching transformer is more relevant to the video clip. The matching transformer can make comments about the soup dumplings that are exactly the key point of the video clip, while the unified transformer can only produce comments about how to process the dirty soup dumplings that fell on the ground, which do not appear in the video. In conclusion, the comments made by our matching transformer are more relevant and correct than that of other baselines.

5 Related Work

Automatic live commenting aims to comment on a video clip based on the surrounding comments, the video clip itself and the corresponding audio clip. This task is similar to the image captioning and video captioning, both of which attract much attention for a long time.

Image Captioning Image captioning involves taking an image, analyzing its visual content, and generating a textual description [5]. [26] try to adopt a retrieval-based model to produce a description of an image from a multimodal space. [27] propose a retrieval approach based on the features extracted by VGG. [23] use a CNN-based model to encode the image and an LSTM to generate the description. [15] try to utilize a merging gate to merge the information in the image and the topics.

Video Captioning Video captioning aims to automatic generate natural language sentences that describe the content of a video [1]. [22] present a CNN-LSTM architecture for generating natural language description of videos. [19] use one LSTM to extract features from video frames and then pass the feature vector through another LSTM for decoding. [24] propose a different neural network architecture based on reinforcement learning for video captioning. [25] release a knowledge-rich video captioning dataset and proposed a new knowledge-aware video description network.

Matching Model A matching model aims to compute the relation between two objects. For text-to-text, [6] adopt an LSTM-based model with cross-attention to predict the relation between two sentences. Inspired by transformer, [28] use self-attention and cross-attention to encode two sentences and model the relation between them. For text-to-image, [11] propose a multi-step reasoning model for visual dialog, which measures the similarity between text-image pairs. [2] present a bottom-up and top-down attention for image captioning and visual question answering. For text-to-audio, [3] try to use transformer learning to learn aligned representations for image, sound and text. [10] propose a framework that learns joint embeddings from a shared lexico-acoustic space for text and audio.

Despite the similarity to image captioning and video captioning, automatic live commenting has its own characteristics. Compared to existing research, it has more diverse contexts including textual

context, visual context and audio context, which is more difficult to tackle. To this end, we propose a multimodal matching transformer model. It can jointly learn the representation of three modalities and the relations among them. Therefore, the proposed model can better integrate information from different angles.

6 Conclusion and Future Work

In this paper, we propose a multimodal matching transformer model for automatic live commenting. It can jointly learn the representation of visual context, audio context, and textual context. In addition, the matching transformer model also explicitly leverages the relations among three modalities to enrich the representation of each one. We evaluate our model on a publicly available live commenting dataset. Experiments show that the proposed multimodal matching transformer can significantly outperforms the state-of-the-art approaches.

For future research, we will further investigate the multimodal interactions among vision, audio, and text in the real-world applications. Moreover, we believe the multimodal pre-training will be a promising direction to explore, where tasks like image captioning and video captioning will benefit from pre-trained models.

Acknowledgments

We would like to thank three anonymous reviewers for many invaluable comments and suggestions to improve our paper. The work of this paper is funded by the project of National key research and development program of China (No. 2017YFB1002102 and 2018YFC0830700).

REFERENCES

- [1] Nayyer Aafaq, Syed Zulqarnain Gilani, Wei Liu, and Ajmal Mian, ‘Video description: a survey of methods, datasets and evaluation metrics’, *arXiv preprint arXiv:1806.00186*, (2018).
- [2] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang, ‘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*, pp. 6077–6086, (2018).
- [3] Yusuf Aytar, Carl Vondrick, and Antonio Torralba, ‘See, hear, and read: Deep aligned representations’, *arXiv preprint arXiv:1706.00932*, (2017).
- [4] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton, ‘Layer normalization’, *arXiv preprint arXiv:1607.06450*, (2016).
- [5] Raffaella Bernardi, Ruket Cakici, Desmond Elliott, Aykut Erdem, Erkut Erdem, Nazli Ikizler-Cinbis, Frank Keller, Adrian Muscat, and Barbara Plank, ‘Automatic description generation from images: A survey of models, datasets, and evaluation measures’, *Journal of Artificial Intelligence Research*, **55**, 409–442, (2016).
- [6] Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen, ‘Enhanced lstm for natural language inference’, in *Proc. ACL*, (2017).
- [7] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio, ‘Empirical evaluation of gated recurrent neural networks on sequence modeling’, *arXiv preprint arXiv:1412.3555*, (2014).
- [8] Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra, ‘Visual dialog’, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 326–335, (2017).
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, ‘Bert: Pre-training of deep bidirectional transformers for language understanding’, *arXiv preprint arXiv:1810.04805*, (2018).

- [10] Benjamin Elizalde, Shuayb Zarar, and Bhiksha Raj, 'Cross modal audio search and retrieval with joint embeddings based on text and audio', in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4095–4099. IEEE, (2019).
- [11] Zhe Gan, Yu Cheng, Ahmed El Kholy, Linjie Li, Jingjing Liu, and Jianfeng Gao, 'Multi-step reasoning via recurrent dual attention for visual dialog', *arXiv preprint arXiv:1902.00579*, (2019).
- [12] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter, 'Audio set: An ontology and human-labeled dataset for audio events', in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 776–780. IEEE, (2017).
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, 'Deep residual learning for image recognition', in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, (2016).
- [14] Diederik P Kingma and Jimmy Ba, 'Adam: A method for stochastic optimization', *arXiv preprint arXiv:1412.6980*, (2014).
- [15] Fenglin Liu, Xuancheng Ren, Yuanxin Liu, Houfeng Wang, and Xu Sun, 'simnet: Stepwise image-topic merging network for generating detailed and comprehensive image captions', *arXiv preprint arXiv:1808.08732*, (2018).
- [16] Shuming Ma, Lei Cui, Damai Dai, Furu Wei, and Xu Sun, 'Livebot: Generating live video comments based on visual and textual contexts', *arXiv preprint arXiv:1809.04938*, (2018).
- [17] Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky, 'The stanford corenlp natural language processing toolkit', in *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, pp. 55–60, (2014).
- [18] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, 'Dropout: a simple way to prevent neural networks from overfitting', *The journal of machine learning research*, **15**(1), 1929–1958, (2014).
- [19] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhudinov, 'Unsupervised learning of video representations using lstms', in *International conference on machine learning*, pp. 843–852, (2015).
- [20] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, 'Attention is all you need', in *Advances in neural information processing systems*, pp. 5998–6008, (2017).
- [21] Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko, 'Sequence to sequence-video to text', in *Proceedings of the IEEE international conference on computer vision*, pp. 4534–4542, (2015).
- [22] Subhashini Venugopalan, Huijuan Xu, Jeff Donahue, Marcus Rohrbach, Raymond Mooney, and Kate Saenko, 'Translating videos to natural language using deep recurrent neural networks', *arXiv preprint arXiv:1412.4729*, (2014).
- [23] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan, 'Show and tell: A neural image caption generator', in *Proc. of CVPR*, pp. 3156–3164, (2015).
- [24] Xin Wang, Wenhui Chen, Jiawei Wu, Yuan-Fang Wang, and William Yang Wang, 'Video captioning via hierarchical reinforcement learning', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4213–4222, (2018).
- [25] Spencer Whitehead, Heng Ji, Mohit Bansal, Shih-Fu Chang, and Clare Voss, 'Incorporating background knowledge into video description generation', in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 3992–4001, (2018).
- [26] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio, 'Show, attend and tell: Neural image caption generation with visual attention', in *International conference on machine learning*, pp. 2048–2057, (2015).
- [27] Semih Yagcioglu, Erkut Erdem, Aykut Erdem, and Ruket Cakici, 'A distributed representation based query expansion approach for image captioning', in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, volume 2, pp. 106–111, (2015).
- [28] Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V Le, 'Qanet: Combining local convolution with global self-attention for reading comprehension', *arXiv preprint arXiv:1804.09541*, (2018).