

Selecting Stickers in Open-Domain Dialogue through Multitask Learning

Zhexin Zhang^{1,2*}, Yeshuang Zhu², Zhengcong Fei², Jinchao Zhang², Jie Zhou²

¹The CoAI group, DCST; ¹Institute for Artificial Intelligence; ¹State Key Lab of Intelligent Technology and Systems; ¹Beijing National Research Center for Information Science and Technology; ¹Tsinghua University, Beijing 100084, China. ²Pattern Recognition Center, WeChat AI, Tencent Inc, China.

zx-zhang18@mails.tsinghua.edu.cn, {yshzhu, dayerzhang, withtomzhou}@tencent.com, feizhengcong@ict.ac.cn,

Abstract

With the increasing popularity of online chatting, stickers are becoming important in our online communication. Selecting appropriate stickers in open-domain dialogue requires a comprehensive understanding of both dialogues and stickers, as well as the relationship between the two types of modalities. To tackle these challenges, we propose a multitask learning method comprised of three auxiliary tasks to enhance the understanding of dialogue history, emotion and semantic meaning of stickers. Extensive experiments conducted on a recent challenging dataset show that our model can better combine the multimodal information and achieve significantly higher accuracy over strong baselines. Ablation study further verifies the effectiveness of each auxiliary task. Our code is available at <https://github.com/nonstopfor/Sticker-Selection>.

1 Introduction

With the development of mobile messaging apps (e.g., WhatsApp and Messenger), visual content is getting more and more frequently used in our daily conversation, such as emojis and stickers. Compared with emojis, stickers are larger images consisting of drawing characters, symbolic icons, and text titles, and are hence more expressive and versatile (Konrad et al., 2020). Users send stickers along with text to show intimacy, express strong emotion, and experience the enjoyment of creativity (Tang and Hew, 2019).

Despite the importance of stickers in daily communication, selecting stickers in open-domain dialogue hasn't been widely explored. In this paper, we address the task of selecting an appropriate sticker from a candidate set for an open-domain multi-turn dialogue. This task is a typical setting for various applications, e.g., automatically recommending stickers in messaging apps and building

*Work done during internship at WeChat AI.



Figure 1: An example of the sticker selection task. Given a dialogue history, the model needs to add a sticker to the last textual message which is the most appropriate one among a collection of candidate stickers (the one marked in the red rectangle). The words below in red denote the emotion or meaning of each sticker.

more interesting and human-like chatbots which could respond with stickers. As shown in Figure 1, this task requires an understanding of dialogue context, emotion and semantic meaning of stickers, and a jointly modeling ability for the multimodal information. Only a few previous works have explored this task (Gao et al., 2020; Wang and Jurgens, 2021). However, existing models are only trained on an end-to-end matching objective and lacks finer-grained supervision signals which could guide models to understand multimodal information better.

Considering the challenges of this task and the shortcomings of previous work, we propose a novel multitask learning method to improve sticker selection in open-domain multi-turn dialogue. We design three auxiliary tasks: **1) masked context prediction**, which uses multimodal context to predict masked tokens in the dialogue history, aiming to understand the dialogue in the presence of the sticker;

2) sticker emotion classification, which utilizes the sticker’s contextualized representation to predict its emotion, aiming to improve the model’s understanding of sticker emotion; **3) sticker semantic prediction**, which explicitly instills semantic understanding of stickers by training the model to reconstruct a sticker’s semantic label based on the multimodal inputs. Moreover, all these tasks help improve our model’s joint modeling capability, as both our model architecture and task design require multimodal inputs and deep interactions between them. We evaluate the performance of our method on a recently proposed and challenging dataset. Extensive experiments show that our multi-task method achieves state-of-the-art performance.

There are two contributions of this paper:

- We propose a multitask learning method to help select appropriate stickers in open-domain multi-turn dialogue.
- Experiment results on a challenging dataset demonstrate the effectiveness of each auxiliary task and combining all the tasks achieves state-of-the-art performance.

2 Related Work

Sticker selection. Previous works proposed to recommend emojis in dialogue systems based on textual or multimodal context (Barbieri et al., 2018; Xie et al., 2016; Barbieri et al., 2017). However, emojis are limited in variety and are much less expressive than stickers. Laddha et al. (2020) retrieved stickers for generated text utterances by simply matching the text tags of stickers. Several works have proposed improved matching methods for stickers. Gao et al. (2020) utilized co-attention to capture the interaction between a sticker and each utterance, and used a fusion network to combine the features. Wang and Jurgens (2021) followed the matching framework of CLIP (Radford et al., 2021) and designed a multimodal encoder for animated GIFs. Fei et al. (2021) proposed to generate special sticker tokens along with text utterances using one single GPT (Wang et al., 2020) for emotion prediction and retrieval of stickers. However, existing models are only trained on an end-to-end matching objective that implicitly guides the models to understand multimodal information. In our work, we design finer-grained auxiliary tasks that instill knowledge of stickers and their contextualized usage in a more efficient way.

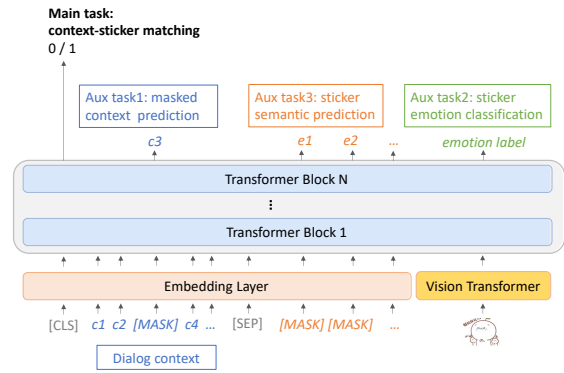


Figure 2: An overview of our training task design. The base model architecture is a multimodal BERT that learns to predict whether the candidate sticker is appropriate given the dialogue context. Three auxiliary tasks are proposed to enhance the model’s ability to understand multimodal input. c_i and e_i represent tokens of dialogue context and semantic label respectively.

Visual Dialogue. Visual dialogue is a task to answer questions about the factual content of the real-world image (Liang et al., 2021; Das et al., 2017a,b). In contrast, selecting appropriate stickers in open-domain dialogue requires understanding sentiment and semantic expression of user-generated, artistic style images.

3 Method

3.1 Task Definition

We assume there is a multi-turn dialogue context $C = \{u_1, \dots, u_N\}$, and a candidate sticker set $S = \{s_1, \dots, s_M\}$, where u_i represents the i -th utterance in the dialogue, and s_i represents the i -th candidate sticker. N is the number of utterances in the dialogue and M is the number of candidate stickers. In this work, we suppose that there is only one appropriate sticker $s^* \in S$, and s^* and u_N belong to the same speaker. The goal is to train a model that can select the right sticker s^* among all candidates S given the dialogue history C .

3.2 Method Overview

An overview of the design of our training tasks is shown in Figure 2. Our main task is to decide whether the candidate sticker is appropriate given the dialogue context. To accomplish this task, we concatenate the embedded dialogue context and the sticker embedding as inputs to BERT. Then we apply a binary classification layer on top of the hidden state of the [CLS] token. In order to enhance the model’s ability to understand the multimodal

input, we design three auxiliary tasks: **1) masked context prediction**, which improves the model’s understanding of dialogue context; **2) sticker emotion classification**, which aims to make the model better understand sticker’s emotion; **3) sticker semantic prediction**, which instills semantic information of stickers to the model. Next, we will introduce our three auxiliary tasks in detail.

3.3 Task 1: Masked Context Prediction

The masked context prediction task follows the masked language modeling (MLM) task in BERT (Devlin et al., 2019). One difference is that we additionally append the embedding of the appropriate sticker to the input embeddings. In this way, the model can learn to utilize stickers for dialogue reconstruction, and thus the interaction between the two modalities is enhanced. The loss for this task is denoted as \mathcal{L}_{ctx} , and takes the same form of cross-entropy loss as in the original MLM task.

3.4 Task 2: Sticker Emotion Classification

In the dataset we used, stickers are annotated with one context-dependent emotion, which means one sticker could have different emotions in different dialogue contexts. Therefore, we design a sticker emotion classification task to enable the model to utilize the text and sticker information simultaneously for understanding sticker emotion. Specifically, we take the hidden state corresponding to the sticker and apply a softmax layer with cross-entropy loss on top of it for emotion classification. The loss for this task is denoted as \mathcal{L}_{emo} .

3.5 Task 3: Sticker Semantic Prediction

Task 1 and Task 2 emphasize learning the implicit meaning of stickers and their correlation with dialogue text. However, many stickers express a clear intention that indicates their proper usage context, e.g., greetings and declines. We believe empowering our model to predict and utilize the semantic meaning of stickers is beneficial for our task. Hence, we further design a semantic label prediction task. We modify our model’s inputs by inserting a fixed-length sequence of [MASK] tokens after the dialogue. The model is trained to recover the label text from the hidden states of the [MASK] tokens. The loss is formulated as the sum of cross-entropy loss for each token in label and is denoted as \mathcal{L}_{sem} . Since the dataset we used has no ground truth semantic labels for stickers, we take the textual information recognized by an OCR tool

as semantic labels for stickers. Note that \mathcal{L}_{sem} is only applied for those stickers with text recognized.

3.6 Total Loss

Besides the above three auxiliary tasks, our main task is a binary classification of whether a candidate sticker is appropriate given the dialogue context. We take all dialogue-sticker pairs in the dataset as positive samples and randomly sample stickers to create an equal number of negative samples. The cross-entropy loss is denoted as \mathcal{L}_{main} .

Our final loss is a combination of the four loss:

$$\mathcal{L} = \mathcal{L}_{main} + \alpha\mathcal{L}_{ctx} + \beta\mathcal{L}_{emo} + \gamma\mathcal{L}_{sem} \quad (1)$$

where α, β, γ are manually tuned hyperparameters.

4 Experiments

4.1 Dataset

We use the Chinese version of the MOD dataset from DSTC10-Track1¹. The dataset is grounded in a dialogue scenario and contains various stickers with contextualized emotion annotation. We split each dialogue into several samples, each containing a text sequence of dialogue history and an accompanying sticker. Note that this dataset is revised from the unpublished one used in Fei et al. (2021).

4.2 Baselines

We compare our model with the following baselines from recent related work: **1) SRS** (Gao et al., 2020), which encodes dialogue history and candidate sticker separately, and then employs a deep interaction network and a fusion network to score each candidate sticker; **2) MOD-GPT** (Fei et al., 2021), which uses one single GPT to generate response text and match sticker; **3) CLIP**, which fine-tunes pretrained CLIP (Radford et al., 2021) for sticker selection using the same contrastive loss.

4.3 Results and Analysis

The result is shown in Table 1. Our full model (MMBERT+ctx+emo+sem) outperforms all baselines on two test sets, and achieves the best performance in almost all settings. As expected, all the results get worse on the hard test set and when selecting one amongst all stickers. As only one

¹<https://openai.weixin.qq.com/dstc/DescriptionEN>. See Appendix A for more details.

²We only provide the results with 10 candidate stickers as their public code does.

	$R_{10}@1$	$R_{10}@2$	$R_{10}@5$	MRR_{10}	$R_{ALL}@1$	$R_{ALL}@2$	$R_{ALL}@5$	MRR_{ALL}
<i>easy test</i>								
SRS ²	30.51	54.24	71.28	48.15	-	-	-	-
MOD-GPT	31.20	54.81	72.13	49.20	5.10	9.05	15.57	11.46
CLIP	38.44	56.45	82.27	56.76	6.00	9.39	16.61	12.69
MMBERT	45.44	66.78	90.95	64.03	5.69	10.08	19.44	13.98
MMBERT+ <i>ctx</i>	47.06	67.34	90.76	65.00	5.91	10.26	20.72	14.44
MMBERT+ <i>ctx+emo</i>	48.80	70.67	92.29	66.88	6.07	11.26	22.02	15.22
MMBERT+ <i>ctx+emo+sem</i>	49.14**	69.46**	91.76**	66.67**	7.40**	12.07**	22.08**	15.99**
<i>hard test</i>								
SRS	23.85	45.30	63.52	40.33	-	-	-	-
MOD-GPT	25.50	49.22	64.03	40.51	3.52	6.12	12.76	9.23
CLIP	32.81	48.55	76.17	51.28	5.79	8.91	15.14	11.55
MMBERT	32.47	50.40	78.32	51.88	3.90	6.62	13.15	9.71
MMBERT+ <i>ctx</i>	33.11	51.04	78.98	52.60	4.21	7.71	13.82	10.38
MMBERT+ <i>ctx+emo</i>	35.39	52.26	78.14	53.65	4.87	8.18	14.66	11.06
MMBERT+ <i>ctx+emo+sem</i>	36.64**	55.48**	80.78**	55.40**	6.06	9.65*	15.79	12.40**

Table 1: Performance of the models on DSTC10 dataset. All the numbers are scaled by 100. The easy test set contains only the same stickers seen during training, while the hard test set has unseen stickers. The footnotes ₁₀ and _{ALL} indicate the numbers of candidate stickers considered for each train and test case, which are 10 (ground truth sticker plus 9 randomly sampled stickers) or all available stickers respectively. $R@k$ is the recall rate of top-k predicted stickers and MRR the Mean Reciprocal Rank of ground truth stickers. The abbreviations *ctx*, *emo* and *sem* correspond to the auxiliary task 1, 2 and 3 respectively in Section 3. A paired t-test is conducted between the full model (MMBERT+*ctx+emo+sem*) and CLIP (*: $p < 0.05$, **: $p < 0.01$).

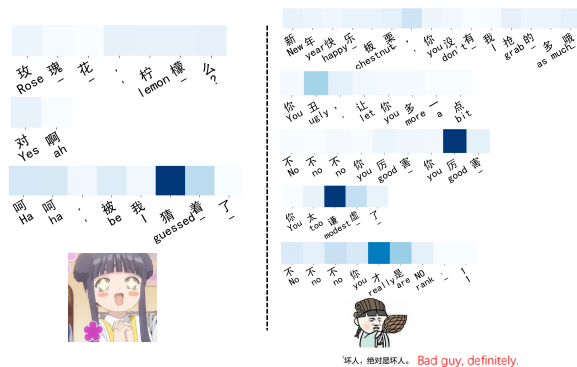


Figure 3: Examples of word saliency in the dialogue history. Word saliency is computed as Frobenius norm of its gradient with regard to the main task loss. Darker color indicates the word is more important. The words in red denote the text in the sticker.

out of the numerous and various online stickers is considered correct, the task is inherently challenging. We find that CLIP is a strong baseline due to its better generalization ability on the hard set, compared with our base model which has no auxiliary task (MMBERT). This may be because CLIP is pretrained on a large number of image-text pairs. However, with multitask learning, our full model outperforms CLIP, although BERT has never seen images during pretraining. Thus, we conclude that our multitask learning method can improve sticker selection by explicitly guiding the model to understand multimodal information.




Figure 4: The diversity of the predicted and ground truth stickers in the easy test set.

We also perform an ablation study to verify the effect of each auxiliary task. A clear trend emerges that the performance improves as each auxiliary training task is added to MMBERT, verifying the efficacy of our task design. One exception is that MMBERT+*ctx+emo* performs slightly better than our full model in terms of $R_{10}@2$, $R_{10}@5$, and MRR_{10} . However, the inconsistency disappears when considering all stickers as candidates. Furthermore, our full model performs significantly better on the hard test set which contains unseen stickers. Hence, we conclude that introducing semantic information improves the model’s generalization ability. We also find that our full model achieves 60% accuracy on the validation set for the auxiliary



Figure 5: A failing case of our model. The leftmost sticker is selected by our model among the four candidate stickers. The appropriate sticker is marked with a red rectangle. The red words explain the stickers’ emotions or meanings.



Figure 6: A case in which our model’s prediction is not the same as the answer but also appropriate. The leftmost sticker is selected by our model among the four candidate stickers. The appropriate sticker is marked with a red rectangle. The red words explain the stickers’ emotions or meanings.

sticker emotion classification task with 52 emotion labels in total, which is reasonable and confirms our model can learn from the auxiliary tasks.

We visualize the saliency of different words in the dialogue history in Figure 3, which shows that the more relevant words (e.g., *guessed*, *good* and *modest*) in the dialogue history contribute more to our model’s prediction. Notably, our model could

attend to some distant words (e.g., *good*), not just the words inside the previous utterance.

We also analyze the prediction diversity of our full model. As shown in Figure 4, the predictions of our model are diverse in general, covering almost all stickers in the whole candidate set. We note that a few stickers are predicted significantly more times than other stickers, which is because they appear much more frequently than other stickers in the training set. We leave addressing the imbalance problem of the training set as our future work.

4.4 Case Study

We present a successful case in Figure 1, where the ground truth sticker has no OCR information, making it challenging for the model to understand its semantic meaning. Moreover, the model needs to understand that the dialogue is in a delighted context, and the stickers’ emotions and meanings in order to distinguish the most appropriate sticker from the others. This case suggests our model has a good understanding of dialogue history and sticker emotion and semantic meaning with the help of auxiliary tasks.

We show a failing case of our model in Figure 5. In this case, the appropriate sticker never appears in the training set. Considering the hard test set is more challenging than the easy test set, improving the generalization ability of our model is thus an important direction of future work. The same is true for baselines.

In the dataset we used, only one sticker is considered correct. However, we observe cases where the model’s selection is not the same as the answer but is also appropriate. An example is shown in Figure 6. Therefore, the results in Table 1 indicate a lower bound performance and our model may perform better in practice.

5 Conclusion

In this paper, we address the challenging task of selecting appropriate stickers in open-domain multi-turn dialogue. We propose a multitask learning method with three auxiliary tasks to enhance the understanding of dialogues and stickers. Experiments show that our model outperforms strong baselines, confirming the effectiveness of our multitask learning method for sticker selection. Although our experiments are conducted on a Chinese dataset, our methods are expected to work for other languages.

References

- Francesco Barbieri, Miguel Ballesteros, Francesco Ronzano, and Horacio Saggion. 2018. Multimodal emoji prediction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 679–686.
- Francesco Barbieri, Miguel Ballesteros, and Horacio Saggion. 2017. Are emojis predictable? In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 105–111.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. 2017a. Visual dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 326–335.
- Abhishek Das, Satwik Kottur, José MF Moura, Stefan Lee, and Dhruv Batra. 2017b. Learning cooperative visual dialog agents with deep reinforcement learning. In *Proceedings of the IEEE international conference on computer vision*, pages 2951–2960.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zhengcong Fei, Zekang Li, Jinchao Zhang, Yang Feng, and Jie Zhou. 2021. Towards expressive communication with internet memes: A new multimodal conversation dataset and benchmark. *arXiv preprint arXiv:2109.01839*.
- Shen Gao, Xiuying Chen, Chang Liu, Li Liu, Dongyan Zhao, and Rui Yan. 2020. Learning to respond with stickers: A framework of unifying multi-modality in multi-turn dialog. In *Proceedings of The Web Conference 2020*, pages 1138–1148.
- Artie Konrad, Susan C Herring, and David Choi. 2020. Sticker and emoji use in facebook messenger: Implications for graphicon change. *Journal of Computer-Mediated Communication*, 25(3):217–235.
- Abhishek Laddha, Mohamed Hanoosh, Debdoot Mukherjee, Parth Patwa, and Ankur Narang. 2020. Understanding chat messages for sticker recommendation in messaging apps. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13156–13163.
- Zujie Liang, Huang Hu, Can Xu, Chongyang Tao, Xiubo Geng, Yining Chen, Fan Liang, and Daxin Jiang. 2021. [Maria: A visual experience powered conversational agent](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5596–5611, Online. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*.
- Nils Reimers and Iryna Gurevych. 2020. Making monolingual sentence embeddings multilingual using knowledge distillation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4512–4525.
- Ying Tang and Khe Foon Hew. 2019. Emoticon, emoji, and sticker use in computer-mediated communication: A review of theories and research findings. *International Journal of Communication*, 13:27.
- Xingyao Wang and David Jurgens. 2021. An animated picture says at least a thousand words: Selecting gif-based replies in multimodal dialog. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3228–3257.
- Yida Wang, Pei Ke, Yinhe Zheng, Kaili Huang, Yong Jiang, Xiaoyan Zhu, and Minlie Huang. 2020. A large-scale chinese short-text conversation dataset. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 91–103. Springer.
- Ruobing Xie, Zhiyuan Liu, Rui Yan, and Maosong Sun. 2016. Neural emoji recommendation in dialogue systems. *arXiv preprint arXiv:1612.04609*.

A Dataset Details

Statistics of the dataset are shown in Table 3. There are 307 stickers in total and 228 out of them have textual information extracted by OCR. For stickers without emotion labels or semantic labels, we simply ignore the emotion classification loss or the semantic prediction loss. A better way to deal with the missing labels is left as future work.

For each dialogue sample, we ignore stickers in the middle of the dialogue history, as we found in preliminary experiments that removing them has no significant impact on the performance.

B Implementation Details

For all the models implemented by ourselves in our experiments, we set the batch size to 8 and use AdamW optimizer with cosine scheduler. For the CLIP baseline, as there is no available CLIP

	$R_{10}@1$	$R_{10}@2$	$R_{10}@5$	MRR_{10}	$R_{ALL}@1$	$R_{ALL}@2$	$R_{ALL}@5$	MRR_{ALL}
<i>easy test</i>								
MMBERT+ctx+emo	48.80	70.67	92.29	66.88	6.07	11.26	22.02	15.22
MMBERT+ctx+emo+sem	49.14	69.46	91.76	66.67	7.40	12.07	22.08	15.99
MMBERT+ctx+emo+sem-OCR	49.95	70.89	92.13	67.39	6.63	12.07	22.18	15.80
MMBERT+ctx+emo+sem-OCR+data	47.06	67.50	91.07	65.12	6.03	9.74	19.75	14.23
<i>hard test</i>								
MMBERT+ctx+emo	35.39	52.26	78.14	53.65	4.87	8.18	14.66	11.06
MMBERT+ctx+emo+sem	36.64	55.48	80.78	55.40	6.06	9.65	15.79	12.40
MMBERT+ctx+emo+sem-OCR	32.87	50.03	76.07	51.51	4.58	7.74	13.70	10.44
MMBERT+ctx+emo+sem-OCR+data	33.42	50.94	78.78	52.49	4.75	7.80	14.31	10.75

Table 2: Effect of incorporating semantic label prediction and OCR feature on DSTC10 dataset. All the numbers are scaled by 100. The easy test set only contains stickers ever seen in the training set, while the hard test set contains stickers unseen during training. $R_{10}@k$ and $R_{ALL}@k$ mean recall rate of ground truth stickers from top-k stickers chosen by the models, given a candidate set of 10 or all available stickers respectively. MRR_{10} and MRR_{ALL} represent Mean Reciprocal Rank of ground truth stickers among 10 or all available stickers. MMBERT+ctx+emo+sem-OCR means not using OCR information for other tasks except sticker semantic prediction. MMBERT+ctx+emo+sem-OCR+data means not using OCR information for other tasks except sticker semantic prediction and adds extra sticker-description pairs for sticker semantic prediction task.

	Train	Valid	Easy test	Hard test
# samples	211575	3542	3215	7028
# emo samples	209890	3495	-	-
# utterances	1666208	26040	25447	59773
# tokens	10400	2718	2780	3818
# stickers	283	249	239	278
Avg. # utterances	7.88	7.35	7.92	8.50
Avg. # tokens	18.42	12.47	12.91	14.54

Table 3: Dataset statistics. Easy test set’s stickers all appear in the train set, while the hard test set contains stickers which don’t appear in the train set. One original dialogue could be split into several samples, each containing one sticker response. The token num is computed by the tokenizer of BERT. # emo samples means the number of samples containing emotion annotation. Avg. # tokens means the average number of tokens for each utterance.

model especially pretrained in Chinese, we use a multilingual version adapted via knowledge distillation (Reimers and Gurevych, 2020)³. We cut the dialogue history to take only the last sentence in order to fit the length limit of CLIP’s text encoder. We also tried to use the last two or more sentences, but found that the performance decreased. All BERT-based models and MOD-GPT use the image encoder in the CLIP baseline. We set the CLIP image encoder’s learning rate to $5e-7$ and the text encoder’s learning rate to $9e-6$. For BERT-based models, we set the learning rate to $9e-6$ and

³https://www.sbert.net/docs/pretrained_models.html#image-text-models

fix the image encoder following (Fei et al., 2021). The maximum epoch is set to 10. For the total loss, α is set to 0.05, β is set to 0.2 and γ is set to 0.1. All the hyperparameters are selected based on the validation set. The maximum training time for one epoch is about 5 hours on one single V100 GPU.

C Effect of Semantic Information

In our full model, we also added semantic labels to other tasks’ inputs, i.e., the main task of context-sticker matching, the masked context prediction task and the sticker emotion classification task. It raises an interesting question of how the performance will change if we remove this information. The result is shown in Table 2. As we can see, the sticker semantic prediction task is more beneficial for the easy test set, while adding OCR information to other tasks is more beneficial for the hard test set. We conjecture that because of the relatively small number of stickers (less than 300), it could be easier for the model to memorize the meaning of all stickers in the dataset, which potentially damages the model’s generalization ability on unseen stickers in the hard test set. Adding OCR information for other tasks greatly alleviates this phenomenon because it could offer semantic labels for unseen stickers and enhance the model’s generalization ability.

We also tried to enhance the model’s generalization ability by incorporating additional sticker-description pairs from another source into the

	$R_{ALL}@1$	$R_{ALL}@2$	$R_{ALL}@5$	MRR_{ALL}
<i>easy test</i>				
CLIP	7.49/2.95/ 4.54	11.23/6.09/ 5.14	18.21/13.23/ 4.98	14.29/9.54/ 4.75
MMBERT+ <i>ctx+emo+sem</i>	8.69/4.85/3.84	12.80/10.75/2.05	19.69/27.12/-7.43	15.89/16.28/-0.39
<i>hard test</i>				
CLIP	6.82/3.58/ 3.24	10.23/6.46/ 3.77	16.60/12.66/ 3.94	12.74/9.26/ 3.48
MMBERT+ <i>ctx+emo+sem</i>	6.95/4.50/2.45	10.68/7.77/2.91	16.48/14.62/1.86	13.24/10.91/2.33

Table 4: Performance of our full model and CLIP on the divided dataset. **The three values divided by / correspond to the performance on the sub dataset where stickers have text, the performance on the sub dataset where stickers don’t have text and their difference.** All the numbers are scaled by 100. The easy test set contains only the same stickers seen during training, while the hard test set has unseen stickers. $R@k$ is the recall rate of top-k predicted stickers and MRR the Mean Reciprocal Rank of ground truth stickers. The abbreviations *ctx*, *emo* and *sem* correspond to the auxiliary task 1, 2 and 3 respectively in Section 3.

sticker semantic prediction task. As Table 2 shows, this method could increase the performance on the hard test set as expected, but the performance on the easy test set drops significantly, which may be attributed to the distribution difference between the additional data and our original data.

D Analysis of Sensitivity to the Text in Stickers

To explore whether stickers have text or not could affect the model’s performance, we split each test set into two parts, i.e., stickers with recognized text labels versus those without text labels. The number of samples in each part is 2164 and 1051 for the easy set, and 4429 and 2599 for the hard set. We compare the performance of our full model with that of CLIP on the divided test sets in Table 4. To avoid randomness in candidate set construction, we only compare the two models with the whole candidate set. In general, our full model and CLIP work better when the stickers have text, which suggests that the text in the sticker could help the model better understand the sticker.⁴ However, our model is less sensitive to whether stickers have text or not according to the smaller difference value compared with CLIP. This implies our model is more robust to different kinds of candidate stickers.

⁴Strictly speaking, there may be other factors that make the two parts inherently different in difficulty, but a fair comparison is difficult to make. Intuitively, stickers without text labels are generally harder to understand for the models.