

MMChat: Multi-Modal Chat Dataset on Social Media

Yinhe Zheng^{♣*}, Guanyi Chen^{♣*}, Xin Liu[♡], Jian Sun^{♠†}

[♠] Alibaba Group, [♣] Utrecht University, [♡] Samsung Research China - Beijing (SRC-B)
zhengyinhe1@163.com, g.chen@uu.nl, jian.sun@alibaba-inc.com

Abstract

Incorporating multi-modal contexts in conversation is important for developing more engaging dialogue systems. In this work, we explore this direction by introducing MMCHAT: a large-scale Chinese multi-modal dialogue corpus (32.4M raw dialogues and 120.84K filtered dialogues). Unlike previous corpora that are crowd-sourced or collected from fictitious movies, MMCHAT contains image-grounded dialogues collected from real conversations on social media, in which the *sparsity* issue is observed. Specifically, image-initiated dialogues in common communications may deviate to some non-image-grounded topics as the conversation proceeds. To better investigate this issue, we manually annotate 100K dialogues from MMCHAT and further filter the corpus accordingly, which yields MMCHAT-HF. We develop a benchmark model to address the sparsity issue in dialogue generation tasks by adapting the attention routing mechanism on image features. Experiments demonstrate the usefulness of incorporating image features and the effectiveness of handling the sparsity of image features.

Keywords: Dialogue Systems, Multi-Modality, Social Media

1. Introduction

The ability to converse like a human is one of the desiderata for building open-domain dialogue systems (Adiwardana et al., 2020). Current attempts to build human-like open-domain dialogue systems generally follow two angles: 1) Enriching the dialogue system with textual or structural contexts such as knowledge (Madotto et al., 2018) or personalities (Zhang et al., 2018a; Zheng et al., 2019); 2) Enabling the dialogue systems to perceive multi-modality contexts beyond text such as vision, voice, or even gesture (Shuster et al., 2020b; Shuster et al., 2020c; Liao et al., 2018; Ju et al., 2019). Systems built following the second angle are also known as Multi-Modal Dialogue Systems (MMDSs).

To facilitate the development of data-driven MMDSs, a few dialogue datasets containing visual information have been constructed (Mostafazadeh et al., 2017; Mogadala et al., 2019; AlAmri et al., 2019; Kottur et al., 2019; Pasunuru and Bansal, 2018; Shuster et al., 2020a; Meng et al., 2020). For instance, Shuster et al. (2020a) introduced a crowd-sourced image grounded dialogue corpus IMAGE-CHAT, in which annotators are employed to chat in accordance with given images. Meng et al. (2020) proposed OPENVIDIAL by directly extracting dialogues and their visual contexts from movies and TV series. There are also works on visual question answering (Das et al., 2017) that focus on the question answering tasks involving image inputs.

A significant drawback of existing datasets is that they postulated every utterance in a dialogue to be grounded on the given image. Nevertheless, this is not always true in our daily communications. Concretely, the topic triggered by the image may drift in the conversation flow so that not every utterance in a dialogue session is



Figure 1: Example dialogues from MMCHAT (translated from Chinese).

image-grounded. Taking the right dialogue in Figure 1 as an example, it is initialized by objects shown in the associated image (i.e., “*paintbox*” and “*dried paint*”) but the focus of the following dialogue move to the speaker’s own experience as a painter, which is not image-related anymore. A similar pattern is also observed in the left example of Figure 1. We refer to this phenomenon as the issue of *sparsity* and dialogues that exhibit this phenomenon as the *sparse image-grounded dialogues*.

To tackle the above issue, we introduce MMCHAT: a large-scale dataset that contains sparse image-grounded dialogues in Chinese. We first collected 32.4M sessions of raw dialogues and 8.41M associated images from social media. Based on these raw dialogues, we design an elaborate data filtering process and construct MMCHAT, which contains 120.84K sessions of filtered high-quality dialogues and 204.32K images. Dialogues that are incoherent or involve offensive content are filtered. Two example dialogue ses-

* Equal Contribution

† Corresponding Author

sions are shown in Figure 1. Unlike previous multi-modal dialogue datasets that only provide a single raw image per dialogue session, each filtered dialogue session in MMCHAT corresponds to one or multiple images. The semantic information of each image is further revealed in our dataset using a pre-trained image caption model. Specifically, a set of detected object labels and a generated descriptive caption are released for each image in MMCHAT.

To further improve the quality of dialogues in MMCHAT, we sample 100K dialogue sessions from MMCHAT, and manually check the quality of images and whether these dialogues are strongly correlated with the associated images. This yields a “human filtered” dataset MMCHAT-HF that contains 19.90K dialogue sessions and 52.24K images.

Building on both MMCHAT and MMCHAT-HF, We provide a strong benchmark model to tackle the image-sparsity issue in open-domain dialogue generation tasks based on the attention routing mechanism (Zheng et al., 2020). Evaluation results on both datasets suggest that incorporating visual contexts contributes positively to dialogue modeling, and the approach used in our benchmark model helps alleviate the sparsity issue. Besides enlightening advanced models for realistic multi-modal conversations, MMCHAT is also built to help understand how Chinese multi-modal communications are conducted from the aspect of social science (Jovanovic and Van Leeuwen, 2018). The vast amount of dialogues and images in MMCHAT can also benefit the study of multi-modal pretraining models.

In what follows, we summarize our main contributions:

- We construct a large multi-modal dialogue dataset MMCHAT, addressing the issue of “sparsity”. A dedicated automatic filtering process is proposed to clean the dataset.
- We offer a human filtered dataset MMCHAT-HF based on 100K dialogue sessions sampled from MMCHAT.
- We build benchmark models on MMCHAT. The results indicate that incorporating visual contexts contribute positively to dialogue modeling, and our benchmark model can better tackle the sparsity issue.

Our dataset and code are available in <https://github.com/silverriver/MMChat>

2. Dataset Construction: MMCHAT

MMCHAT originates from a Chinese social media on which users can share their daily lives through images and texts. This section starts by introducing how MMCHAT is constructed (Section 2.1) and cleaned (Section 2.2). We then sketch how to manually annotate MMCHAT to produce MMCHAT-HF (Section 2.3). At length, we report an analysis of the MMCHAT dataset.

#(Dialogues)	120.84K
#(Total Images)	204.32K
#(Total Utterances)	314.13K
#(Dialogue Sessions) Longer than 4	17.32K
#(Image) per Dialogue	2.91
#(Utterance) per Dialogue	2.59
#(Character) per Utterance	8.52
#(Raw Dialogues)	32.4M

Table 1: Statistics of MMCHAT.

2.1. Data Collection


A two-phase pipeline is used to construct raw dialogues in MMCHAT: the first phase aims to collect seed users who are active on social media. We start this phase with a few hand-collected mass media accounts. Professionals maintain these accounts and are committed to posting daily news on broad topics. The users who comment under this news are collected as our *seed users*. The second phase starts from the seed users collected above. Specifically, the images posted by these seed users are obtained, and the comments under these images are collected. Dialogues along these images are constructed by restoring the reply relationship between these comments.

The two-phase data collection approach used in our study effectively avoids spammers’ noises since most spammers will not bother to follow and reply to daily news. Moreover, we also filter out seed users that are not active to make the data collection process more effective. Finally, we collected a corpus containing about 32.4M sessions of raw dialogues.

2.2. Data Filtering and Post-processing

To improve the quality of MMCHAT, a set of rules is carefully designed to filter out low-quality images and dialogues from the raw corpus collected in Section 2.1. Specifically, images with extremely low resolution (fewer than 500 pixels) or high aspect ratios (larger than 10) are abandoned, and dialogues that contain extremely long utterances (longer than 200 tokens) are filtered. Moreover, we only retain dialogues that contain more than 2 utterances. The offensive contents are also filtered using an offensive word list and a pre-trained offensive content classifier (Wang et al., 2020). To ensure the dialogue contents in MMCHAT are related to the corresponding images in the first few turns of the conversation, we only retain images that are uploaded through the direct-share mode. This mode allows users to share images without providing textual content. We argue that the initial few turns of the dialogues following these image-only posts are usually triggered by the visual information because there are no previous textual contexts except for the uploaded images.


Note that eliminating posts that are not uploaded through the direct-share mode filters out a vast major-



Caption: A store on the street during night.
Detected Objects: building, light, red signature, red calligraphy, plants, flowers, lamps, table



Caption: A brick house with guidepost.
Detected Objects: building, wall, black signature, red signature, calligraphy, lamp



Caption: Flags on the street.
Detected Objects: building, blue sky, red flags, roof, people, trees

- A: 去王家大院看看。
 (Will you have a visit to the Wang Family Courtyard?)
 B: 跟古城差不多吧。
 (Isn't there having the same scene with the ancient city?)
 A: 比古城好，有时间可以去。
 (Much better than the ancient city. You can have a visit if you have spare time.)
 B: 好的 古城过度开发了都成商业街了。
 (Sure. The ancient city has been over-commercialized.)
 A: 嗯嗯，王家大院很震撼我。
 (Uh-huh. The Wang Family Courtyard impressed me a lot.)

Table 2: An example dialogue (translated from Chinese) and its associated images from MMCHAT. Each image’s semantic information (including object labels, attributes, bounding boxes and image captions) is provided.

ity of collected raw dialogues. However, this rule is adopted not because these filtered dialogues are of low quality but because we only have limited computation resources. We want our model to focus on dialogues that are more closely related to its multi-modal contexts. We believe these filtered raw dialogues are useful in building large-scale multi-modal dialogue models or multi-modal pre-trained models. We will release all the collected raw dialogues to facilitate further studies in this direction.

The statistics of the resulting MMCHAT dataset are shown in Table 1. Each dialogue session is associated with at least one image (9 images maximum), and a considerable amount of sessions (more than 17.32K) in MMCHAT contain at least 4 utterances (i.e., 2 turns). Note that different dialogues may share the same post (i.e., the same set of images). To protect data privacy, MMCHAT is released under strict terms for academic users only. More details for the data release protocol can be found in the Broader Consideration section.

2.3. Human Filtering

To facilitate further studies on MMCHAT, we recruit annotators to manually filter MMCHAT, and construct a dataset MMCHAT-HF with higher quality. Concretely, we random sample 100k dialogue sessions from MMCHAT and ask annotators to annotate each session from the following three aspects: 1) *Whether*

#(Dialogues)	19.90K
#(Total Images)	52.66K
#(Total Utterances)	81.06K
#(Dialogue Sessions) Longer than 4	8.91K
#(Image) per Dialogue	2.70
#(Utterance) per Dialogue	4.07
#(Character) per Utterance	11.93

Table 3: Statistics of MMCHAT-HF.

the associated images are qualified. The associated images of a dialogue session are not qualified if any of the images is overlong/flat (i.e., depth-width ratio > 10 or < 0.1) or is a screenshot of texts (e.g., email, news, etc.). We also identify selfies and offensive images as disqualified; 2) *Whether the dialogue contents are non-offensive.* Though we have filtered offensive contents automatically, we ask annotators to check them further manually; 3) *Whether the dialogue contents are strongly correlated with the associated images.* A dialogue session is annotated as “true” in this aspect if its content contains mentions of any object/person/background of its associated images.

In MMCHAT-HF, we only retain dialogue sessions that are annotated as “true” in all the above three aspects. This results in 19.90K dialogue sessions as well as 52.66K images. More statistics are shown in Table 3.

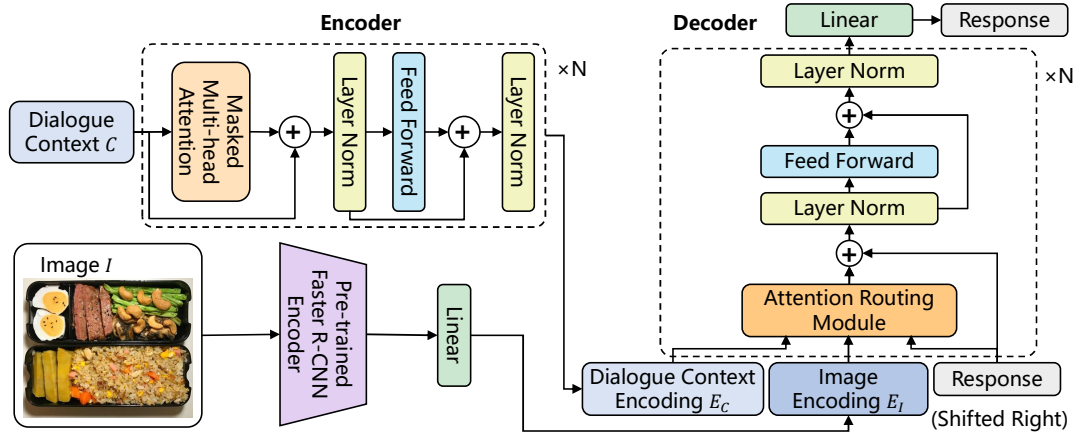


Figure 2: Overview of multi-modal dialogue generation model.

In MMCHAT-HF, dialogues contains 4.07 utterances on average and utterances contains 11.93 characters on average, which are longer than dialogues in MMCHAT (see Table 1). All these annotations will be released under our data release protocol.

2.4. Data Analysis

We try to reveal the semantics of images contained in MMCHAT. Specifically, a Faster R-CNN model (Ren et al., 2015; Lu et al., 2019) trained with attribute head on the Visual Genome (Krishna et al., 2017) dataset is used to detect objects in each image¹. The regions where any class detection confidence exceeds a threshold (0.2 in our case following (Anderson et al., 2018)) are selected to further detect the specific object labels. We follow the work of Anderson et al. (2018) to use an object and attribute vocabulary with the size of 1600 and 400, respectively. An average count of 11.42 objects is detected in each image. This indicates that images in our dataset contain rich semantic information and thus are informative. An example dialogue session, together with its associated images, is shown in Table 2.

3. Dialogue Generation on MMCHAT

This section starts by formally defining the task of MMDS (Section 3.1). Moreover, we propose to use the attention routing mechanism (Zheng et al., 2020) to capture the issue of sparsity (Section 3.2).

3.1. Task Definition

The task of MMDS is to learn a function f that can map textual contexts \mathcal{C} (e.g., dialogue histories) and multi-modal contexts \mathcal{I} (e.g., images, audio or video) into dialogue responses Y , i.e., learn $f : \{\mathcal{C}, \mathcal{I}\} \mapsto Y$. In this study, we focus on the image modality in \mathcal{I} , i.e., \mathcal{I} is composed of a set of images $\{I_n\}_{n=1}^N$.

¹We use the pre-trained model provided by <https://github.com/peteanderson80/bottom-up-attention>

3.2. Dialogue Generation Model

The Seq2Seq architecture is used as our backbone to build a multi-modal dialogue generation model. As shown in Figure 2, two encoders are used to respectively encode the textual context \mathcal{C} and image context \mathcal{I} into encoded representations E_C and E_I . An attention routing module is utilized to merge E_C and E_I in the decoder, and the response Y is decoded autoregressively.

3.2.1. Encoder

The encoder for the textual context \mathcal{C} is parameterized with the Transformer architecture (Vaswani et al., 2017) (12 layers, 12 attention heads, and 768 hidden states). To improve the generation quality, we initialize its weights using a pre-trained GPT model (Radford et al., 2018). Utterances in the dialogue history are concatenated using a special token “[SEP]”, and E_C is obtained by feeding the concatenated token sequence into the textual encoder.

The encoder for the image context \mathcal{I} is implemented as the Faster R-CNN model with ResNet-101 backbone. The weights of this encoder are pre-trained on the Visual Genome dataset and fixed in the training process². Specifically, a feature vector with the size of 2048 is extracted from each image region. The top-50 high confidence regions are used to produce E_I with a linear layer to adjust the feature-length, i.e., the resulting E_I contains 50 features, each has a length of 768.

3.2.2. Decoder

We implement the dialogue decoder with the Transformer architecture and share its weights with our textual encoder. To tackle the sparsity issue, we equip the dialogue decoder with the attention routing mechanism (Zheng et al., 2020) to balance the contribution of each region feature. Specifically, given the encoding of the dialogue context E_C , image context E_I , and previous decoded tokens E_{pre} , three attention routes are

²We have released the pre-trained weights

Model	BLEU-2	BLEU-3	BLEU-4	Dist-1	Dist-2	Ent-1	Ent-2
Seq2Seq	2.830	1.376	0.805	2.63	33.92	6.00	9.47
Seq2Seq+PIMG	2.928	1.469	0.888	2.73	34.34	6.01	9.45
	(+3.46%)	(+6.76%)	(+10.31%)	(+3.80%)	(+1.24%)	(+0.17%)	(-0.21%)
Seq2Seq+IMG	3.001	1.588	1.006	2.82	35.38	6.07	9.52
	(+6.04%)	(+15.41%)	(+24.97%)	(+7.22%)	(+4.30%)	(+1.17%)	(+0.53%)
Human Reference	N/A	N/A	N/A	9.09	48.77	6.69	9.64

Table 4: Evaluation Results on MMCHAT. Relative improvements compared to the Seq2Seq baseline is shown in parentheses.

Model	BLEU-2	BLEU-3	BLEU-4	Dist-1	Dist-2	Ent-1	Ent-2
Seq2Seq	3.779	2.405	1.641	5.35	45.62	6.11	9.26
Seq2Seq+PIMG	4.576	3.094	2.230	5.04	42.61	6.01	9.14
	(+21.09%)	(+28.65%)	(+35.89%)	(-5.79%)	(-6.60%)	(-1.64%)	(-0.21%)
Seq2Seq+IMG	4.818	3.381	2.541	5.75	45.35	6.05	9.15
	(+27.49%)	(+40.58%)	(+54.84%)	(+7.48%)	(-0.59%)	(-0.98%)	(+1.19%)
Human Reference	N/A	N/A	N/A	6.17	47.43	5.98	9.00

Table 5: Evaluation Results on MMCHAT-HF. Relative improvements compared to the Seq2Seq baseline is shown in parentheses.

computed as:

$$O_C = \text{MHA}(E_{\text{pre}}, E_C, E_C), \quad (1)$$

$$O_I = \text{MHA}(E_{\text{pre}}, \gamma E_I, \gamma E_I), \quad (2)$$

$$O_{\text{pre}} = \text{MMHA}(E_{\text{pre}}, E_{\text{pre}}, E_{\text{pre}}), \quad (3)$$

where $\gamma \in [0, 1]$ is a hyper-parameter to re-scale E_I . MHA and MMHA represent masked and unmasked multi-head attention, respectively, in which E_{pre} serves as the query. The results of each attention operation are averaged before proceeding to the next sub-module:

$$O_{\text{merge}} = \frac{O_C + O_I + O_{\text{pre}}}{3}. \quad (4)$$

Note that the attention route on image features (i.e., Eq. 2) assigns different weights to different image regions. This facilitates more flexible control over image features in the decoding process and thus helps ease the sparsity issue.

4. Experiments

Experiments are performed to assess both our model and datasets. Specifically, we train our model on both MMCHAT and MMCHAT-HF. For MMCHAT, we sample 4.0K and 2.0K dialogue sessions for testing and validation, respectively, and for MMCHAT-HF, we sample 1.0K and 1.0K dialogue sessions for testing and validation, respectively. Two baselines are also implemented in our study.

4.1. Implementation Details

In our proposed dialogue model (referred to as **Seq2Seq+IMG**), the encoder and decoder are 12-layer

transformers with 768-dimensional hidden states and 12 attention heads. For the position-wise feed-forward networks, 3,072-dimensional inner states are used. The Adam optimizer is used to train our model with $\beta_1 = 0.9$, $\beta_2 = 0.98$ and $\epsilon = 10^{-9}$. The maximum learning rate is set to $1.0e-4$. The training starts with a warmup step of 1,000, and the learning rate is annealed proportionally to the inverse square root of the step number. The batch size is set to 360, and the training iterates 60 epochs. A character-level vocabulary of size 13,084 is used. Other settings of our Transformer model follow the work of Radford et al. (2018).

We share the weights of the encoder and decoder in the dialogue model and initialize these weights using a pre-trained GPT model (Wang et al., 2020). The pre-training corpus contains about 0.5 billion tokens, and the pre-training process lasts for a week on 8 GTX1080Ti GPUs. We use the top-K ($K = 20$) decoding scheme with beam search for the dialogue model in the inference phase. The beam size and length penalty are 4 and 2.0, respectively. The training of our model lasted for about 24 hours on 2 NVIDIA Tesla V100 GPUs. The number of parameters for our model (textual encoder and decoder, not including the image encoder) is 191.01M. Note that due to the large computation loads to utilize the pre-training based method in our model and our baselines, we inherit most of the hyper-parameter settings from the previous studies, such as Radford et al. (2018), and skip the hyper-parameter tuning process. Moreover, for fair comparisons, we use a fixed set of hyper-parameters in all our experiments (including all the baselines).

Input Image	Dialogue Contexts and Generated Responses
	<p>Context: “在拜托啦学长那里看到你，欢迎来我们福建玩，坐标福建漳州。不过现在在外上学” (“I saw you on the channel ”Please! Seniors.”. Welcome to Fujian! I am from Zhangzhou, Fujian, but studying out of Fujian.”)</p> <p>Human: 厦门。 (Xiamen.)</p> <p>Seq2Seq: 啊，那还真的好啊，我已经回来了。 (Aha! That’s great! I have been back.)</p> <p>Seq2Seq+PIMG: 我是漳州的学妹。 (I am your younger female schoolmate from Zhangzhou.)</p> <p>Seq2Seq+IMG: 原来是这样，下次有机再来福建找你玩哈。 (I see. I’ll look you up, if I ever have a chance to visit Fujian again.)</p>
	<p>Context: “你又睡那么晚”，“美甲！”“你有空了就来撒” (“Why you go to sleep so late again?”, “Just back from a nail salon.”, “Shall we meet if you have spare time?”)</p> <p>Human: 好。 (Okay.)</p> <p>Seq2Seq: 我就在四小时以左右开始睡不着觉！其他时候都是两三点。 (I found it is hard to fall asleep at 4 o’clock this morning. This used to be 2 or 3 o’clock.)</p> <p>Seq2Seq+PIMG: 上完课就去上，没时间了。 (I’ll go there after class. There is no time.)</p> <p>Seq2Seq+IMG: 好滴！国庆之后就能见到了。 (Okay. I suppose to meet you soon after the national day.)</p>

Table 6: Example outputs (with English translations) generated by our multi-modal dialogue systems.

4.2. Baselines

We also implement two baselines to validate our dataset and model: **Seq2Seq**: a transformer-based Seq2Seq model is built with only textual inputs; **Seq2Seq+PIMG**: an image-grounded dialogue model is built with a single pooled image representation. Specifically, a max-pooling operation is applied to E_I , and the pooled vector is added to each representation vector in E_C . The attention route to E_I (i.e., Eq. 2) is not applied. Note that the first baseline does not use image contexts, and the second baseline does not model the sparsity phenomenon.

For fair comparisons, all baselines employ the same architecture, hyper-parameter setting and initialization scheme with our model **Seq2Seq+IMG**.

4.3. Metrics

We use the following metrics: **BLEU** (Papineni et al., 2002) measures the n-gram (n=2,3,4) overlaps between generated and reference responses; **Distinct (Dist)** (Li et al., 2016) measures the proportion of unique n-gram in the generated response (n=1,2); **Entropy (Ent)** (Zhang et al., 2018b) measures how evenly the empirical n-gram (n=1,2) distribution is:

$$\text{Ent} = \frac{1}{\sum_w F(w)} \sum_{w \in V} F(w) \log \frac{F(w)}{\sum_w F(w)}, \quad (5)$$

where V is the set of all n-grams and $F(w)$ is the frequency of n-gram w . Note that both distinct and the entropy measure the diversity of the generated responses.

4.4. Results

Table 4 and Table 5 shows the results on MMCHAT and MMCHAT-HF, respectively. Table 6 lists some example outputs of our dialogue generation models and baselines. Generally speaking, our Seq2Seq+IMG outperforms both baselines on most metrics. The exception is on the entropy metric: the difference between Seq2Seq and Seq2Seq+IMG on both datasets is marginal (approximately 1%).

Based on the above results, we can observe that: 1) Incorporating image contexts in dialogue models helps to produce better responses. Specifically, our model obtains 24.97% and 54.84% relative improvements on the BLEU-4 score compared to the text-only baseline Seq2Seq on MMCHAT and MMCHAT-HF, respectively. Meanwhile, similar improvements are also identified when comparing Seq2Seq+PIMG and Seq2Seq. These results validate our motivation to incorporate multi-modal features in the dialogue generation model and prove that MMCHAT (as well as MMCHAT-HF) can be used to build image-grounded dialogue models. 2) Our model, Seq2Seq+IMG, obtains greater relative improvement on BLEU than Seq2Seq+PIMG. This indicates that explicitly modeling the sparsity phenomenon helps to further improve the dialogue generation performance, and MMCHAT/MMCHAT-HF facilitates the study of such a phenomenon.

Moreover, by comparing results on MMCHAT and MMCHAT-HF, we find that: 1) Models trained on

MMCHAT-HF generally receive higher BLEU and distinct scores than those trained on MMCHAT-HF; 2) Incorporating image information on MMCHAT-HF introduce higher improvement on the BLEU score comparing to the improvement observed on MMCHAT (e.g., BLEU-4 is improved 54.84% on MMCHAT-HF while improved 24.97% on MMCHAT). These results indicate that filtering out low-quality images and dialogues that are irrelevant to their associated images (see Section 2.3) do help build a better dataset.

Note that the diversity improvement of our model Seq2Seq+IMG is not significant compared to the baseline Seq2Seq, particularly on the MMCHAT-HF dataset. This may be because the generated responses are bounded by more contexts (i.e., images).³

5. Conclusion

We introduce MMCHAT, a large-scale multi-modal dialogue corpus that reveals the image-sparsity phenomenon in real conversations. Our dataset contains 120.84K dialogue sessions filtered from 32.4M sessions of raw multi-modal dialogues. Building on 100K dialogues from MMCHAT, we further conduct human filtering, yielding MMCHAT-HF, in which there are 19.9K high-quality multi-modal dialogue sessions. A dialogue model is proposed to tackle the image-sparsity issue utilizing MMCHAT. Experiment results indicate that both MMCHAT and MMCHAT-HF help to develop image-grounded dialogue systems and facilitate further study of the image-sparsity issue. Besides the filtered dialogues in MMCHAT, we will also release all the raw dialogues obtained in the data collection process to facilitate further studies.

6. Broader Considerations

Our dataset MMCHAT originates from a Chinese social media. The dataset collection and release protocols are carefully designed to avoid violating the privacy of each user on that social media. Specifically, each user’s permission setting is strictly respected so that only publicly visible contents are collected. Rules are designed to filter out dialogues that may potentially expose users’ private information, such as phone numbers or emails. Moreover, we will not host these images in MMCHAT on our own server. Only the URLs to these images will be released along with the download scripts.

To further enforce the data privacy, MMCHAT is released under strict terms for academic uses only, in which they promise no abuse of MMCHAT besides academic purposes.

In addition to the privacy issues, there might also be toxic or biased texts in MMCHAT or be generated by MMDSs trained on MMCHAT. Although we take the

³Also note that models receive higher distinct scores on MMCHAT-HF than MMCHAT, which could, to a large extent, because the test set of MMCHAT-HF is smaller than that of MMCHAT.

responsibility to remove toxic texts (using an offensive word list, an offensive content classifier, and human filtering), we cannot guarantee that there are no offensive contents left. However, as offensive and abusive content recognition is a rapidly developing area (Vidgen et al., 2019), we would deploy more advanced filters once the new state-of-the-art offensive and abusive classifiers are proposed in the future.

Regarding the potential biases, except those from the dataset itself (Henderson et al., 2018) (which always exists in dialogue datasets), biases might be introduced by the pre-trained language model (Bender et al., 2021) and the pre-trained image encoder (Steed and Caliskan, 2020) used in this work. In the future, we head to apply and develop corresponding mitigation techniques (following works such as Dinan et al. (2020) and Liu et al. (2020)).

During annotation, we pay each annotator 0.5 CNY per item. This results in approximately 60 CNY per hour, which is 4.5 times the minimum wage standard in China.

Besides, we also note that the goal of our work is to facilitate further work on multi-modal dialogue systems. Although the model used in this work is still far from realistic, our dataset can be regarded as an initial step toward the sparsity issue in real-world multi-modal conversations.

7. Bibliographical References

- Adiwardana, D., Luong, M.-T., So, D. R., Hall, J., Fiedel, N., Thoppilan, R., Yang, Z., Kulshreshtha, A., Nemade, G., Lu, Y., et al. (2020). Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977*.
- AlAmri, H., Cartillier, V., Das, A., Wang, J., Cherian, A., Essa, I., Batra, D., Marks, T. K., Hori, C., Anderson, P., Lee, S., and Parikh, D. (2019). Audio visual scene-aware dialog. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 7558–7567. Computer Vision Foundation / IEEE.
- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., and Zhang, L. (2018). Bottom-up and top-down attention for image captioning and visual question answering. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6077–6086. IEEE Computer Society.
- Bender, E. M., Gebru, T., McMillan-Major, A., and Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big. *Proceedings of FAccT*.
- Das, A., Kottur, S., Gupta, K., Singh, A., Yadav, D., Moura, J. M. F., Parikh, D., and Batra, D. (2017). Visual dialog. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 1080–1089. IEEE Computer Society.

- Dinan, E., Fan, A., Williams, A., Urbanek, J., Kiela, D., and Weston, J. (2020). Queens are powerful too: Mitigating gender bias in dialogue generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8173–8188, Online, November. Association for Computational Linguistics.
- Henderson, P., Sinha, K., Angelard-Gontier, N., Ke, N. R., Fried, G., Lowe, R., and Pineau, J. (2018). Ethical challenges in data-driven dialogue systems. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 123–129.
- Jovanovic, D. and Van Leeuwen, T. (2018). Multimodal dialogue on social media. *Social Semiotics*, 28(5):683–699.
- Ju, D., Shuster, K., Boureau, Y.-L., and Weston, J. (2019). All-in-one image-grounded conversational agents. *arXiv preprint arXiv:1912.12394*.
- Kottur, S., Moura, J. M. F., Parikh, D., Batra, D., and Rohrbach, M. (2019). CLEVR-dialog: A diagnostic dataset for multi-round reasoning in visual dialog. 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 582–595, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.-J., Shamma, D. A., et al. (2017). Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73.
- Li, J., Galley, M., Brockett, C., Gao, J., and Dolan, B. (2016). A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California, June. Association for Computational Linguistics.
- Liao, L., Ma, Y., He, X., Hong, R., and Chua, T.-s. (2018). Knowledge-aware multimodal dialogue systems. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 801–809.
- Liu, H., Wang, W., Wang, Y., Liu, H., Liu, Z., and Tang, J. (2020). Mitigating gender bias for neural dialogue generation with adversarial learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 893–903, Online, November. Association for Computational Linguistics.
- Lu, J., Batra, D., Parikh, D., and Lee, S. (2019). Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In Hanna M. Wallach, et al., editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13–23.
- Madotto, A., Wu, C.-S., and Fung, P. (2018). Mem2Seq: Effectively incorporating knowledge bases into end-to-end task-oriented dialog systems. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1468–1478, Melbourne, Australia, July. Association for Computational Linguistics.
- Meng, Y., Wang, S., Han, Q., Sun, X., Wu, F., Yan, R., and Li, J. (2020). Openvidual: A large-scale open-domain dialogue dataset with visual contexts. *arXiv preprint arXiv:2012.15015*.
- Mogadala, A., Kalimuthu, M., and Klakow, D. (2019). Trends in integration of vision and language research: A survey of tasks, datasets, and methods. *arXiv preprint arXiv:1907.09358*.
- Mostafazadeh, N., Brockett, C., Dolan, B., Galley, M., Gao, J., Spithourakis, G., and Vanderwende, L. (2017). Image-grounded conversations: Multimodal context for natural question and response generation. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 462–472, Taipei, Taiwan, November. Asian Federation of Natural Language Processing.
- Papineni, K., Roukos, S., Ward, T., and 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*, pages 311–318, Philadelphia, Pennsylvania, USA, July. Association for Computational Linguistics.
- Pasunuru, R. and Bansal, M. (2018). Game-based video-context dialogue. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 125–136, Brussels, Belgium, October-November. Association for Computational Linguistics.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. (2018). Improving language understanding by generative pre-training.
- Ren, S., He, K., Girshick, R. B., and Sun, J. (2015). Faster R-CNN: towards real-time object detection with region proposal networks. In Corinna Cortes, et al., editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 91–99.
- Shuster, K., Humeau, S., Bordes, A., and Weston, J. (2020a). Image-chat: Engaging grounded conversations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2414–2429, Online, July. Association for Computational Linguistics.
- Shuster, K., Ju, D., Roller, S., Dinan, E., Boureau, Y.-L., and Weston, J. (2020b). The dialogue do-

- decaathlon: Open-domain knowledge and image grounded conversational agents. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2453–2470, Online, July. Association for Computational Linguistics.
- Shuster, K., Smith, E. M., Ju, D., and Weston, J. (2020c). Multi-modal open-domain dialogue. *arXiv preprint arXiv:2010.01082*.
- Steed, R. and Caliskan, A. (2020). Image representations learned with unsupervised pre-training contain human-like biases. *arXiv preprint arXiv:2010.15052*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. In Isabelle Guyon, et al., editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Vidgen, B., Harris, A., Nguyen, D., Tromble, R., Hale, S., and Margetts, H. (2019). Challenges and frontiers in abusive content detection. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 80–93, Florence, Italy, August. Association for Computational Linguistics.
- Wang, Y., Ke, P., Zheng, Y., Huang, K., Jiang, Y., Zhu, X., and Huang, M. (2020). A large-scale chinese short-text conversation dataset. In *Natural Language Processing and Chinese Computing*.
- Zhang, S., Dinan, E., Urbanek, J., Szlam, A., Kiela, D., and Weston, J. (2018a). Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia, July. Association for Computational Linguistics.
- Zhang, Y., Galley, M., Gao, J., Gan, Z., Li, X., Brockett, C., and Dolan, B. (2018b). Generating informative and diverse conversational responses via adversarial information maximization. In Samy Bengio, et al., editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 1815–1825.
- Zheng, Y., Chen, G., Huang, M., Liu, S., and Zhu, X. (2019). Personalized dialogue generation with diversified traits. *arXiv preprint arXiv:1901.09672*.
- Zheng, Y., Zhang, R., Huang, M., and Mao, X. (2020). A pre-training based personalized dialogue generation model with persona-sparse data. In *AAAI*, pages 9693–9700.