

# TopicDiff: A Topic-enriched Diffusion Approach for Multimodal Conversational Emotion Detection

Jiamin Luo\*, Jingjing Wang\*, Guodong Zhou†

School of Computer Science and Technology, Soochow University, China  
No.1, Shizi Street, Suzhou City, Jiangsu Province, China  
20204027003@stu.suda.edu.cn, {djingwang, gdzhou, }@suda.edu.cn

## Abstract

Multimodal Conversational Emotion (MCE) detection, generally spanning across the acoustic, vision and language modalities, has attracted increasing interest in the multimedia community. Previous studies predominantly focus on learning contextual information in conversations with only a few considering the topic information in single language modality, while always neglecting the acoustic and vision topic information. On this basis, we propose a model-agnostic **Topic-enriched Diffusion** (TopicDiff) approach for capturing multimodal topic information in MCE tasks. Particularly, we integrate the diffusion model into neural topic model to alleviate the diversity deficiency problem of neural topic model in capturing topic information. Detailed evaluations demonstrate the significant improvements of TopicDiff over the state-of-the-art MCE baselines, justifying the importance of multimodal topic information to MCE and the effectiveness of TopicDiff in capturing such information. Furthermore, we observe an interesting finding that the topic information in acoustic and vision is more discriminative and robust compared to the language.

**Keywords:** Multimodal Conversational Emotion Detection, Multimodal Topic, Diffusion Model

## 1. Introduction

Multimodal Conversational Emotion (MCE) detection constitutes a crucial task in the fields of natural language processing (NLP) (Poria et al., 2017; Majumder et al., 2019; Ghosal et al., 2019) and multimodal processing (Hazarika et al., 2018a; Hu et al., 2021b), with the objective of automatically detecting emotions in each utterance within a conversation. As exemplified in Figure 1, an utterance (including acoustic spectrum, video frame and language “Guess what, I got an audition”) generated by a speaker could be labeled with emotion *joy*. Over the past decade, MCE has a significant impact on the development of empathetic systems and various potential applications, including opinion mining, healthcare and intelligent assistants (Jiao et al., 2019; Cao et al., 2019). Existing studies have primarily focused on employing RNN-variants (e.g., LSTM or GRU) (Poria et al., 2017; Majumder et al., 2019), graph-based (Ghosal et al., 2019) and transformer-based (Shen et al., 2021) models to capture the contextual information of individual utterances during a conversation.

Recently, a few studies realize the importance of topic information in conversation tasks (Wang et al., 2020a; Zhu et al., 2021). They assert that the utilization of topic information in conversational contexts contributes to fully mining the global clues of utterances during the whole conversation. For instance, conversations regarding the *funeral* topic tend to elicit *sadness*, while those about the *wedding* topic

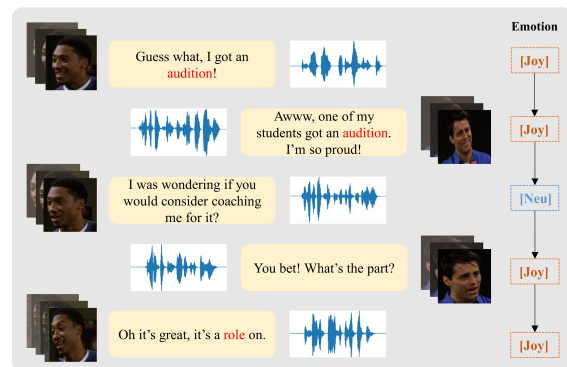


Figure 1: A multimodal conversational example from MELD dataset to illustrate the importance of multimodal topic information, where each utterance contains acoustic spectrum, video frame, language and corresponding emotion label.

tend to elicit *happiness*. Therefore, considering the topic information in a conversation could improve our comprehension of content, context and intent within the conversation. However, existing studies on incorporating the topic information into conversation tasks solely rely on single language modality, without taking the topic information present in both acoustic and vision into account.

In this paper, we argue that the multimodal topic information is crucial for boosting the performance of emotion detection in multimodal conversation. By analyzing various modal topic information, such as tones, facial expressions, and body postures, we can better detect the emotions. As illustrated in Figure 1, we observe a multimodal conversation between two speakers, where the language clues “get-

\*Equal Contribution.

†Corresponding Author: Guodong Zhou.

ting an audition” and “role”, indicating the topic related to *good job opportunities* that could potentially aid in detecting the emotions of speakers. Moreover, the acoustic spectrums with *excited* tones and video frames with *smile* facial expression generated by two speakers indicate a positive (i.e., *joy*) emotion throughout the conversation, which are also useful to discriminate the emotions. Besides, as previously mentioned conversations related to *funeral* and *wedding* topics, the tones (frequency and intensity variations in tones) and facial expressions (smiling or crying face in video frames) can precisely detect the emotions in the whole conversations. Therefore, we believe that a well-behaved approach should consider topic information within multiple modalities in multimodal conversation to enable the precise detection of emotions.

However, traditional neural topic models (Wang et al., 2020b; Jin et al., 2021) generally adopt the architecture of variational auto-encoder (VAE) (Kingma and Welling, 2014) to capture the topic information inside language, which may suffer from the problem of semantic diversity deficiency (Gao et al., 2019). More seriously, acoustic and vision semantics are often more sparse compared to language (He et al., 2022). This may further exacerbate the problem of diversity deficiency in topic information. With this in mind, traditional neural topic models are potentially not well-suited for capturing multimodal topic information. Recently, the diffusion model (Song and Ermon, 2019) has attracted increasing attention due to their ability to sufficiently capture the diverse information from multimodal semantic spaces, achieving remarkable performance in various multimodal tasks (Ruan et al., 2023; Huang et al., 2023). Therefore, we believe that a better-behaved topic mining approach should consider integrating the diffusion model to capture the multimodal topic information.

To tackle the aforementioned challenges, this paper proposes a new model-agnostic **Topic-enriched Diffusion** (TopicDiff) approach to capture the multimodal topic information for addressing MCE. Specifically, we first integrate the diffusion model into neural topic model to alleviate the problem of semantic diversity deficiency. Then, we leverage three TopicDiff modules to capture the topic information inside acoustic, vision, and language modalities, respectively. Finally, we jointly train the basic MCE approaches alongside three TopicDiff modules under the architecture of multi-task learning. Detailed evaluations demonstrate that our TopicDiff approach achieves significant performance improvements compared to the state-of-the-art MCE baselines. Overall, the main contributions of this paper are summarized as follows:

- We are the first to consider the multimodal

topic information for boosting the performance of conversational emotion detection.

- We propose a TopicDiff approach to capture the multimodal topic information in MCE, which first attempts to integrate the diffusion model into neural topic models for alleviating the diversity deficiency problem of them in capturing the topic information. Since TopicDiff is model-agnostic, it can be easily extended to MCE approaches.
- We conduct detailed experiments on our topic-density M3ED\* and two public topic-sparsity MELD, IEMOCAP datasets, and the results on these datasets demonstrate that TopicDiff achieves significant improvements compared to state-of-the-art MCE baselines. This fully justifies the importance of multimodal topic information and the effectiveness of our TopicDiff approach in capturing such information.
- We observe an interesting finding that the topic information in acoustic and vision is more discriminative and robust than that in language.

## 2. Related Work

### Multimodal Conversational Emotion Detection.

Previous studies in MCE focus on capturing contextual information due to the unique structure of conversations, which can be broadly categorized into RNN-, graph-, and knowledge-based approaches. RNN-based approaches leverage RNN variant networks to model contextual information (Poria et al., 2017; Hazarika et al., 2018b) or track emotional states (Hazarika et al., 2018a; Majumder et al., 2019) in conversations. Graph-based approaches leverage graph networks to model the interaction of speakers (Ghosal et al., 2019; Ishiwatari et al., 2020), model multimodal dependencies and speaker dependencies (Hu et al., 2021b; Joshi et al., 2022; Lian et al., 2023), and understand the conversational contexts (Hu et al., 2021a, 2022). Knowledge-based approaches leverage external knowledge to guide contextual modeling, such as commonsense knowledge (Ghosal et al., 2020), psychological knowledge (Li et al., 2021), knowledge from pre-trained language models (Shen et al., 2021; Zhu et al., 2021), while these approaches only consider language modality.

In summary, all the above studies ignore the multimodal topic information, which can be leveraged to boost the performance of emotion detection in multimodal conversation.

**Neural Topic Model.** Based on the architecture of variational auto-encoder (VAE) (Kingma and Welling, 2014), Miao et al. (2016) propose neural topic models to capture the topic information within languages. Then, some improved studies attempt to employ various strategies to address problems

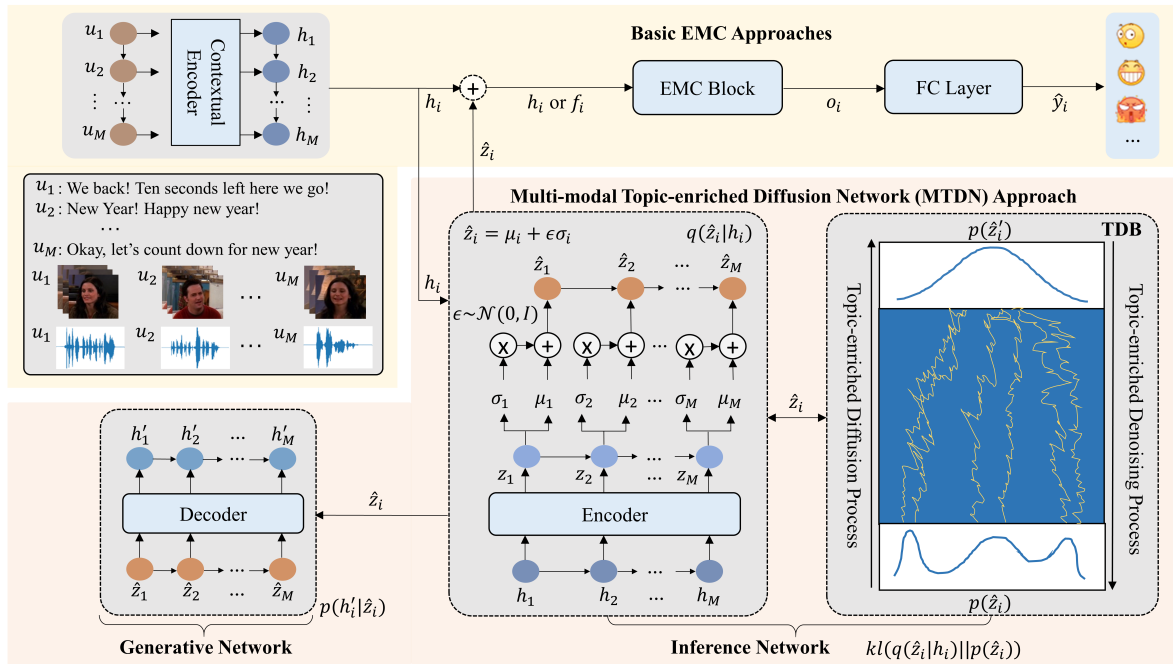


Figure 2: The overall architecture of our model-agnostic Topic-enriched Diffusion (TopicDiff) approach for MCE, where TDB represents Topic-enriched Diffusion Block consisting of Topic-enriched Diffusion Process and Topic-enriched Denoising Process.

in language topic modeling, such as adversarial training (Wang et al., 2020b), pre-trained semantic embeddings (Jin et al., 2021), reinforcement learning (Gui et al., 2019) and multi-task learning (Wang et al., 2020a). Besides, a few studies (Pu et al., 2020; An et al., 2020) consider incorporating the multimodal topic information, while these efforts are not designed to conversation scenarios and hence are not appropriate for our MCE tasks.

Different from all above studies, we focus on multimodal conversation scenarios and design an TopicDiff approach to mine multimodal topic information, which could mitigate the diversity deficiency problem of traditional neural topic models in capturing topic information.

**Diffusion Model.** Recently, diffusion models are the state-of-the-art latent variable models in generative modeling (Yang et al., 2022), which have two main paradigms, i.e., Denoising Diffusion Probabilistic Models (DDPM) (Ho et al., 2020) and Score-based Generative Models (SGM) (Song and Ermon, 2019). At present, there are many studies on diffusion models (Vahdat et al., 2021; Baranchuk et al., 2022; Rombach et al., 2022), of which Vahdat et al. (2021) train SGMs with VAE to explore the effects of the diffusion model on image generation. Ruan et al. (2023) and Huang et al. (2023) leverage the diffusion model to address the multimodal tasks and achieve promising results.

Overall, the diffusion model has shown a remarkable ability in capturing diverse information from latent spaces. Inspired by this, we first attempt to

take advantage of the diffusion model and integrate it into neural topic model to mitigate the diversity deficiency problem, thereby sufficiently capturing the multimodal topic information in MCE.

### 3. Approach

In this section, we formulate the MCE task as follows. A conversation can be defined as a sequence of utterances  $\mathbf{u} = \{u_1, u_2, \dots, u_M\}$ , where  $M$  is the number of utterances. Each utterance involves three sources of utterance-aligned data corresponding to acoustic ( $a$ ), vision ( $v$ ) and language ( $l$ ), which can be formulated as:  $\mathbf{u}_i = \{u_i^a, u_i^v, u_i^l\}$ , where  $u_i^a$ ,  $u_i^v$  and  $u_i^l$  denote the raw feature representation of  $u_i$  from acoustic, vision and language, respectively. The MCE task aims to predict the emotional status label for each utterance  $u_i$  in the conversation based on the available information from all three modalities.

In this paper, we propose a model-agnostic **Topic-enriched Diffusion** (TopicDiff) approach for MCE. Figure 2 shows the overall architecture of TopicDiff, consisting of an inference network with a topic-enriched diffusion block (TDB), and a generative network. Before introducing TopicDiff, we first give an overview of the basic MCE approaches.

#### 3.1. Basic MCE Approaches

The basic MCE approaches mainly focus on leveraging RNN-variants (e.g., GRU), graph-based or

transformer-based models to model contextual information of each conversation, and the general workflow of them is described as follows. Firstly, each utterance  $u_i$  is encoded by different modal contextual encoders to obtain context-aware representation  $h_i$ . Various MCE approaches employ different contextual encoders for each modality, which we denote as CTEncoder for uniformity. Thus, the context-aware representation  $h_i$  of each utterance  $u_i$  is formulated as follows:

$$h_i = \text{CTEncoder}(u_i) \quad (1)$$

where  $h_i = \{h_i^a, h_i^v, h_i^l\}$  is serving as the uniform representation of acoustic, vision, and language

Subsequently,  $h_i$  is fed into various MCE blocks (e.g., MMGCN employs a spectral domain graph convolutional network to encode the multimodal contextual representations), where  $h_i$  is transformed into an intermediate representation  $o_i$ . For clarity, we use ‘‘MCE Block’’ to represent the core and complex components of various MCE approaches as shown in Figure 2.

Finally, the emotion prediction  $\hat{y}_i$  is made by passing  $o_i$  through a fully connected (FC) layer. The basic MCE approaches are trained via the cross-entropy loss function, which can be defined as:

$$\mathcal{L}_{mce} = -\frac{1}{\sum_{n=1}^N c(n)} \sum_{j=1}^N \sum_{i=1}^{c(n)} y_{j,i}^n \log \hat{y}_{j,i}^n \quad (2)$$

where  $N$  is the total number of conversations in the training set.  $c(n)$  is the number of utterances in conversation  $n$ .  $y_{j,i}^n$  and  $\hat{y}_{j,i}^n$  are the expected class label and the predicted emotion label of utterance  $i$  in conversation  $j$ , respectively.

### 3.2. Topic-enriched Diffusion Approach

To capture the topic information inside acoustic, vision, and language, we propose a Topic-enriched Diffusion (TopicDiff) approach. TopicDiff is essentially a variant of neural topic model (Miao et al., 2016), and thus also consists of two main components, i.e., inference network and generative network as illustrated in Figure 2. Particularly, we design a **Topic-enriched Diffusion Block (TDB)**, which is then integrated into the inference network for alleviating the diversity deficiency problem of traditional vae-based neural topic model. Specifically, we leverage three TopicDiff modules to capture the topic information inside acoustic, vision, and language on the basis of basic MCE approaches, respectively.

**Inference Network** is leveraged to map the context-aware representation  $h_i$  in Eq.(1) to a low-dimension latent topic representation  $\hat{z}_i$ .<sup>1</sup> Specifically, an MLP encoder layer is first employed to

extract features from  $h_i$  to obtain the output  $z_i$ . And then  $z_i$  can be fed into two different fully connected layers  $f_\mu$  and  $f_\sigma$  to estimate the mean  $\mu_i$  and standard deviation  $\sigma_i$  of a Gaussian distribution, respectively, denoted as  $\mu_i = f_\mu(z_i)$  and  $\sigma_i = \log f_\sigma(z_i)$ . Finally, we sample  $\hat{z}_i$  from the posterior topic distribution  $q(\hat{z}_i|h_i)$  by using a reparameterization trick as described in Kingma and Welling (2014), i.e.,  $\hat{z}_i = \mu_i + \epsilon\sigma_i$ . Here,  $\epsilon$  is sampled from  $\mathcal{N}(0, \mathbf{I})$ .

Traditional vae-based neural topic models directly feed latent topic representation  $\hat{z}_i$  into generative network to reconstruct context-aware representation  $h_i$ , which may suffer from the problem of diversity deficiency (Gao et al., 2019) in capturing multimodal topic information. Therefore, we design a topic-enriched diffusion block (TDB) integrated into the inference network, consisting of topic-enriched diffusion process and topic-enriched denoising process as illustrated in Figure 2. In the following, we will formulate the TDB integrated into the inference network of TopicDiff in detail.

• **Topic-enriched Diffusion Process** is leveraged to perturb the latent topic representation  $\hat{z}_i$  with an infinite number of noise scales. Indexed by a continuous time variable  $t \in [0, T]$ ,  $\hat{z}_i \sim p(\hat{z}_i)$  and  $\hat{z}'_i \sim p(\hat{z}'_i)$ , where  $p(\hat{z}_i)$  and  $p(\hat{z}'_i)$  represent the latent topic distribution and prior distribution in TDB, respectively. Therefore, we have a tractable form to generate representations efficiently, and the topic-enriched diffusion process can be modeled as the solution to a stochastic differential equation (SDE), denoted as:

$$d\hat{z}_i = f(\hat{z}_i, t)dt + g(t)d\mathbf{w} \quad (3)$$

where  $\mathbf{w}$  is the standard Wiener process<sup>2</sup>.  $f(\hat{z}_i, t)$  and  $g(t)$  represent the vector-valued and scalar function called the drift and diffusion coefficient of  $\hat{z}_i$ , respectively. The SDE has a unique strong solution as long as the coefficients are globally Lipschitz<sup>3</sup> both in state and time.

• **Topic-enriched Denoising Process** is leveraged to synthesize the latent topic representations  $\hat{z}_i$  from the prior distribution  $p(\hat{z}'_i)$ , which also can be seen as a diffusion process, running backwards in time and given by the reverse-time SDE, formulated as:

$$d\hat{z}_i = [f(\hat{z}_i, t) - g(t)^2 \nabla_{\hat{z}_i} \log p(\hat{z}_i)]dt + g(t)d\hat{\mathbf{w}} \quad (4)$$

where  $\hat{\mathbf{w}}$  is the standard Wiener process when time flows backwards from  $T$  to 0, and  $dt$  is an infinitesimal negative time step. Once the score of each

we incorporate latent topic representation  $\hat{z}_i$  via TopicDiff, without modifying the overall architecture of MCE approaches. Therefore, we say that our TopicDiff approach is model-agnostic.

<sup>2</sup>[https://en.wikipedia.org/wiki/Wiener\\_process](https://en.wikipedia.org/wiki/Wiener_process)

<sup>3</sup>[https://en.wikipedia.org/wiki/Lipschitz\\_continuity](https://en.wikipedia.org/wiki/Lipschitz_continuity)

<sup>1</sup>After obtaining the context-aware representation  $h_i$ ,

marginal distribution,  $\nabla_{\hat{z}_i} \log p(\hat{z}_i)$ , is known for all  $t$ , we can derive the denoising process from Eq.(4) and simulate it to sample from  $p(\hat{z}_i)$ .

To estimate  $\nabla_{\hat{z}_i} \log p(\hat{z}_i)$ , we can train a time-dependent score-based model  $s_\theta(\hat{z}_i, t)$  via a continuous generalization, formulated as:  $\theta^* = \arg \min \mathbb{E}_{\hat{z}_i} \mathbb{E}_{\hat{z}'_i | \hat{z}_i} \left\| s_\theta(\hat{z}_i, t) - \nabla_{\hat{z}_i} \log p(\hat{z}_i)(\hat{z}'_i | \hat{z}_i) \right\|_2^2$  where  $t$  is uniformly sampled over  $[0, T]$ . With sufficient data and model capacity, score matching ensures that the optimal solution for this formulation, denoted as  $s_{\theta^*}(\hat{z}_i, t)$ , which equals  $\nabla_{\hat{z}_i} \log p(\hat{z}_i)$  for almost all  $\hat{z}_i$  and  $t$ .

**Generative Network** is leveraged to reconstruct the context-aware representation  $\mathbf{h}_i$  from the latent topic representation  $\hat{z}_i$  at each time step. Similar to inference network, we employ an MLP decoder to take the sampled latent topic representation  $\hat{z}_i$  as input and generate reconstructed  $\mathbf{h}'_i$  for the context-aware representation  $\mathbf{h}_i$ . Besides, generative network also uses Gaussian distributions for both generative prior and variational distribution, while it defines an independent diagonal Gaussian distribution over the embedding space of different modalities. Therefore, generative network simultaneously learns to approximate the conditional probability distribution  $p(\mathbf{h}'_i | \hat{z}_i)$ , which represents the likelihood of generating the reconstructed context-aware representation  $\mathbf{h}'_i$  given the latent topic representation  $\hat{z}_i$ .

### 3.3. Topic-enriched Learning

After obtaining the latent topic representation  $\hat{z}_i$  of each modality, we compute the output representation  $\mathbf{f}_i$  via concatenating  $\hat{z}_i$  and context-aware representation  $\mathbf{h}_i$  as illustrated in Figure 2, i.e.,  $\mathbf{f}_i = \mathbf{h}_i \oplus \hat{z}_i$ , which is used for final emotion detection following basic MCE approaches.

To train three TopicDiff modules, we minimize the variational upper bound on negative data log-likelihood, consisting of two components: reconstruction loss and Kullback-Leibler (KL) loss. The reconstruction loss is used to minimize the reconstruction error between the context-aware representation  $\mathbf{h}_i$  and the reconstructed  $\mathbf{h}'_i$ , thereby improving the quality of the reconstructed context-aware representation  $\mathbf{h}'_i$ , formulated as follows:

$$\mathcal{L}_{rec} = \mathbb{E}_{q(\hat{z}_i | \mathbf{h}_i)} \left[ \log p(\mathbf{h}'_i | \hat{z}_i) \right] \quad (5)$$

And the KL loss is used to minimize the degree of difference between the posterior topic distribution  $q(\hat{z}_i | \mathbf{h}_i)$  and the latent topic distribution target distribution  $p(\hat{z}_i)$ , which can be formulated as follows:

$$\mathcal{L}_{kl} = kl(q(\hat{z}_i | \mathbf{h}_i) || p(\hat{z}_i)) \quad (6)$$

Therefore, we jointly train the basic MCE approaches alongside three TopicDiff modules, and

Dataset	Conversations			Utterances		
	Train+Val	Test	Total	Train+Val	Test	Total
M3ED*	809	181	990	19,702	4,747	24,449
MELD	1,153	280	1,433	11,098	2,610	13,708
IEMOCAP	120	31	151	5,810	1,623	7,433

Table 1: The statistics of our constructed topic-density M3ED\* and two public topic-sparsity MELD, IEMOCAP datasets. \* denotes this dataset is different from the original.

the total training objective can be denoted as:

$$\mathcal{L}_{total} = \mathcal{L}_{mce} + \alpha \sum_{(a,v,l)} \mathcal{L}_{rec} + \beta \sum_{(a,v,l)} \mathcal{L}_{kl} \quad (7)$$

where  $\sum_{(a,v,l)}$  represents the sum of losses from corresponding three modalities. Besides,  $\alpha$  and  $\beta$  are hyper-parameters of weight to balance the losses between basic MCE approaches and three TopicDiff modules.

## 4. Experimental Settings

### 4.1. Datasets and Baselines

For datasets, we empirically evaluate our TopicDiff approach on our constructed topic-density M3ED\* and two public topic-sparsity MELD, IEMOCAP datasets. **MELD** (Poría et al., 2019) is a topic-sparsity dataset containing multi-speaker conversations from the only one popular TV series *Friends*. **IEMOCAP** (Busso et al., 2008) is also a topic-sparsity dataset containing dyadic conversations where actors perform improvisations or scripted scenarios. **M3ED\*** is a topic-density dataset containing dyadic conversations from 56 different TV series, thus surpassing MELD and IEMOCAP in topic richness. We construct this topic-density M3ED\* dataset based on the original M3ED dataset<sup>4</sup> to evaluate the effectiveness of multimodal topic information. Specifically, for the construction of M3ED\*, we partition the TV series into 56, 10, and 20 for the training, validation and test sets respectively, and maintain the 7:1:2 ratio of utterances, which is different from the original M3ED dataset in Zhao et al. (2022a). To focus on evaluating the effectiveness of TopicDiff towards mining multi-topic information, we ensure that the TV domains in test set are seen in the training set, which furthest excludes the influence of adaptation problem stemming from different TV domains. All the three datasets include aligned acoustic, vision and language three modalities, and a comprehensive statistical summary is reported in Table 1.

For baselines, we choose DialogueCRN (Hu et al., 2021a), MMGCN (Hu et al., 2021b), COGMEN (Joshi et al., 2022), MM-DFN (Hu et al., 2022)

<sup>4</sup><https://github.com/AIM3-RUC/RUCM3ED>

Approach	M3ED*							MELD	IEMOCAP	
	Happy	Neutral	Sad	Disgust	Angry	Fear	Surprise	W-F1	W-F1	W-F1
DialogueCRN	54.38	67.75	54.27	34.59	70.74	12.10	55.55	61.32	54.32 <sup>†</sup>	65.04 <sup>‡</sup>
+ TopicDiff	56.22(↑)	72.21(↑)	55.06(↑)	38.92(↑)	73.41(↑)	28.45(↑)	55.08(↓)	64.49(↑)	55.36(↑)	66.05(↑)
+ TopicDiff w/o TDB	54.13	68.60	54.45	37.94	72.24	27.83	53.21	62.36	54.43	65.31
MMGCN	58.83	69.00	56.68	34.31	69.61	23.47	54.17	62.51	57.26 <sup>†</sup>	66.22 <sup>‡</sup>
+ TopicDiff	62.70(↑)	73.03(↑)	57.80(↑)	38.98(↑)	72.08(↑)	33.02(↑)	55.95(↑)	65.72(↑)	58.26(↑)	67.02(↑)
+ TopicDiff w/o TDB	60.52	72.10	58.11	36.55	71.39	0.8	43.46	63.94	57.63	66.47
COGMEN	59.25	71.20	56.98	40.20	73.50	22.94	58.93	64.88	52.29 <sup>†</sup>	64.56 <sup>†</sup>
+ TopicDiff	60.95(↑)	72.84(↑)	60.180(↑)	38.18(↓)	74.32(↑)	25.63(↑)	60.86(↑)	66.39(↑)	53.54(↑)	65.48(↑)
+ TopicDiff w/o TDB	59.45	71.64	57.29	39.83	73.98	20.37	61.56	65.26	52.76	64.91
MM-DFN	62.29	76.81	60.72	43.58	74.99	14.77	61.88	68.58	57.54 <sup>†</sup>	65.66 <sup>†</sup>
+ TopicDiff	63.69(↑)	77.78(↑)	61.60(↑)	45.66(↑)	76.47(↑)	38.02(↑)	62.140(↑)	70.06(↑)	58.42(↑)	66.52(↑)
+ TopicDiff w/o TDB	62.78	77.57	59.903	44.41	75.76	24.52	60.55	69.10	57.97	65.85
GCNet	46.65	72.24	47.09	27.40	66.77	3.73	38.40	59.02	-	56.18 <sup>‡</sup>
+ TopicDiff	51.54(↑)	71.09(↓)	51.21(↑)	36.46(↑)	71.42(↑)	8.92(↑)	45.63(↑)	61.71(↑)	-	57.80(↑)
+ TopicDiff w/o TDB	50.04	70.97	49.64	24.53	69.39	4.68	41.52	59.78	-	56.78

Table 2: Comparison of several MCE approaches with and without our TopicDiff approach on our constructed topic-density M3ED\* and two public topic-sparsity MELD, IEMOCAP datasets. The grey rows highlights the results of TopicDiff applied to MCE approaches, where  $\uparrow$  and  $\downarrow$  represent improvement and reduction respectively, and  $\uparrow$  indicates better performance for all metrics. Since TopicDiff is model-agnostic, for a fair comparison, we reproduce all the baselines under the same experimental environment and independently run five times with random seeds, taking their average as final results to report. Particularly, these reproduced results with symbol  $\ddagger$  are close to those reported in Hu et al. (2021b) and Lian et al. (2023), while those reproduced results with  $\dagger$  are slightly lower than results reported in their original papers due to different experimental environments. The symbol - denotes the results are not available, because GCNet is not suitable to multiple speakers in MELD dataset as reported in (Lian et al., 2023).

and GCNet (Lian et al., 2023) five MCE approaches to evaluate the effectiveness of our TopicDiff approach. We aim to compare the performance differences of these MCE approaches with and without TopicDiff. It is worth noting that these approaches are all multimodal approaches, rather than only using language modality. Since existing single-text conversation approaches, such as COSMIC (Ghosal et al., 2020) and MuCDN (Zhao et al., 2022b), finetune pre-trained models like RoBERTa (Liu et al., 2019) and incorporate external knowledge (e.g., commonsense and events), we consider it is unfair to compare with single-text conversation approaches. Because the above approaches have different experimental settings (e.g., different features or partitions in datasets, different evaluation metrics, different devices etc.), for a fair and thorough comparison, we reproduce these approaches in our experiments.

## 4.2. Implementation Details and Metrics

Since our TopicDiff approach is model-agnostic, for a fair comparison, we reproduce all the baselines on our topic-density M3ED\* and two public topic-sparsity MELD, IEMOCAP datasets under the same experimental environment (i.e., all the approaches are reproduced with PyTorch on a machine with NVIDIA GeForce RTX 3090, Intel(R) Xeon(R) E5-2650 v4 CPU (2.20 GHz), CUDA version 11.7, and PyTorch 1.7.1 library with python

3.6.13 on Ubuntu 20.04.1 LTS.). Besides, the hyper-parameters of these baselines reported by their public papers are still adopting the same setting, and the others are tuned according to the validation set. Furthermore, we conduct five independent runs with random seeds for each baseline with and without TopicDiff, taking their average as the final results to report. Particularly, the reproduced results of DialogueCRN, MMGCN and GCNet on IEMOCAP could achieve close results as reported in their corresponding papers. However, the reproduced results of other baselines on two public topic-sparsity MELD, IEMOCAP datasets are slightly lower than those results reported in their corresponding papers. For our TopicDiff approach in experimental settings, we use a simple VAE consisting of MLP with linear layers as inference and generative networks, and leverage NCSN (Song and Ermon, 2019) as score-based diffusion backbone. Specifically, we adopt Adam as the optimizer with two initial learning rates (1e-4, 1e-5) and L2 weight decays (1e-4, 1e-4) for VAE and NCSN to update training parameters, respectively. The dropout rate is set to 0.25 for VAE. The dimension of hidden state after VAE is set to 20, which also represents the number of topics. The implementation details of NCSN are referenced in their publicly available code<sup>5</sup>. In addition, we set the weights  $\alpha$  and  $\beta$  to (0.5, 0.5) to balance the losses between MCE

<sup>5</sup><https://github.com/ermongroup/ncsn>




Figure 3: Four line charts to study the robustness of our TopicDiff approach and different modal topic information with the change of topic-density degree via maintaining the total number of training set and varying the numbers of TV series, where the x-axis represents the numbers of TV series. Line charts of (a) and (b) show the performance trend on two MCE approaches (with/without TopicDiff). And line charts of (c) and (d) illustrate the performance trend on acoustic, vision and language topic information via two MCE approaches with TopicDiff. All the line charts are conducted on our constructed topic-density M3ED\* dataset and evaluated on W-F1 metric.

and TopicDiff. All the approaches are trained for a maximum of 200 epochs, and stopped if the validation loss does not decrease for 20 consecutive epochs. To facilitate the corresponding research in this direction, all codes together with datasets are released<sup>6</sup>.

Following Hu et al. (2021a), the performance is evaluated via *Weighted-Average F1* (W-F1) on our constructed topic-density M3ED\* and two public topic-sparsity MELD, IEMOCAP datasets. Besides, we report the F1-score of each emotion on our topic-density M3ED\* dataset. Moreover, *t-test*<sup>7</sup> is used to evaluate the significance of performance differences by following Chen et al. (2020).

## 5. Results and Discussions

### 5.1. Experimental Results

Table 2 illustrates the comparative results for MCE. From the table, we can see that our TopicDiff approach applied to various MCE approaches significantly outperforms those without TopicDiff on our constructed topic-density M3ED\* and two public topic-sparsity MELD, IEMOCAP datasets. For instance, TopicDiff yields an improvement of 3.21% in terms of W-F1 over the MMGCN on our constructed topic-density M3ED\* dataset ( $p$ -value  $< 0.01$ ), and 1% and 1.2% on two public topic-sparsity MELD, IEMOCAP datasets respectively ( $p$ -value  $< 0.05$ ). This indicates that our TopicDiff approach is more effective on topic-density scenarios, and further encourages us to consider incorporating multimodal topic information in MCE.

Furthermore, we report the results for individual emotions on our constructed topic-density M3ED\* dataset, where TopicDiff boosts the performance of most individual emotions. Impressively, TopicDiff applied to MMGCN improves the performance of all

Language	Acoustic	Vision	DialogueCRN	MMGCN
✓			62.34	63.69
	✓		62.78	64.03
		✓	62.82	64.09
✓	✓		63.13	64.43
✓		✓	63.18	64.42
	✓	✓	63.55	64.76
✓	✓	✓	64.49	65.72

Table 3: The effectiveness study of various modal topic information in MCE, where ✓ means that we capture the current modal topic information. The grey columns show the difference of MCE approaches with and without various modal topic information via TopicDiff. All the experiments are conducted on our constructed topic-density M3ED\* dataset and evaluated on W-F1 metric.

individual emotions, with a significant improvement up to 9.55% for the *fear* emotion ( $p$ -value  $< 0.01$ ). This highlights the capability of TopicDiff to boost the detection of challenging emotions.

### 5.2. Ablation Study of TDB

In order to demonstrate the effectiveness of integrating the diffusion model into neural topic model for fully capturing multimodal topic information, we conduct ablation studies on our constructed topic-density M3ED\* and two public topic-sparsity MELD, IEMOCAP datasets. Specifically, we remove the topic-enriched diffusion block (TDB) from our TopicDiff approach, and apply the remaining components (i.e., neural topic model) to various MCE approaches. We present the results (i.e., + TopicDiff w/o TDB) in Table 2 for a clear and intuitive comparison with our primary results. The results show that while the removal of TDB could improve performance, it is not as significant as that achieved with TopicDiff. This indicates that integrating the diffusion model into neural topic model could mitigate diversity deficiency problem, thereby fully capturing multimodal topic information in MCE.

<sup>6</sup><https://github.com/Pomelo518/TopicDiff>

<sup>7</sup><https://docs.scipy.org/doc/scipy/reference/stats.html>

### Multimodal Conversational Sample: A Conversation about Wedding Topic




Figure 4: A multimodal conversational sample includes utterances comprising acoustic spectrums, video frames, language, and ground-truth emotions, alongside the probabilities of ground-truth emotion *joy* on V4 predicted by various approaches. Language/Acoustic/Vision Topic denotes the utilization of TopicDiff to capture the corresponding modal topic information.

### 5.3. Effectiveness Study of Modal Topics

Table 3 presents a comprehensive evaluation of TopicDiff in exploring the contributions of various modal topic information on our constructed topic-density M3ED\* dataset. Specifically, we randomly choose DialogueCRN and MMGCN with TopicDiff to capture various combinations of modal topic information, and calculate the performance improvements between baselines with and without the incorporation of various modal topic information. From the table, we can see that the incorporation of topic information from various modalities still yields improvements in emotion detection, and the use of three modal topic information outperforms that of one or two modal topic information. This demonstrates the importance of topic information in conversation tasks, and further highlights the potential effectiveness of multimodal topic information.

Moreover, we have observed an interesting experimental phenomenon wherein incorporating acoustic or vision topic information leads to better performance than language. For instance, the incorporation of acoustic and vision topic information yields improvements of 0.44% and 0.48% compared to language via using DialogueCRN respectively. This observation demonstrates that compared to language, acoustic and vision topic information is more discriminative, which is reasonable since acoustic and vision signals are more natural and direct than language, making their associated topic information more discriminative to emotion.

### 5.4. Robustness Study of Modal Topics

To investigate the impact of topic-density degree on the effectiveness of our TopicDiff approach, we conduct experiments that involved sampling different numbers of TV series and analyzing the resulting changes in performance. Specifically, we randomly select 5,000 utterances from the training set, while varying the number of TV series to 20, 27, 34, 41, 48, and 56. Subsequently, we randomly choose DialogueCRN and MMGCN, with

and without TopicDiff, to analyze the performance trend concerning changes in the topic-density degree. As illustrated in Figure 3 (a) and (b), all approaches show a decrease in performance due to the reduction in the number of samples from the same TV domain, while maintaining the total number (i.e., 5,000) of the training set<sup>8</sup> and increasing topic-density, necessitating the use of multimodal topic information to sustain performance. We observe a significant decline in performance for the two MCE approaches, while TopicDiff exhibits a slight and stable decline. This demonstrates that our TopicDiff approach is robust since it incorporates additional multimodal topic information.

To further explore the performance difference of various modal topic information under the change of topic-density degree, we randomly choose DialogueCRN and MMGCN with TopicDiff to capture acoustic, vision or language topic information. We plot the trend of performance changes in Figure 3 (c) and (d). The results reveal that TopicDiff, with either acoustic or vision topic information, maintains a more stable performance compared to language, even with changes in topic-density degree. This interesting observation indicates that acoustic and vision topic information is more robust than language, which is reasonable since the emotional clues expressed by acoustic and vision tend to be more consistent across different topics than language. For example, the facial expressions of *smile* are always similar in different topics, thus making them have better robustness.

### 5.5. Qualitative Study of TopicDiff

As shown in Figure 4, we randomly choose DialogueCRN with and without TopicDiff and leverage TopicDiff to capture the topic information of language, acoustic and vision to predict the probability of ground-truth emotion *joy* on V4. From this fig-

<sup>8</sup>To furthest exclude the influence of domain adaptation, similar to the construction of M3ED\*, we ensure that TV domains in test set are seen in the training set.



ure, we can see that: 1) Predicting the emotion of V4 in the absence of context is challenging. Moreover, the use of global language topic clues (i.e., “worried about”, “nervous”) may lead to negative predictions. Thus, it is essential to combine acoustic and vision topic clues (i.e., *joy tones*, *smile face*) for predicting the emotion of V4 precisely. This highlights the importance of multimodal topic information in conversations. 2) DialogueCRN with TopicDiff provides a higher probability of *joy* emotion on V4 compared to DialogueCRN, indicating the effectiveness of TopicDiff in capturing multimodal topic information. 3) TopicDiff without TDB provides a lower probability compared to TopicDiff, indicating the effectiveness of integrating the diffusion model to capture multimodal topic information. 4) The prediction probabilities of acoustic or vision topic are obviously higher than language topic, indicating that acoustic and vision topic information is more discriminative compared to language.

## 6. Conclusion

In this paper, we propose a model-agnostic **Topic-enriched Diffusion** (TopicDiff) approach to MCE, which integrates the diffusion model into neural topic model for mitigating the problem of semantic diversity deficiency in capturing multimodal topic information. Detailed experiments on our topic-density M3ED\* and two public topic-sparsity MELD, IEMOCAP datasets demonstrate the effectiveness of our TopicDiff approach in MCE. Furthermore, we observe an interesting finding that acoustic and vision topic information is more discriminative and robust compared to language. In our future work, we would like to improve the emotion detection by incorporating more information, such as speaker personality and posture information. Besides, we would like to leverage large-scale multimodal video pre-training to enhance the representation ability of basic modalities. Moreover, we intend to transfer TopicDiff to other conversation based multimodal tasks, such as multimodal psychological counseling and conversational depression detection.

## Ethics Statement

**Data Disclaimer.** We construct a topic-density M3ED\* dataset based on M3ED, which is available at <https://github.com/AIM3-RUC/RUCM3ED>. M3ED dataset is widely used by other academics and are typically accessible to the public. Therefore, our proposed dataset does not involve any sensitive information that may harm others.

**Human Construction.** When recruiting constructors for this study, we claim that all potential constructors are free to choose whether they want to participate, and they can withdraw from the

study anytime without any negative repercussions. Thus, the establishment of our topic-density M3ED\* dataset is compliant with ethics.

## Acknowledgments

We thank our anonymous reviewers for their helpful comments. This work was supported by three NSFC grants, i.e., No.62006166, No.62376178 and No.62076175. This work was also supported by a Project Funded by the Priority Academic Program Development of Jiangsu Higher Education Institutions (PAPD).

## References

- Minghui An, Jingjing Wang, Shoushan Li, and Guodong Zhou. 2020. [Multimodal topic-enriched auxiliary learning for depression detection](#). In *Proceedings of COLING 2020*.
- Dmitry Baranchuk, Andrey Voynov, Ivan Rubachev, Valentin Khruikov, and Artem Babenko. 2022. Label-efficient semantic segmentation with diffusion models. In *Proceedings of ICLR 2022*.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeanette N. Chang, Sungbok Lee, and Shrikanth S. Narayanan. 2008. [IEMOCAP: interactive emotional dyadic motion capture database](#). *Lang. Resour. Evaluation*, 42(4):335–359.
- Jie Cao, Michael Tanana, Zac E. Imel, Eric Poitras, David C. Atkins, and Vivek Srikumar. 2019. [Observing dialogue in therapy: Categorizing and forecasting behavioral codes](#). In *Proceedings of ACL 2019*, pages 5599–5611.
- Xiao Chen, Changlong Sun, Jingjing Wang, Shoushan Li, Luo Si, Min Zhang, and Guodong Zhou. 2020. [Aspect sentiment classification with document-level sentiment preference modeling](#). In *Proceedings of ACL 2020*, pages 3667–3677.
- Jun Gao, Wei Bi, Xiaojiang Liu, Junhui Li, Guodong Zhou, and Shuming Shi. 2019. [A discrete CVAE for response generation on short-text conversation](#). In *Proceedings of EMNLP 2019*, pages 1898–1908.
- Deepanway Ghosal, Navonil Majumder, Alexander F. Gelbukh, Rada Mihalcea, and Soujanya Poria. 2020. [COSMIC: commonsense knowledge for emotion identification in conversations](#). In *Findings of EMNLP 2020*.

- Deepanway Ghosal, Navonil Majumder, Soujanya Poria, Niyati Chhaya, and Alexander F. Gelbukh. 2019. [Dialoguecn: A graph convolutional neural network for emotion recognition in conversation](#). In *Proceedings of EMNLP 2019*.
- Lin Gui, Jia Leng, Gabriele Pergola, Yu Zhou, Ruifeng Xu, and Yulan He. 2019. [Neural topic model with reinforcement learning](#). In *Proceedings of EMNLP 2019*.
- Devamanyu Hazarika, Soujanya Poria, Rada Mihalcea, Erik Cambria, and Roger Zimmermann. 2018a. [ICON: interactive conversational memory network for multimodal emotion detection](#). In *Proceedings of EMNLP 2018*.
- Devamanyu Hazarika, Soujanya Poria, Amir Zadeh, Erik Cambria, Louis-Philippe Morency, and Roger Zimmermann. 2018b. [Conversational memory network for emotion recognition in dyadic dialogue videos](#). In *Proceedings of NAACL 2018*.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. 2022. [Masked autoencoders are scalable vision learners](#). In *Proceedings of CVPR 2022*, pages 15979–15988.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. [Denosing diffusion probabilistic models](#). In *Proceedings of NeurIPS 2020*.
- Dou Hu, Xiaolong Hou, Lingwei Wei, Lian-Xin Jiang, and Yang Mo. 2022. [MM-DFN: multimodal dynamic fusion network for emotion recognition in conversations](#). In *Proceedings of ICASSP 2022*, pages 7037–7041.
- Dou Hu, Lingwei Wei, and Xiaoyong Huai. 2021a. [Dialoguecrn: Contextual reasoning networks for emotion recognition in conversations](#). In *Proceedings of ACL 2021*.
- Jingwen Hu, Yuchen Liu, Jinming Zhao, and Qin Jin. 2021b. [MMGCN: multimodal fusion via deep graph convolution network for emotion recognition in conversation](#). In *Proceedings of ACL 2021*.
- Ziqi Huang, Kelvin C.K. Chan, Yuming Jiang, and Ziwei Liu. 2023. [Collaborative diffusion for multimodal face generation and editing](#). In *Proceedings of CVPR 2023*.
- Taichi Ishiwatari, Yuki Yasuda, Taro Miyazaki, and Jun Goto. 2020. [Relation-aware graph attention networks with relational position encodings for emotion recognition in conversations](#). In *Proceedings of EMNLP 2020*.
- Wenxiang Jiao, Haiqin Yang, Irwin King, and Michael R. Lyu. 2019. [Higr: Hierarchical gated recurrent units for utterance-level emotion recognition](#). In *Proceedings of NAACL 2019*, pages 397–406.
- Yuan Jin, He Zhao, Ming Liu, Lan Du, and Wray L. Buntine. 2021. [Neural attention-aware hierarchical topic model](#). In *Proceedings of EMNLP 2021*.
- Abhinav Joshi, Ashwani Bhat, Ayush Jain, Atin Vikram Singh, and Ashutosh Modi. 2022. [COGMEN: contextualized GNN based multimodal emotion recognition](#). In *Proceedings of NAACL 2022*, pages 4148–4164.
- Diederik P. Kingma and Max Welling. 2014. [Auto-encoding variational bayes](#). In *Proceedings of ICLR 2014*.
- Jiangnan Li, Zheng Lin, Peng Fu, and Weiping Wang. 2021. [Past, present, and future: Conversational emotion recognition through structural modeling of psychological knowledge](#). In *Findings of EMNLP 2021*.
- Zheng Lian, Lan Chen, Licai Sun, Bin Liu, and Jianhua Tao. 2023. [Gcnet: Graph completion network for incomplete multimodal learning in conversation](#). *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(7):8419–8432.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pre-training approach](#). *CoRR*, abs/1907.11692.
- Navonil Majumder, Soujanya Poria, Devamanyu Hazarika, Rada Mihalcea, Alexander F. Gelbukh, and Erik Cambria. 2019. [Dialoguecrn: An attentive RNN for emotion detection in conversations](#). In *Proceedings of AAAI 2019*.
- Yishu Miao, Lei Yu, and Phil Blunsom. 2016. [Neural variational inference for text processing](#). In *Proceedings of ICML 2016*, volume 48, pages 1727–1736.
- Soujanya Poria, Erik Cambria, Devamanyu Hazarika, Navonil Majumder, Amir Zadeh, and Louis-Philippe Morency. 2017. [Context-dependent sentiment analysis in user-generated videos](#). In *Proceedings of ACL 2017*.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. [MELD: A multimodal multi-party dataset for emotion recognition in conversations](#). In *Proceedings of ACL 2019*.

- Shi Pu, Yijiang He, Zheng Li, and Mao Zheng. 2020. [Multimodal topic learning for video recommendation](#). *CoRR*, abs/2010.13373.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. [High-resolution image synthesis with latent diffusion models](#). In *Proceedings of CVPR 2022*, pages 10674–10685.
- Ludan Ruan, Yiyang Ma, Huan Yang, Huiguo He, Bei Liu, Jianlong Fu, Nicholas Jing Yuan, Qin Jin, and Baining Guo. 2023. Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation. In *Proceedings of CVPR 2023*.
- Weizhou Shen, Junqing Chen, Xiaojun Quan, and Zhixian Xie. 2021. Dialogxl: All-in-one xlnet for multi-party conversation emotion recognition. In *Proceedings of AAAI 2021*.
- Yang Song and Stefano Ermon. 2019. Generative modeling by estimating gradients of the data distribution. In *Proceedings of NeurIPS 2019*.
- Arash Vahdat, Karsten Kreis, and Jan Kautz. 2021. Score-based generative modeling in latent space. In *Proceedings of NeurIPS 2021*.
- Jiancheng Wang, Jingjing Wang, Changlong Sun, Shoushan Li, Xiaozhong Liu, Luo Si, Min Zhang, and Guodong Zhou. 2020a. Sentiment classification in customer service dialogue with topic-aware multi-task learning. In *Proceedings of AAAI 2020*.
- Rui Wang, Xueming Hu, Deyu Zhou, Yulan He, Yuxuan Xiong, Chenchen Ye, and Haiyang Xu. 2020b. [Neural topic modeling with bidirectional adversarial training](#). In *Proceedings of ACL 2020*.
- Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Yingxia Shao, Wentao Zhang, Ming-Hsuan Yang, and Bin Cui. 2022. [Diffusion models: A comprehensive survey of methods and applications](#). *CoRR*, abs/2209.00796.
- Jinming Zhao, Tenggao Zhang, Jingwen Hu, Yuchen Liu, Qin Jin, Xinchao Wang, and Haizhou Li. 2022a. [M3ED: multi-modal multi-scene multi-label emotional dialogue database](#). In *Proceedings of ACL 2022*.
- Weixiang Zhao, Yanyan Zhao, and Bing Qin. 2022b. Mucdn: Mutual conversational detachment network for emotion recognition in multi-party conversations. In *Proceedings of COLING 2022*.
- Lixing Zhu, Gabriele Pergola, Lin Gui, Deyu Zhou, and Yulan He. 2021. [Topic-driven and knowledge-aware transformer for dialogue emotion detection](#). In *Proceedings of ACL 2021*.