

Xiaolin Chen cxlicd@gmail.com School of Software, Joint SDU-NTU Centre for Artificial Intelligence Research, Shandong University Xuemeng Song* sxmustc@gmail.com School of Computer Science and Technology, Shandong University Yinwei Wei weiyinwei@hotmail.com School of Computing, National University of Singapore

Liqiang Nie* nieliqiang@gmail.com School of Computer Science and Technology, Harbin Institute of Technology (Shenzhen) Tat-Seng Chua dcscts@nus.edu.sg School of Computing, National University of Singapore

ABSTRACT

Textual response generation is an essential task for multimodal task-oriented dialog systems. Although existing studies have achieved fruitful progress, they still suffer from two critical limitations: 1) focusing on the attribute knowledge but ignoring the relation knowledge that can reveal the correlations between different entities and hence promote the response generation, and 2) only conducting the cross-entropy loss based output-level supervision but lacking the representation-level regularization. To address these limitations, we devise a novel multimodal task-oriented dialog system (named MDS-S²). Specifically, MDS-S² first simultaneously acquires the context related attribute and relation knowledge from the knowledge base, whereby the non-intuitive relation knowledge is extracted by the *n*-hop graph walk. Thereafter, considering that the attribute knowledge and relation knowledge can benefit the responding to different levels of questions, we design a multi-level knowledge composition module in MDS-S² to obtain the latent composed response representation. Moreover, we devise a set of latent query variables to distill the semantic information from the composed response representation and the ground truth response representation, respectively, and thus conduct the representation-level semantic regularization. Extensive experiments on a public dataset have verified the superiority of our proposed MDS-S². We have released the codes and parameters to facilitate the research community.

CCS CONCEPTS

 \bullet Computing methodologies \rightarrow Natural language generation; Discourse, dialogue and pragmatics.

SIGIR '23, July 23–27, 2023, Taipei, Taiwan

KEYWORDS

Multimodal Task-oriented Dialog Systems; Dual Semantic Knowledge; Representation-level Regularization

ACM Reference Format:

Xiaolin Chen, Xuemeng Song*, Yinwei Wei, Liqiang Nie*, and Tat-Seng Chua. 2023. Dual Semantic Knowledge Composed Multimodal Dialog Systems. In Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23), July 23–27, 2023, Taipei, Taiwan. ACM, New York, NY, USA, 10 pages. https: //doi.org/10.1145/3539618.3591673

1 INTRODUCTION

In recent years, task-oriented dialog systems have penetrated into many aspects of our daily life, such as restaurant reserving and ticket booking. According to the report of Salesforce¹, roughly 68% of customers tend to interact with the intelligent dialog agents for their quick responses rather than waiting for the human services. Considering its value, a surge of researches are dedicated to developing task-oriented dialog systems. Early studies in this research line focus on the pure text-based dialog system [8, 28], overlooking that both the user and the agent may need to express themselves with certain visual clues (*i.e.*, images). For example, as shown in Figure 1, the user needs to utilize the image to express his/her desired shopping mall in the utterance u_7 , while the agent needs to use images to illustrate special dishes for the user in u_4 . Therefore, recent research attention has been swifted to the multimodal task-oriented dialog systems.

In fact, multimodal task-oriented dialog systems mainly contain two tasks [2]: the textual response generation and the image response selection. Considering that the former is more challenging and its performance is still far from satisfactory, many researchers focus on this task for multimodal task-oriented dialog systems [16, 19]. Despite the favorable performance obtained by existing efforts [2, 4, 16, 19, 21, 27, 37], they mainly suffer from two critical limitations. 1) **Ignoring the relation knowledge.** In the context of multimodal task-oriented dialog systems, there is always a knowledge base containing abundant attribute-value pairs as well as images of a large number of entities. Previous studies focus on exploiting the attribute knowledge of entities, but neglecting the relation knowledge residing in the knowledge

^{*}Corresponding authors: Xuemeng Song and Liqiang Nie.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

^{© 2023} Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-9408-6/23/07...\$15.00 https://doi.org/10.1145/3539618.3591673

¹https://startupbonsai.com/chatbot-statistics.

SIGIR '23, July 23-27, 2023, Taipei, Taiwan



Figure 1: Illustration of a multimodal dialog system between a user and an agent. "u" refers to the utterance.

base, which can capture the relations between entities, and benefit the response reasoning and generation. For example, as shown in Figure 1, the agent can generate the appropriate response (*i.e.*, u_8) only conditioned on the relation knowledge of Inaniwa Yosuke \xrightarrow{near} Wisma Atria \xrightarrow{domain} mall. 2) **Lacking the representation-level regularization.** Previous studies only adopt the token-level cross-entropy loss to regulate the generated response to be similar to the ground truth response. This may be insufficient for the task whose input (*i.e.*, text and image) and output (*i.e.*, text) present apparent heterogeneity. In fact, they ignore the potential representation-level regularization between the context-knowledge composed response representation and the ground truth response representation, which can enhance the composed response representation learning and hence improve the response generation performance.

To address these two limitations, in this work, we aim to conduct the research of textual response generation in multimodal task-oriented dialog systems by integrating the dual semantic knowledge (i.e., attribute and relation knowledge) and the representation-level regularization. However, this is non-trivial due to the following three challenges. 1) Different from the attribute knowledge, the relation one is not straightforwardly provided by the knowledge base that only contains attribute-value pairs and images of entities. Hence, how to mine the related relation knowledge with respect to the given multimodal context is a crucial challenge. 2) In a sense, the intuitive attribute knowledge is beneficial to response generation of simple questions (e.g., "Can you get their phone number for me?"), while the relation knowledge is helpful for responding relatively more complicated questions (e.g., "Can you help me look for a hotel nearby Singapore River?"). Accordingly, how to effectively compose the multimodal context with the dual semantic knowledge and thus generate the proper response is another vital challenge. And 3) ideally, we expect that the representation-level regularization can project the context-knowledge composed response representation

and the ground truth response representation into the same meaningful semantic space. In this way, we can yield meaningful composed response representation that can further enhance the response generation. Therefore, how to fulfil the meaningful representation-level semantic regularization is another challenge.

To address the aforementioned challenges, we devise a novel dual semantic knowledge composed multimodal dialog system, MDS-S² for short, where the generative pretrained language model BART is adopted as the backbone. As demonstrated in Figure 2, the proposed model consists of three pivotal components: dual semantic knowledge acquisition, multi-level knowledge composition, and representation-regularized response generation. To be specific, the first component aims to acquire the context related dual semantic knowledge: attribute knowledge and relation knowledge. In particular, the related relation knowledge is uncovered by the *n*-hop graph walk over the whole knowledge base. Thereafter, the second component is devised to compose the multimodal context and the acquired dual semantic knowledge to obtain the latent composed response representation. Specifically, considering that the attribute knowledge and relation knowledge can facilitate the responding to questions with different complexity levels of user intentions, the attribute knowledge is first composed at the input token level, and the relation knowledge is then adaptively composed at the intermediate representation level. Subsequently, the last component targets at enhancing the proper textual response generation with the additional representation-level regularization. In particular, we design a set of to-be-learned latent query variables to project the composed response representation and the ground truth response representation into the same semantic space with the cross-attention mechanism. Moreover, to fully utilize the representation-level regularization, we also design a semantic-enhanced response decoder. Notably, the decoder can adaptively incorporate the regularized composed response semantic representation, apart from the original multi-level knowledge composed response representation. Extensive experiments on a public dataset have demonstrated the superiority of our proposed MDS-S². Our main contributions can be summarized as follows:

- We propose a novel dual semantic knowledge composed multimodal dialog system. To the best of our knowledge, we are among the first to exploit the relation knowledge residing in the knowledge base and integrate the representation-level semantic regularization for the textual response generation in multimodal task-oriented dialog systems.
- We present the dual semantic knowledge acquisition component to select the context related knowledge from both the attribute and relation perspectives. Moreover, we devise the multi-level knowledge composition component, which can compose the attribute and relation knowledge at the input token level and the intermediate representation level, respectively.
- We devise a set of to-be-learned latent variables to conduct the representation-level semantic regularization, and the semantic-enhanced response decoder to promote the textual response generation. As a byproduct, we release the codes and involved parameters to facilitate the research community².

²https://sigir2023.wixsite.com/anonymous7357.





2 RELATED WORK

Traditional task-oriented dialog systems [11, 29] resort to a pipeline structure and mainly contain four functional components: natural language understanding, dialogue state tracking, policy learning, and natural language generation. To be more specific, the natural language understanding component is used to classify the user intention, and then the dialogue state tracking component aims to track the immediate state, based on which the policy learning component can predict the following action. Thereafter, the natural language generation component exhibits the response through generation methods [8, 14] or predefined templates. Although pipeline methods have attained impressive results, they may suffer from error propagation [10] on the sequential components.

With the flourishing development of deep neural networks, recent studies are centered on exploring end-to-end task-oriented dialog systems. Early end-to-end studies focus on single-modal (*i.e.*, textual modality) task-oriented dialog systems. Although these studies have made tremendous strides, they neglect that both the user and agent may need to leverage certain images to convey their needs or services. Accordingly, Saha et al. [27] investigated the multimodal task-oriented dialog systems with two critical tasks: textual response generation and image response selection, and presented a multimodal hierarchical encoder-decoder model (MHRED). In addition, they released a large-scale multimodal dialog dataset in the fashion domain, which considerably stimulates the progress on multimodal task-oriented dialog systems. Beyond this, several studies further probe the semantic relation in the multimodal dialog context and integrate the knowledge based on the framework of MHRED [1, 4, 16, 21, 22, 37]. More recently, several studies draw on Transformer [30] to propel the development of multimodal dialog systems [2, 9, 20]. Although these studies achieve remarkable performance, they only utilize the attribute knowledge and overlook the representation-level regularization. Beyond that, in this paper, we worked on investigating the dual semantic knowledge composition and representation-level semantic regularization to improve the response generation performance.

3 MODEL

3.1 **Problem Formulation**

Suppose we have a set of N training dialog pairs $\mathcal{D} = \{(C_1, \mathcal{R}_1), (C_2, \mathcal{R}_2), \cdots, (C_N, \mathcal{R}_N)\}$. Thereinto, each dialog pair consists of a multimodal dialog context C_i and a ground truth response \mathcal{R}_i . In particular, each utterance in C_i may involve both textual and visual modalities, as the user/agent may utilize certain related images to promote the request/response expression. In light of this, each multimodal dialog context C_i can be further represented by two modalities: the sequence of tokens $\mathcal{T}_i = [t_g^i]_{g=1}^{N_I^t}$ derived by concatenating all the textual utterances in the context and a set of images $\mathcal{V}_i = \{v_j^i\}_{j=1}^{N_V^i}$ involved in the context, where t_g^i refers to the g-th token and v_j^i represents the j-th image of C_i . N_T^i and N_V^i are the total number of tokens and images, respectively. Notably, $N_V^i = 0$ (*i.e.*, $\mathcal{V}_i = \emptyset$), if there is no image involved in C_i .

where r_n^i stands for the *n*-th token and N_R^i is the number of tokens in the response. In addition, we have a knowledge base including both the semantic and visual knowledge of N_K entities $\mathcal{K} = \{e_p\}_{p=1}^{N_K}$ to assist in response generation. To be specific, each entity e_p is associated with a set of attribute-value pairs \mathcal{R}_p (*e.g.*, {-location: Orchard Road>, -domain: food>}) and a set of images I_p that exhibit the visual information of the entity (*e.g.*, showing the appearance of Esplanade Park).

In a sense, we aim to devise a novel model \mathcal{F} which can generate the appropriate textual response based on the given multimodal dialog context and the knowledge base as follows,

$$\mathcal{F}(C_i, \mathcal{K}|\Theta_F) \to \mathcal{R}_i, \tag{1}$$

where Θ_F denotes the model parameters.

3.2 Dual Semantic Knowledge Acquisition

Since knowledge plays a vital role in the response generation of task-oriented dialog systems, we first conduct the knowledge acquisition for the given multimodal context. Considering the semantic knowledge is pivotal to capturing the user's intentions [33, 36, 37], we focus on selecting two kinds of semantic knowledge: attribute knowledge and relation knowledge. Thereinto, the attribute knowledge, which is widely used, refers to the attribute-value pairs of entities mentioned directly in the context. In a sense, the attribute knowledge can be useful for responding the simple questions, like "Can you get their phone number for me?". In addition to the attribute knowledge, beyond previous methods, we also incorporate the relation knowledge contained in the knowledge base, which helps to uncover the correlation between entities and respond to relatively more complicated questions, like "Can you help me look for a hotel nearby Esplanade Park to stay at?". To answer this complicated question, we need to first identify which entities are near the entity "Esplanade Park", and then recognize which nearby entity is a "hotel" with the attribute knowledge of nearby entities. Therefore, we devise the dual semantic knowledge acquisition component with two modules: attribute knowledge acquisition and relation knowledge acquisition.

3.2.1 Attribute Knowledge Acquisition. Due to the multimodal nature of the given dialog context, following the existing method [2], we retrieve the related attribute knowledge according to both the textual and visual context.

For the textual context, we directly judge which knowledge entity is mentioned to obtain the textual context related knowledge entities. Namely, if the knowledge entity e_p appears in the textual context, we will select the set of attribute-value pairs \mathcal{A}_p of it as the relevant knowledge. In this way, we can collect the related knowledge of textual context $\mathcal{K}_t^A = \mathcal{A}_1^t \cup \mathcal{A}_2^t \cup \cdots \cup \mathcal{A}_{N_k^t}^t$, where \mathcal{A}_m^t is the set of attribute-value pairs of the *m*-th mentioned knowledge entity and N_k^t is the total number of textual context related knowledge entities.

Pertaining to the visual context, we resort to mining its visual features to obtain the related entities in the knowledge base. Specifically, we first utilize ViT-B/32 [6] to derive the visual features of images in context C_i and that of each entity in the knowledge base \mathcal{K} . Notably, similar to the visual context, each entity in the

knowledge base can be associated with multiple images. Thereafter, for each image v_j in the visual context \mathcal{V}_i , we calculate its visual similarity³ with each image of each entity in the knowledge base and regard the maximum image similarity as the entity similarity with the given image. Notably, to guarantee the quality of the retrieved knowledge, we only regard the entities whose similarity with the given context image is larger than the threshold ϵ as the related entities. Ultimately, by merging the attribute-value pairs of all the visual context related entities, we obtain the visual context related attribute knowledge as $\mathcal{K}_v^A = \mathcal{A}_1^v \cup \mathcal{A}_2^v \cup \cdots \cup \mathcal{A}_{N_k^v}^v$, where \mathcal{A}_n^v denotes the attribute-value pairs set of the *n*-th related knowledge entity and N_k^v is the total number of visual context related knowledge entities.

3.2.2 Relation Knowledge Acquisition. To obtain the related relation knowledge, we first cast the whole knowledge base \mathcal{K} into a directed knowledge graph $\mathcal{G}_a = \{\mathcal{E}_a, \mathcal{R}_a\}$. $\mathcal{E}_a = \{e_q\}_{q=1}^{N_a^K}$ is the set of nodes, including two types of nodes (*i.e.*, head and tail nodes), where the head node refers to a knowledge entity, while the tail node denotes an attribute value. N_a^K is the total number of nodes. $\mathcal{R}_a = \{r_z\}_{z=1}^{N_a^R}$ is the edge set, where each r_z refers to an attribute type linking the head and tail nodes. N_a^R is the number of attribute types in the knowledge base. Intuitively, each triplet (h, r, t), where $h, t \in \mathcal{E}_a$ and $r \in \mathcal{R}_a$, indicates that the attribute value of the knowledge entity h regarding the attribute type r is t. For example, the triplet (*Wisma Atria, location, Orchard Road*) indicates that the "location" of the entity "Wisma Atria" is at "Orchard Road".

Thereafter, for the given multimodal dialog context, we first identify the entities involved in the given context in the same way as the attribute knowledge acquisition module. Then, for each identified entity, we perform the *n*-hop graph walk over G_a to uncover the potential relations it involves. In particular, each hop walk traverses a triplet (*i.e.*, $\langle e, r, \bar{e} \rangle$) and the *n*-hop graph walk process will terminate when the last traversed node of the current walk does not connect to any other nodes, or the number of walks reaches the pre-defined maximum number (i.e., n). Finally, each *n*-hop graph walk yields a sequence of traversed triplets, which can compose a high-order relation of the initial entity node. Formally, we use a tuple to represent each relation, whose entries include the sequence of traversed nodes and edges. For example, the high-order relation $e_1 \xrightarrow{r_1} e_2 \xrightarrow{r_2} e_3$ can be represented as $[e_1, r_1, e_2, r_2, e_3]$. In this vein, we can derive all the context related relation tuples $\mathcal{E}_{H}^{i}=\{h_{q}^{i}\}_{q=1}^{N_{h}^{i}},$ where each relation tuple h_{q}^{i} contains an arbitrary number of entries and N_h^i is the number of relation tuples.

3.3 Multi-level Knowledge Composition

As aforementioned, the attribute knowledge refers to the attribute-value pairs of entities mentioned directly in the context, whose role is straightforward in responding questions that contain shallow user intentions. Meanwhile, the relation knowledge uncovers the correlation between entities, which can benefit the responding to questions that contain complex user intentions. In light of this, we devise the multi-level knowledge composition

³Here, we use the cosine similarity.

component with the shallow attribute knowledge composition and complex relation knowledge composition. For simplicity, we temporally omit the script *i* that indexes the training samples.

3.3.1 Shallow Attribute Knowledge Composition. In this module, we first extract the embeddings of the attribute knowledge and multimodal context, respectively. As for the attribute knowledge, we merge \mathcal{K}_t^A and \mathcal{K}_v^A as a whole $\mathcal{K}_A = [\mathcal{K}_t^A, \mathcal{K}_v^A]$. Thereafter, we treat the set of attribute-value pairs in \mathcal{K}_A as a sequence of tokens, and feed it into the position-wise embedding layer of BART to obtain the attribute knowledge embedding $\mathbf{E}_k \in \mathbb{R}^{N_K \times D}$, where N_K is the number of tokens in \mathcal{K}_A . In the same manner, we can obtain the embedding for the given textual context $\mathcal{T} = [t_1, t_2, \cdots, t_{N_T}]$, denoted as $\mathbf{E}_t \in \mathbb{R}^{N_T \times D}$. For the given visual context, *i.e.*, the set of images $\mathcal{V} = \{v_1, v_2, \cdots, v_{N_V}\}$, we first utilize ViT-B/32 [6] (*i.e.*, \mathcal{B}_v) pretrained by CLIP [25] to encode each image v_i , and then employ a fully connected layer and the layer normalization to get each image embedding $\bar{\mathbf{v}}_i$ as follows,

$$\begin{cases} \mathbf{v}_j = \mathcal{B}_v(v_j), j = 1, 2, \cdots, N_V, \\ \bar{\mathbf{v}}_j = LN(\mathbf{v}_j^\top \mathbf{W}_v^e + \mathbf{b}_v^e), \end{cases}$$
(2)

where \mathbf{v}_i is the visual representation, extracted by ViT-B/32, of the image v_j , and $LN(\cdot)$ refers to the layer normalization. \mathbf{W}_n^e and \mathbf{b}_{n}^{e} are the to-be-learned weight matrix and bias vector of the fully connected layer, respectively. Finally, let $\mathbf{E}_v = [\bar{\mathbf{v}}_1; \bar{\mathbf{v}}_2; \cdots; \bar{\mathbf{v}}_{N_V}]^\top \in$ $\mathbb{R}^{N_V \times D}$ denote the embedding of the visual context.

Thereafter, to derive the attribute knowledge composed response representation, we feed the embeddings of the attribute knowledge and the multimodal context into the encoder \mathcal{B}_{e} of BART as follows,

$$\mathbf{T}_t = \mathcal{B}_e([\mathbf{E}_k, \mathbf{E}_t, \mathbf{E}_v]),\tag{3}$$

where $\mathbf{T}_t \in \mathbb{R}^{N_b \times D}$ is the attribute knowledge composed response representation. $N_b = (N_K + N_T + N_V)$ is the total number of tokens.

3.3.2 Complex Relation Knowledge Composition. In this module, we further integrate the complex relation knowledge to refine the above composed response representation. To be specific, we regard each related relation tuple (*i.e.*, h_q^i) as a sequence of words and feed it into the position-wise embedding layer of BART to get its embedding $\mathbf{E}_q^h \in \mathbb{R}^{N_q^e \times D}$, where N_q^e is the number of tokens in h_q^i . Thereafter, we employ the encoder \mathcal{B}_e of BART to derive the relation tuple representation as follows,

$$\mathbf{T}_q^h = \mathcal{B}_e(\mathbf{E}_q^h),\tag{4}$$

where $\mathbf{T}_q^h \in \mathbb{R}^{N_q^e \times D}$ is the representation of the relation tuple h_q^i . In fact, different relation tuples tend to play different roles in enhancing the textual response generation. For example, as shown in Figure 1, the relation tuple (Inaniwa Yosuke \xrightarrow{near} Wisma Atria \xrightarrow{domain} mall) conveys more vital clues in generating the response u_8 than (Inaniwa Yosuke \xrightarrow{near} Wisma Atria $\xrightarrow{creditcards}$ yes). Therefore, we resort to the cross-attention mechanism [30] to emphasize relation tuples that are highly related to the composed response representation and obtain the reorganized relation tuples representation for the given attribute composed representation, due to its superior performance in capturing the interaction relation between two items [17, 18, 31, 32, 34]. Specifically, we treat the composed response representation T_t obtained in Eqn.(3) as the query, and the relation tuple representations T_q^h as the key and value as follows,

$$\begin{cases} Q_m = T_t W_Q^m, K_m = T_h W_K^m, V_m = T_h W_V^m, \\ \bar{T}_h = softmax(Q_m K_m^-) V_m, \end{cases}$$
(5)

where $\mathbf{T}_h = [\mathbf{t}_h^1, \mathbf{t}_h^2, \cdots, \mathbf{t}_h^{N_h}] \in \mathbb{R}^{N_h \times D}$ denotes the initial representation of all the related relation tuples, and $\mathbf{t}_h^q =$ $avg(\mathbf{T}_q^h)$ refers to the average pooling over the *q*-th relation tuple representation \mathbf{T}_q^h obtained in Eqn.(4). \mathbf{W}_Q^m , \mathbf{W}_K^m , and \mathbf{W}_V^m are weight matrices. $\mathbf{Q}_m \in \mathbb{R}^{N_b \times D}$, $\mathbf{K}_m \in \mathbb{R}^{N_h \times D}$ and $\mathbf{V}_m \in \mathbb{R}^{N_h \times D}$ $\mathbb{R}^{N_h \times D}$ are the query, key, value matrices, respectively. softmax(·) represents the softmax activation function, and $\bar{\mathbf{T}}_h \in \mathbb{R}^{N_b \times D}$ stands for the reorganized relation tuples representation for the given attribute composed representation.

Next, to fulfill the relation knowledge composition, instead of merely using the common residual operation [13], similar to [7], we adopt the attention mechanism to adaptively fuse the attribute knowledge composed response representation and the reorganized relation tuples representation. The reason behind is that they may contribute differently towards the ground truth response generation. According to the attention mechanism, we have,

$$\begin{cases} \mathbf{H}_{t} = tanh(\mathbf{T}_{t}\mathbf{W}_{t} + \mathbf{B}_{t}), \\ \mathbf{H}_{h} = tanh(\bar{\mathbf{T}}_{h}\mathbf{W}_{h} + \mathbf{B}_{h}), \\ [\mathbf{r}_{t}, \mathbf{r}_{h}] = softmax([\mathbf{H}_{t}, \mathbf{H}_{h}]\mathbf{a}), \end{cases}$$
(6)

where W_t and W_h are weight matrices, while B_t and B_h are bias matrices. $\mathbf{r}_t \in \mathbb{R}^{N_b}$ and $\mathbf{r}_h \in \mathbb{R}^{N_b}$ denote the normalized confidence vectors for the attribute knowledge composed response representation and the relation knowledge representation. Here, the to-be-learned vector $\mathbf{a} \in \mathbb{R}^D$ can be interpreted as the query "which part contributes more to the response generation". Ultimately, we reach the final multi-level knowledge composed response representation $\mathbf{T}_c \in \mathbb{R}^{N_b \times D}$ as follows,

$$\mathbf{T}_c = \mathbf{r}_t \odot \mathbf{T}_t + \mathbf{r}_h \odot \bar{\mathbf{T}}_h,\tag{7}$$

where \odot represents the element-wise multiplication operation.

3.4 **Representation-regularized Response** Generation

Towards the final response generation, one straightforward and commonly used solution is to feed the composed response representation into the BART decoder, and utilize the cross-entropy loss to perform the output-level supervision. Although the method is feasible, it only considers the output-level supervision, but neglects the potential representation-level regularization. In fact, we can guide the composed response representation learning with the ground truth response representation. Therefore, we devise the representation-regularized response generation component with two key modules: representation-level semantic regularization and semantic-enhanced response generation. The former aims to promote the composed response representation learning with a semantic regularization between the composed response representation and the ground truth response representation. The latter targets at generating the response with not only the original

multi-level knowledge composed response representation but also the regularized composed response semantic representation.

3.4.1 Representation-level Semantic Regularization. To conduct the representation-level semantic regularization, we first introduce a set of to-be-learned latent variable vectors as queries to interact with the multi-level knowledge composed response representation and the ground truth response representation, respectively, with the goal of projecting them into the same semantic space and deriving their semantic representations. Let $\mathbf{P}_g = \{\mathbf{p}_q^1, \mathbf{p}_q^2, \cdots, \mathbf{p}_q^{N_P}\} \in$ $\mathbb{R}^{N_P \times D}$ denote the latent variable matrix with N_P variable vectors. Then for deriving the multi-level knowledge composed response semantic representation, we first employ the cross-attention mechanism to distinguish informative representation dimensions, where P_q is regarded as the query, while the composed response representation T_c is treated as the key and value. Subsequently, we utilize the multi-layer perceptron (MLP) [38] and the residual operation to further enhance the representation generalization and get the final composed response semantic representation as follows,

$$\begin{aligned} \left(\mathbf{Q}_c = \mathbf{P}_g \mathbf{W}_Q^c, \mathbf{K}_c = \mathbf{T}_c \mathbf{W}_K^c, \mathbf{V}_c = \mathbf{T}_c \mathbf{W}_V^c, \\ \bar{\mathbf{T}}_c = softmax(\mathbf{Q}_c \mathbf{K}_c^\top) \mathbf{V}_c, \\ \bar{\mathbf{T}}_c = \bar{\mathbf{T}}_c + f(\bar{\mathbf{T}}_c), \end{aligned}$$

$$\end{aligned}$$

$$\tag{8}$$

where the query \mathbf{Q}_c is projected from the \mathbf{P}_g , while the key \mathbf{K}_c and the value \mathbf{V}_c are projected from the multi-level knowledge composed response representation. $\bar{\mathbf{T}}_c$ is the intermediate composed response semantic representation. \mathbf{W}_Q^c , \mathbf{W}_K^c , and \mathbf{W}_V^c are to-be-learned weight matrices. $f(\cdot)$ refers to the MLP network. $\tilde{\mathbf{T}}_c$ is the final composed response semantic representation.

As for the ground truth response, we first obtain its embedding matrix $\mathbf{E}_r \in \mathbb{R}^{N_R \times D}$ by the position-wise embedding layer of BART, where N_R is the total number of its tokens. We then extract its representation with the BART encoder \mathcal{B}_e as follows,

$$\mathbf{T}_r = \mathcal{B}_e(\mathbf{E}_r),\tag{9}$$

where $\mathbf{T}_r \in \mathbb{R}^{N_R \times D}$ stands for the representation of the ground truth response. Thereafter, similar to the composed response semantic representation extraction in Eqn.(8), we resort to the cross-attention mechanism, where we treat \mathbf{P}_g as the query, and the ground truth response representation \mathbf{T}_r as both key and value. Let $\tilde{\mathbf{T}}_r$ be the obtained semantic representation of the ground truth response via the cross-attention mechanism.

Subsequently, to promote the composed response representation learning towards the response generation, we adopt the Frobenius norm to regularize the composed response semantic representation and the ground truth response semantic representation to be as similar as possible as follows,

$$\mathcal{L}_r = \left| \left| \tilde{\mathbf{T}}_r - \tilde{\mathbf{T}}_c \right| \right|_F^2, \tag{10}$$

where $|| \cdot ||_F^2$ is the Frobenius norm.

3.4.2 Semantic-enhanced Response Generation. Considering that the user may particularly pay more attention to the entity's attributes to obtain the desired response [4], we adopt the revised BART decoder (*i.e.*, \overline{B}_d) introduced by [2] as our decoder, which can distinguish the informative tokens of the related attribute knowledge and adaptively utilize the knowledge to promote the

textual response generation. Compared with the standard BART decoder, the revised BART decoder additionally introduces a dot-product knowledge-decoder sub-layer between the masked multi-head self-attention mechanism sub-layer and multi-head encoder-decoder attention mechanism sub-layer. To be specific, we feed the knowledge composed response representation T_c and the attribute knowledge embedding E_k into $\overline{\mathcal{B}}_d$ as follows,

$$\bar{\mathbf{z}}_{j}^{dec} = \bar{\mathcal{B}}_{d}(\mathbf{T}_{c}, \mathbf{E}_{k}, \tilde{y}_{1}, \tilde{y}_{2}, \cdots, \tilde{y}_{j-1}), \tag{11}$$

where \bar{z}_{j}^{dec} is the latent representation for generating the *j*-th token learned by the revised BART decoder $\bar{\mathcal{B}}_{d}$.

Thereafter, different from previous studies that directly predict the response token distribution based on \bar{z}_j^{dec} , we further incorporate the regularized composed response semantic representation to promote the textual response generation. Considering that different semantic dimensions can contribute differently towards the response generation, we also employ the cross-attention mechanism to obtain the reorganized composed response semantic representation for \bar{z}_j^{dec} . Specifically, we treat the latent representation \bar{z}_j^{dec} as the query, and the composed response semantic representation \tilde{T}_c as the key and value as follows,

$$\begin{cases} \mathbf{q}_{d} = \bar{\mathbf{z}}_{j}^{dec^{\top}} \mathbf{W}_{Q}^{d}, \mathbf{K}_{d} = \tilde{\mathbf{T}}_{c} \mathbf{W}_{K}^{d}, \mathbf{V}_{d} = \tilde{\mathbf{T}}_{c} \mathbf{W}_{V}^{d}, \\ \hat{\mathbf{t}}_{c} = softmax(\mathbf{q}_{d}^{\top} \mathbf{K}_{d}^{\top}) \mathbf{V}_{d}, \end{cases}$$
(12)

where \mathbf{q}_d is the query vector, while \mathbf{K}_d and \mathbf{V}_d are the key and value matrices, respectively. \mathbf{W}_Q^d , \mathbf{W}_K^d , and \mathbf{W}_V^d are weight matrices. $\hat{\mathbf{t}}_c$ is the refined composed response semantic representation. Thereafter, we can obtain the semantic-enhanced latent representation by fusing $\bar{\mathbf{z}}_i^{dec}$ and $\hat{\mathbf{t}}_c$ as follows,

$$\hat{\mathbf{z}}_{j}^{dec} = LN(\bar{\mathbf{z}}_{j}^{dec} + \hat{\mathbf{t}}_{c}), \tag{13}$$

where \hat{z}_{j}^{dec} is the semantic-enhanced latent representation for generating the *j*-th response token. Specifically, we can derive the *j*-th token probability distribution based on \hat{z}_{i}^{dec} as follows,

$$\hat{\mathbf{y}}_j = softmax(\hat{\mathbf{z}}_j^{dec\,\top} \mathbf{W}_y + \mathbf{b}_y), \tag{14}$$

where \mathbf{W}_{y} and \mathbf{b}_{y} denote the weight matrix and bias vector, respectively. $\hat{\mathbf{y}}_{j}$ represents the predicted probability distribution for the *j*-th token of the response. Notably, to avoid error accumulation, in practice, we utilize the semantic representation of ground truth response $\tilde{\mathbf{T}}_{r}$ instead of the composed response semantic representation $\tilde{\mathbf{T}}_{c}$ during the training phase.

Ultimately, following existing methods [16, 22], we use the cross entropy loss [12] to fulfill the output-level supervision as follows,

$$\mathcal{L}_{CE} = -\frac{1}{N_R} \sum_{n=1}^{N_R} log(\hat{\mathbf{y}}_n[t*]),$$
(15)

where $\tilde{\mathbf{y}}_n[t^*]$ denotes the element of $\tilde{\mathbf{y}}$ that corresponds to the *n*-th token of the ground truth response \mathcal{R} , and N_R is the number of tokens in \mathcal{R} . Notably, the loss is defined for a single sample. In the end, we can reach the final objective function for the textual response generation as follows,

$$\mathcal{L} = \lambda \mathcal{L}_{CE} + \gamma \mathcal{L}_r + \beta ||\Theta_F||_F^2, \tag{16}$$

where λ , γ , and β are the non-negative hyper-parameters. Θ_F denotes the set of parameters of the proposed MDS-S².

4 EXPERIMENT

4.1 Experiment Setting

Dataset. Previous studies mainly evaluate their models on the public dataset MMD constructed by Saha et al. [27] from the fashion domain. However, similar to [2], we did not evaluate our model on this dataset owing to the fact that MMD only allows the attribute knowledge acquisition. In contrast, we utilized the public dataset MMConv [15], which supports the knowledge acquisition from both the attribute and relation perspectives. The MMConv dataset consists of 5, 106 conversations between users and agents covering five domains: *Food, Hotel, Nightlife, Shopping mall,* and *Sightseeing.* Thereinto, the number of single-modality and multi-modality dialogues are 751 and 4, 355, respectively, where the average number of turns are 7.1 and 7.9. Besides, the knowledge base contains 1, 771 knowledge entities, where each entity is equipped with a set of attribute-value pairs and a few images. The average number of attribute-value pairs and images are 13.7 and 64.3, respectively.

Implementation Details. Following the original setting in MMConv, we divided dialogues into three chunks: 3,500 for training, 606 for validation, and 1,000 for testing. Similar to existing studies [2, 4], we regarded each utterance of agents as a ground truth response and employed its former two-turn utterances as the given context. We adopted the pretrained BART-large⁴ [35] model with 12 layers for the encoder and decoder, respectively. As for optimization, we utilized the adaptive moment estimation optimizer and settled the learning rate as 1*e*-5. In addition, we fine-tuned the proposed MDS-S² based on the training and validation dataset with 100 epochs, and reported the performance on the testing dataset. Besides, we implemented our model by Pytorch [24] and deployed all experiments on a server equipped with 8 NVIDIA 3090 GPUs. Following existing studies [2, 22], we utilized BLEU-*N* [23] where *N* ranges from 1 to 4, and Nist [5] as evaluation metrics.

4.2 Model Comparison (RQ1)

To verify the effectiveness of our proposed MDS-S², we chose the following state-of-the-art models on multimodal task-oriented dialog systems as baselines.

- MHRED [27] is the first study on the multimodal task-oriented dialog systems with a hierarchical encoder and a decoder. Thereinto, the hierarchical encoder consists of two levels of the gated recurrent units (GRU) [3], modeling the utterance and context, respectively. This baseline does not consider the knowledge as well as the representation-level regularization.
- **KHRED** [2] is extended from MHRED by integrating the attribute knowledge. Specifically, this method employs the memory network to encode the attribute knowledge, and then uses the GRU-based decoder to generate the textual response.
- LARCH [21] employs a multimodal hierarchical graph to encode the given dialog context, where each word, image, sentence, utterance, dialog pair, and the session is regarded as a node. In addition, considering the pivotal role of knowledge in multimodal

Table 1: Performance comparison among different methods in terms of BLEU-N (%) and Nist. "Improve.[†]": the relative improvement by our model over the best baseline. The best results are in **boldface**, and the second best are underlined.

| Models | BLEU-1 | BLEU-2 | BLEU-3 | BLEU-4 | Nist |
|--------------------|--------|--------|--------|--------|--------|
| MHRED | 15.02 | 6.66 | 4.24 | 2.94 | 0.9529 |
| KHRED | 18.29 | 8.28 | 4.98 | 3.36 | 1.1189 |
| LARCH | 20.86 | 11.33 | 7.58 | 5.58 | 1.3400 |
| MATE | 30.45 | 22.06 | 17.05 | 13.41 | 2.3426 |
| UMD | 31.14 | 21.87 | 17.12 | 13.82 | 2.5290 |
| TREASURE | 34.75 | 24.82 | 18.67 | 14.53 | 2.4398 |
| DKMD | 39.59 | 31.95 | 27.26 | 23.72 | 4.0004 |
| MDS-S ² | 41.40 | 32.91 | 27.74 | 23.89 | 4.2142 |
| Improve. | 4.57% | 3.00% | 1.76% | 0.72% | 5.34% |

dialog systems, this method integrates the attribute knowledge with a memory network.

- MATE [16] first resorts to the Transformer network to explore the semantic relation between the textual context and the visual context, and thus enhances the context representation. Thereafter, the method utilizes the Transformer-based decoder to generate the textual response.
- UMD [4] utilizes a hierarchy-aware tree encoder to capture the taxonomy-guided attribute-level visual representation, and a multimodal factorized bilinear pooling layer to obtain the utterance representation.
- **TREASURE** [37] presents an attribute-enhanced textual encoder, which integrates the attribute knowledge into the utterance representation. Besides, the method adopts a graph attention network to capture the semantic relation among utterances and obtain the context representation.
- **DKMD** [2] presents a dual knowledge-enhanced generative pretrained language model, where BART is adopted as the backbone. In particular, the method only explores the textual and visual context related attribute knowledge, overlooking the relation knowledge and the representation-level regularization.

Table 1 summarizes the performance comparison among different methods with respect to different evaluation metrics. From this table, we can draw the following observations. 1) Our proposed MDS-S² consistently surpasses all the baselines, exhibiting the superiority of the proposed network. In a sense, this suggests that it is reasonable to integrate the multi-level knowledge composition as well as the representation-level regularization. 2) Our proposed MDS-S² outperforms all the baselines that only consider the attribute knowledge (i.e., DKMD, TREASURE, UMD, LARCH, and KHRED), which indicates the advantage of simultaneously incorporating both the attribute knowledge and relation knowledge in multimodal task-oriented dialog systems. 3) MHRED gets the worst performance compared to other methods. This may be due to the fact that MHRED overlooks both the dual semantic knowledge (i.e., attribute and relation knowledge) and the representation-level regularization. And 4) both MDS-S² and DKMD exceed other baselines, which confirms the benefit of exploiting the generative language model in the context of multimodal task-oriented dialog systems.

⁴https://huggingface.co/facebook/bart-large.



Figure 3: Comparison between our MDS-S² and DKMD on two testing dialog pairs. "GT-R" refers to the ground truth response.

To intuitively verify the effectiveness of the proposed MDS-S², we randomly selected two testing dialog pairs, and exhibited the responses generated by the MDS-S² and the best baseline DKMD due to the space limitation in Figure 3. As can be seen, our proposed MDS-S² outperforms DKMD in the *case 1* that may involve the relation knowledge. Meanwhile, In case 2 that does not need the complicated relation knowledge, we found that our proposed MDS-S² can generate the appropriate response, while DKMD fails. One possible explanation is that our proposed MDS-S² can promote the composed response representation learning with the representation-level regularization.

4.3 On Dual Semantic Knowledge (RQ2)

To explore the roles of dual semantic knowledge in multimodal dialog systems, we designed the following five derivations. 1) w/o-rel. To illustrate the importance of the relation knowledge, we only kept the shallow attribute knowledge composition and utilized the composed response representation T_t in Eqn.(3) as the input for Eqn.(11). 2) w/o-att. To verify the importance of the attribute knowledge, we removed both the shallow attribute knowledge composition and the attribute knowledge from Eqn.(11) by using the standard BART decoder. 3) w/o-att-com. To exhibit the necessity of the shallow attribute knowledge composition, we disabled the attribute knowledge in Eqn.(3). 4) w/o-att-dec. To demonstrate the necessity of incorporating the attribute knowledge into the decoder, we replaced the knowledge revised decoder (i.e., $\bar{\mathcal{B}}_d$) with the original decoder of BART in Eqn.(11). 5) w/o-dual-k. To illustrate the role of knowledge towards the textual response generation, we removed all the attribute and relation knowledge from our proposed MDS-S².

Table 2 demonstrates the performance comparison between our proposed MDS-S² and its above derivations. From this table, we had the following observations. 1) Our proposed MDS-S² outperforms w/o-rel, w/o-att, and w/o-dual-k. Besides, disabling all knowledge (*i.e*, w/o-dual-k) results in the worst performance. It exhibits that removing either the attribute or the relation knowledge will hurt the performance of MDS-S² to some extent. This reconfirms the superiority of considering the dual semantic knowledge in multimodal task-oriented dialog systems. 2) Both w/o-att and w/o-att-com perform worse than w/o-rel, suggesting that the attribute knowledge contributes more to textual response generation than the relation knowledge. This may be due to the fact that in most cases, users want to learn the attribute information of certain known entities, and the attribute knowledge is enough for responding such cases. And 3) our proposed MDS-S² surpasses both w/o-att-com and w/o-att-dec. The rationale behind is that the attribute knowledge integrated in the shallow attribute knowledge composition can facilitate the composed response representation learning, while that in the decoder is able to explicitly incorporate the attribute knowledge in the decoder phrase, both benefiting the final textual response generation.

To obtain deeper insights into the relation knowledge, we performed the case study on the relation knowledge confidence assignment with a testing multimodal dialog in Figure 4. Due to the limited space, we only showed a few essential relation tuples related to the given context. Different color shadows correspond to different relation tuples. As we can see, towards the textual response generation, MDS-S² assigns the highest confidence to the relation tuple "Clarke Quay $\xrightarrow{address}$ river valley rd. (north boat quay)", as compared to other tuples. Checking the given multimodal context, we found that the user provides an image and wants to find a similar place with a scenic view around Singapore River. In this case, the address of the entity (*i.e.*, "river valley rd. (north boat quay)") indicates that it may be near the river, which is instructive for the proper response generation. In light of this, the confidence

Table 2: Ablation study results on the dual semantic knowledge in terms of BLEU-N (%) and Nist.

| Models | BLEU-1 | BLEU-2 | BLEU-3 | BLEU-4 | Nist |
|--------------------|--------|--------|--------|--------|--------|
| w/o-rel | 36.33 | 27.50 | 22.60 | 19.02 | 3.5398 |
| w/o-att | 33.58 | 24.80 | 20.16 | 16.93 | 2.9815 |
| w/o-att-com | 34.82 | 26.28 | 21.62 | 18.28 | 3.3089 |
| w/o-att-dec | 39.54 | 31.68 | 26.97 | 23.46 | 4.0662 |
| w/o-dual-k | 32.65 | 24.30 | 19.85 | 16.69 | 2.9856 |
| MDS-S ² | 41.40 | 32.91 | 27.74 | 23.89 | 4.2142 |



Figure 4: Illustration of the learned confidences for the relation knowledge.

assignment of our model regarding relation tuples for the given multimodal context is reasonable. This suggests that our model can identify the informative relation tuples to enhance the textual response generation in multimodal task-oriented dialog systems.

4.4 On Representation-level Regularization (RQ3)

To thoroughly verify the effectiveness of the representation-level regularization, we designed two derivatives. 1) **w/o-regular**. To verify the importance of the representation-level regularization, we removed it from our model. Consequently, the composed response semantic representation will not be used. 2) **w/o-sem-dec**. To exhibit the benefit of injecting the multi-level knowledge composed response semantic representation into the BART decoder, we kept the representation-level regularization but disabled the composed response semantic representation integration in the decoder.

Table 3 summarizes the performance of our MDS-S² and its derivatives. As can be seen, our proposed MDS-S² outperforms w/o-regular, which suggests the necessity of conducting the representation-level regularization in the context of multimodal task-oriented dialog systems. Besides, we found that w/o-sem-dec performs worse than our proposed MDS-S². This may be due to that the composed response semantic representation integrated into the decoder phrase can directly promote the textual response generation and thus enhance the performance. Thirdly, w/o-regular underperforms w/o-sem-dec, revealing that it is reasonable to conduct the representation-level regularization, which contributes to regularize the composed response semantic representation to be similar to the ground truth response semantic representation.

To intuitively reflect the effectiveness of the representation-level regularization, we randomly sampled 2,000 multimodal dialogs, and visualized the multi-level knowledge composed response representation as well as the ground truth response representation learned by our MDS-S² and w/o-regular with the help of tSNE [26] in Figure 5. The red points illustrate the composed response

Table 3: Ablation study results on the representation-level regularization in terms of BLEU-N (%) and Nist.

| Models | BLEU-1 | BLEU-2 | BLEU-3 | BLEU-4 | Nist |
|--------------------|--------|--------|--------|--------|--------|
| w/o-regular | 39.13 | 31.28 | 26.54 | 23.02 | 3.9602 |
| w/o-sem-dec | 37.63 | 29.39 | 24.50 | 20.91 | 3.6561 |
| MDS-S ² | 41.40 | 32.91 | 27.74 | 23.89 | 4.2142 |

SIGIR '23, July 23-27, 2023, Taipei, Taiwan



Figure 5: Visualization of the composed response representation distribution (red points) as well as the ground truth response representation distribution (blue points) learned by our MDS-S² and its derivative w/o-regular.

representation, and the blue ones denote the ground truth response representation. As we can see, the distribution of the multi-level knowledge composed response representation and that of the ground truth response representation achieved by MDS-S² is more consistent as compared to that obtained by w/o-regular. This validates the effect of our proposed representation-level regularization in capturing the meaningful information from the composed response representation towards the textual response generation. In addition, as for our MDS-S², we found that there is still a separate region where the two representation distributions are not well aligned. Checking these samples, we learned that they tend to involve open questions (*e.g.*, "any tips when visiting there?"), for which the effect of the dual semantic knowledge is limited.

5 CONCLUSION AND FUTURE WORK

In this work, we investigate the textual response generation task in multimodal task-oriented dialog systems and propose a novel multimodal dialog system, named MDS-S². Extensive experiments on a public dataset have validated the effectiveness of the proposed MDS-S². Interestingly, we observe that the attribute knowledge and the relation knowledge are both conducive to the textual response generation. Besides, the representation-level regularization does help in guiding the composed response representation learning with the ground truth response and should be taken into account. As aforementioned, the public MMConv dataset covers dialogs of multiple domains (e.g., Food, Hotel, and Shopping mall). Currently, we did not explore the domain information of each dialog towards the textual response generation. In the future, we plan to explore the semantic transition among different domain topics in the multimodal context and further enhance the response generation performance of multimodal dialog systems.

ACKNOWLEDGMENTS

This work is supported by the National Key Research and Development Project of New Generation Artificial Intelligence, No.:2018AAA0102502, the Shandong Provincial Natural Science Foundation (No.:ZR2022YQ59), the National Natural Science Foundation of China, No.:U1936203 and No.:62236003; Shenzhen College Stability Support Plan, No.: GXWD20220817144428005. Xiaolin Chen was supported by the China Scholarship Council for 1 year study at the National University of Singapore. SIGIR '23, July 23-27, 2023, Taipei, Taiwan

Xiaolin Chen, Xuemeng Song*, Yinwei Wei, Liqiang Nie*, and Tat-Seng Chua

REFERENCES

- Hardik Chauhan, Mauajama Firdaus, Asif Ekbal, and Pushpak Bhattacharyya. 2019. Ordinal and Attribute Aware Response Generation in a Multimodal Dialogue System. In Proceedings of the Conference of the Association for Computational Linguistics. Association for Computational Linguistics, 5437–5447.
- [2] Xiaolin Chen, Xuemeng Song, Liqiang Jing, Shuo Li, Linmei Hu, and Liqiang Nie. 2022. Multimodal Dialog Systems with Dual Knowledge-enhanced Generative Pretrained Language Model. *CoRR* (2022).
- [3] Junyoung Chung, Çaglar Gülçehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. CoRR (2014).
- [4] Chen Cui, Wenjie Wang, Xuemeng Song, Minlie Huang, Xin-Shun Xu, and Liqiang Nie. 2019. User Attention-guided Multimodal Dialog Systems. In Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM, 445–454.
- [5] George Doddington. 2002. Automatic Evaluation of Machine Translation Quality Using N-Gram Co-Occurrence Statistics. In Proceedings of the International Conference on Human Language Technology Research. Morgan Kaufmann Publishers Inc., 138–145.
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In Proceedings of the International Conference on Learning Representations. OpenReview.net.
- [7] Jing Du, Lina Yao, Xianzhi Wang, Bin Guo, and Zhiwen Yu. 2022. Hierarchical Task-aware Multi-Head Attention Network. In Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM, 1933–1937.
- [8] Wanwei He, Yinpei Dai, Min Yang, Jian Sun, Fei Huang, Luo Si, and Yongbin Li. 2022. Unified Dialog Model Pre-training for Task-Oriented Dialog Understanding and Generation. In Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM, 187–200.
- [9] Weidong He, Zhi Li, Dongcai Lu, Enhong Chen, Tong Xu, Baoxing Huai, and Jing Yuan. 2020. Multimodal Dialogue Systems via Capturing Context-Aware Dependencies of Semantic Elements. In Proceedings of the ACM International Conference on Multimedia. ACM, 2755–2764.
- [10] Wenqiang Lei, Xisen Jin, Min-Yen Kan, Zhaochun Ren, Xiangnan He, and Dawei Yin. 2018. Sequicity: Simplifying Task-oriented Dialogue Systems with Single Sequence-to-Sequence Architectures. In Proceedings of the Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics, 1437–1447.
- [11] Wenqiang Lei, Yao Zhang, Feifan Song, Hongru Liang, Jiaxin Mao, Jiancheng Lv, Zhenglu Yang, and Tat-Seng Chua. 2022. Interacting with Non-Cooperative User: A New Paradigm for Proactive Dialogue Policy. In Proceedings of International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM, 212–222.
- [12] Chun Hung Li and C. K. Lee. 1993. Minimum cross entropy thresholding. Pattern Recognition 26, 4 (1993), 617–625.
- [13] Jiao Li, Xing Xu, Wei Yu, Fumin Shen, Zuo Cao, Kai Zuo, and Heng Tao Shen. 2021. Hybrid Fusion with Intra- and Cross-Modality Attention for Image-Recipe Retrieval. In Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM, 244–254.
- [14] Yanran Li, Wenjie Li, and Zhitao Wang. 2021. Graph-Structured Context Understanding for Knowledge-grounded Response Generation. In Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM, 1930–1934.
- [15] Lizi Liao, Le Hong Long, Zheng Zhang, Minlie Huang, and Tat-Seng Chua. 2021. MMConv: An Environment for Multimodal Conversational Search across Multiple Domains. In Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM, 675–684.
- [16] Lizi Liao, Yunshan Ma, Xiangnan He, Richang Hong, and Tat-Seng Chua. 2018. Knowledge-aware Multimodal Dialogue Systems. In Proceedings of the ACM Multimedia Conference on Multimedia Conference. ACM, 801–809.
- [17] Fan Liu, Huilin Chen, Zhiyong Cheng, Anan Liu, Liqiang Nie, and Mohan Kankanhalli. 2022. Disentangled Multimodal Representation Learning for Recommendation. *IEEE Transactions on Multimedia* (2022), 1–11.
- [18] Fan Liu, Zhiyong Cheng, Lei Zhu, Zan Gao, and Liqiang Nie. 2021. Interest-Aware Message-Passing GCN for Recommendation. In Proceedings of the Web Conference 2021. Association for Computing Machinery, 1296–1305.
- [19] Zhiyuan Ma, Jianjun Li, Guohui Li, and Yongjing Cheng. 2022. UniTranSeR: A Unified Transformer Semantic Representation Framework for Multimodal Task-Oriented Dialog System. In Proceedings of the Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics, 103–114.
- [20] Zhiyuan Ma, Jianjun Li, Guohui Li, and Yongjing Cheng. 2022. UniTranSeR: A Unified Transformer Semantic Representation Framework for Multimodal Task-Oriented Dialog System. In Proceedings of the Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics, 103–114.

- [21] Liqiang Nie, Fangkai Jiao, Wenjie Wang, Yinglong Wang, and Qi Tian. 2021. Conversational Image Search. *IEEE Transactions on Image Processing* 30 (2021), 7732–7743.
- [22] Liqiang Nie, Wenjie Wang, Richang Hong, Meng Wang, and Qi Tian. 2019. Multimodal Dialog System: Generating Responses via Adaptive Decoders. In Proceedings of the ACM International Conference on Multimedia. ACM, 1098–1106.
- [23] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In Proceedings of the Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics, 311–318.
- [24] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. Curran Associates Inc.
- [25] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In Proceedings of the International Conference on Machine Learning. PMLR, 8748–8763.
- [26] Paulo E. Rauber, Alexandre X. Falcão, and Alexandru C. Telea. 2016. Visualizing Time-Dependent Data Using Dynamic t-SNE. In Proceedings of the Eurographics Conference on Visualization. Eurographics Association, 73–77.
- [27] Amrita Saha, Mitesh M. Khapra, and Karthik Sankaranarayanan. 2018. Towards Building Large Scale Multimodal Domain-Aware Conversation Systems. In Proceedings of the AAAI Conference on Artificial Intelligence. AAAI Press, 696–704.
- [28] Clemencia Siro, Mohammad Aliannejadi, and Maarten de Rijke. 2022. Understanding User Satisfaction with Task-oriented Dialogue Systems. In Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM, 2018–2023.
- [29] Liqiang Song, Mengqiu Yao, Ye Bi, Zhenyu Wu, Jianming Wang, Jing Xiao, Juan Wen, and Xin Yu. 2021. LS-DST: Long and Sparse Dialogue State Tracking with Smart History Collector in Insurance Marketing. In Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event. ACM, 1960–1964.
- [30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In Proceedings of the Advances in Neural Information Processing Systems. 5998–6008.
- [31] Wenjie Wang, Xinyu Lin, Fuli Feng, Xiangnan He, Min Lin, and Tat-Seng Chua. 2022. Causal Representation Learning for Out-of-Distribution Recommendation. In WWW. 3562–3571.
- [32] Yinwei Wei, Xiang Wang, Weili Guan, Liqiang Nie, Zhouchen Lin, and Baoquan Chen. 2019. Neural multimodal cooperative learning toward micro-video understanding. *IEEE Transactions on Image Processing* 29 (2019), 1–14.
- [33] Yinwei Wei, Xiang Wang, Liqiang Nie, Shaoyu Li, Dingxian Wang, and Tat-Seng Chua. 2022. Causal Inference for Knowledge Graph based Recommendation. IEEE Transactions on Knowledge and Data Engineering (2022).
- [34] Haokun Wen, Xuemeng Song, Xin Yang, Yibing Zhan, and Liqiang Nie. 2021. Comprehensive Linguistic-Visual Composition Network for Image Retrieval. In Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM, 1369--1378.
- [35] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos, Online, November 16-20, 2020. Association for Computational Linguistics, 38–45.
- [36] Shiquan Yang, Rui Zhang, Sarah M. Erfani, and Jey Han Lau. 2021. UniMF: A Unified Framework to Incorporate Multimodal Knowledge Bases intoEnd-to-End Task-Oriented Dialogue Systems. In Proceedings of the International Joint Conference on Artificial Intelligence. ijcai.org, 3978–3984.
- [37] Haoyu Zhang, Meng Liu, Zan Gao, Xiaoqiang Lei, Yinglong Wang, and Liqiang Nie. 2021. Multimodal Dialog System: Relational Graph-based Context-aware Question Understanding. In Proceedings of the ACM Multimedia Conference. ACM, 695–703.
- [38] Zhengyou Zhang, Michael J. Lyons, Michael Schuster, and Shigeru Akamatsu. 1998. Comparison Between Geometry-Based and Gabor-Wavelets-Based Facial Expression Recognition Using Multi-Layer Perceptron. In Proceedings of the International Conference on Face & Gesture Recognition. IEEE Computer Society, 454–461.