

Which One Are You Referring To? Multimodal Object Identification in Situated Dialogue

Holy Lovenia*, Samuel Cahyawijaya*, Pascale Fung
Center for Artificial Intelligence Research (CAiRE),
The Hong Kong University of Science and Technology
{hlovenia, scahyawijaya}@connect.ust.hk

Abstract

The demand for multimodal dialogue systems has been rising in various domains, emphasizing the importance of interpreting multimodal inputs from conversational and situational contexts. One main challenge in multimodal dialogue understanding is multimodal object identification, which constitutes the ability to identify objects relevant to a multimodal user-system conversation. We explore three methods to tackle this problem and evaluate them on the largest situated dialogue dataset, SIMMC 2.1. Our best method, scene-dialogue alignment, improves the performance by $\sim 20\%$ F1-score compared to the SIMMC 2.1 baselines. We provide analysis and discussion regarding the limitation of our methods and the potential directions for future works. Our code is publicly available at <https://github.com/holylovenia/multimodal-object-identification>.

1 Introduction

Recent advancements in multimodal dialogue systems have gained more traction in various domains such as retail, travel, fashion, interior design, and many others. A real-world application of multimodal dialogue systems is situated dialogue, where a dialogue agent shares a co-observed vision or physical space with the user, and is responsible for handling user requests based on the situational context, which are often about the objects in their surroundings. This makes multimodal object identification from a dialogue (i.e., identifying objects that fit a dialogue context) an indispensable skill in multimodal dialogue understanding, built on cross-modal understanding to comprehend the relations between linguistic expressions and visual cues.

Various methods have been proposed to perform multimodal object identification through different paradigms (Yu et al., 2016; Hu et al., 2016; Ilinykh

*Equal contribution.




Figure 1: Multimodal object identification is the fundamental step required to enable multimodal dialogue systems to understand the object referred to by the user. Image is adapted from (Kottur et al., 2021).

et al., 2019; Kamath et al., 2021; Kuo and Kira, 2022). These efforts have established remarkable progress in solving this problem. However, aside from an observed gap between the performance of the existing works and human-level performance in multimodal object identification, prior works also rely on a presumption that the information given by the textual context will only lead to specific (i.e., unambiguous) objects, which does not conform to real-world multimodal conversations where ambiguity exists.

Therefore, in this work, we explore three different solutions to enable multimodal object identification in the situated dialogue system, i.e., dialogue-contextualized object detection, object-dialogue alignment, and scene-dialogue alignment, without adopting the unambiguity assumption. Dialogue-contextualized object detection utilizes the spatial and object understanding capability of a pre-trained object detection model, to generate semantic representation containing both visual cues and the spatial understanding of the object. Object-dialogue alignment incorporates the image-text alignment capability of CLIP (Radford et al., 2021), which has been pre-trained on large image-text corpora to perform multimodal object identification from the given dialogue context. Scene-object alignment

combines the spatial and object understanding capability of a pre-trained object detection model and a pre-trained textual understanding model to produce better semantic vision-language alignment.

Our contributions are three-fold:

- We introduce three different methods for handling multimodal object identification in situated dialogue, i.e., dialogue-contextualized object detection, object-dialogue alignment, and scene-dialogue alignment;
- We show the dialogue-contextualized object detection method fails to outperform even the heuristic baselines despite having an acceptable performance on the object detection task;
- We show the effectiveness of the other two methods which significantly outperform the SIMMC 2.1 baselines by $\sim 5\%$ F1-score for object-dialogue alignment and $\sim 20\%$ F1-score for scene-dialogue alignment;

2 Related Work

Multimodal Dialogue System Multiple studies have attempted to enable the skills required for multimodal dialogue system, e.g., understanding visual (Antol et al., 2015; Das et al., 2017; Kottur et al., 2019) or visual-temporal (Alamri et al., 2019) content to answer user’s questions, grounding conversations to images (Mostafazadeh et al., 2017; Shuster et al., 2020), interpreting multimodal inputs and responding with multimodal output to assist users with their goal (Saha et al., 2018) or as a means to converse (Sun et al., 2022), and perceiving the shared environment to grasp situational context to enable proper navigation, adaptation, and communication (Lukin et al., 2018; Brawer et al., 2018; Kottur et al., 2021).

At the core of these efforts, the ability to understand language and vision, as well as integrate both representations to align the linguistic expressions in the dialogue with the relevant visual concepts or perceived objects, is the key to multimodal dialogue understanding (Landragin, 2006; Loáiciga et al., 2021b,a; Kottur et al., 2018; Utescher and Zarriß, 2021; Sundar and Heck, 2022; Dai et al., 2021).

Multimodal Object Identification Identifying objects or visual concepts related to a linguistic expression is an incremental exploration in vision-language research. It starts with identifying sim-

ple objects in a sanitized environment (Mitchell et al., 2010) based on image descriptions or captions. Then, multimodal object identification has been gradually increasing in complexity and realism by involving visual contexts with cluttered and diverse scenes (Kazemzadeh et al., 2014; Gkatzia et al., 2015; Yu et al., 2016; Mao et al., 2016; Hu et al., 2016; Ilinykh et al., 2019; Kamath et al., 2021; Kuo and Kira, 2022).

While these works base their multimodal object identification on single-turn text contexts, another line of works explores the usage of multi-turn sequences as a textual context to enable identifying objects based on implicit constraints deduced through multi-round reasoning (Seo et al., 2017; Johnson et al., 2017; Liu et al., 2019; Moon et al., 2020). However, they focus on identifying only the specific (i.e., unambiguous) objects, in which only a certain object in the scene fits the corresponding linguistic context. This is quite dissimilar from real-world multimodal object identification, where multiple objects could fit a given textual context and induce ambiguity into the conversation (Kottur et al., 2021). For this reason, existing works are not equipped with the ability to identify all objects that *plausibly* fit those constraints although this skill is required to perform multimodal object identification in situated dialogue.

Multimodal and Cross-Modal Learning Past works have studied multimodal and cross-modal alignment, grounding, and generation to solve various vision-language tasks, e.g., image captioning (Hossain et al., 2019; Sharma et al., 2018), generating stories from image (Min et al., 2021; Love- nia et al., 2022), as well as multimodal object identification (Li et al., 2019; Wang et al., 2022). These attempts become more substantial and extensive after the rise of pre-trained vision-language models such as CLIP (Radford et al., 2021), ALIGN (Jia et al., 2021), and FLAVA (Singh et al., 2022), which allows transfer knowledge obtained from the large-scale pre-training to downstream tasks.

3 Methodology

In this section, we describe the preliminaries of our work (§3.1) and extensively elaborate on each of our approaches, i.e., dialogue-contextualized object detection (§3.2), object-dialogue alignment (§3.3), and scene-dialogue alignment (§3.4).

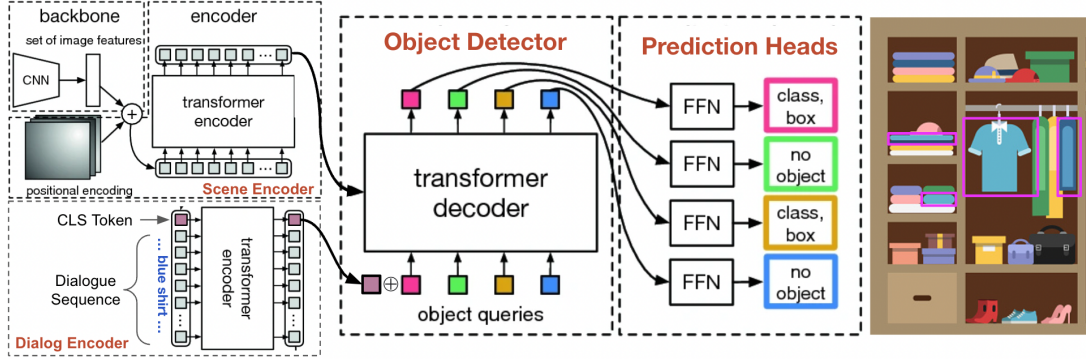


Figure 2: The architecture of SitCoM-DETR. SitCoM-DETR consists of a scene encoder and a dialogue encoder to extract multimodal content, respectively. The dialogue representation is used to guide the object detector module to judiciously filter out unrelated scene objects.

3.1 Preliminaries

The goal of multimodal object identification in situated dialogue is to identify objects from a given scene image that fulfill the user’s request gathered from the user-system interactions. To identify the object(s) that could satisfy a user’s request in a dialogue, it is crucial to match the objects and the implicit constraints interwoven in the dialogue, e.g., S: *“I do! Take a look at these. I have a brown coat towards the far end on the left wall, another brown coat on the left side of the front floor rack, and a black coat on the front of the same rack.”*, U: *“Awesome! Tell me the cost and label on that one.”*. Thus, it is essential for the system to understand the relation between the visual perception of the objects in the scenes and the natural language used to verbalize these constraints, which describe the target object(s) by visual attributes (e.g., color, object category or type, etc.), location (i.e., absolute or relative position), or the combination of both.

We define a dialogue between a user and a system as $D = \{u_1, s_1, u_2, s_2, \dots, u_n, s_n\}$, a scene consisting of images corresponding to multiple viewpoints of the scene as $\{I_1^{scene}, I_2^{scene}, \dots, I_n^{scene}\}$, and a set of objects in the scene as $O^{scene} = \{(b_1, c_1), (b_2, c_2), \dots, (b_n, c_n)\}$, where u_i and s_i respectively denote the user utterance and the system utterance, and c_i and b_i denote the bounding box and the class category of an object. Given a user dialogue turn $D_i^{user} = \{u_1, s_1, u_2, s_2, \dots, u_i\}$, $i \leq n$, and a scene image I_i^{scene} , the goal of the task is to select a subset of scene objects $O^{match} \subseteq O^{scene}$ that could satisfy the referred criteria in D_i^{user} .

3.2 Approach 1: Dialogue-Contextualized Object Detection

For dialogue-contextualized object detection, we frame the task of multimodal object identification as the contextualized object detection task. In object detection, given a scene image I^{scene} , we aim to detect all objects O^{scene} in the scene by predicting their bounding box and class category. While in contextualized object detection, the aim is instead to select only a set of scene objects O^{match} that satisfy a given context.

Our approach for dialogue-contextualized object detection extends a state-of-the-art object detection model, namely DETR (Carion et al., 2020), by injecting dialogue information as the context to guide the detection model to filter out unidentified objects. A similar solution has been proposed by Modulated DETR (MDETR) (Kamath et al., 2021). Despite its strong performance on text-contextualized object detection, MDETR requires an aligned annotation between the text phrase and the visual object for training. Such annotation is not available on SIMMC 2.1, hence we develop a new text-contextualized object detection model namely Situational Context for Multimodal DETR (SitCoM-DETR). Unlike MDETR which concatenates the textual representation along with the visual representation before feeding them into the transformer encoder of DETR (shown in Appendix A), SitCoM-DETR injects a dialogue-level semantic representation vector into the input query of the transformer decoder of DETR in order to guide the model to select objects that match the dialogue context. We incorporate the same loss functions as the original DETR model. The depiction of our SitCoM-DETR model is shown in Figure 2.

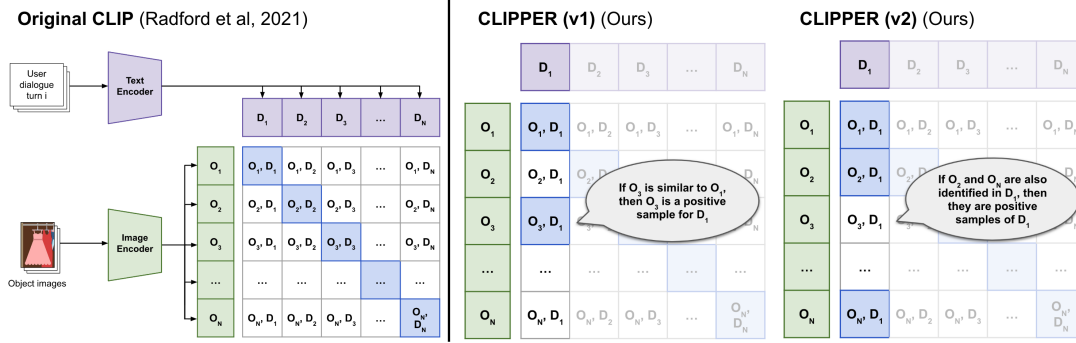


Figure 3: Learning objectives of the original CLIP (Radford et al., 2021), CLIPPER (v1), and CLIPPER (v2) for the object-dialogue alignment approach. The similarities of the positive pairs (blue) are maximized while the similarities of the negative pairs (white) are minimized.

3.3 Approach 2: Object-Dialogue Alignment

For object-dialogue alignment, we frame the task of multimodal object identification as the alignment between a target object O_i^{match} and a user dialogue turn D_i^{user} pair. Given a user dialogue turn D_i^{user} and its corresponding scene image I_i^{scene} , we first preprocess I_i^{scene} to extract the object images of O^{match} . Each of the object images is paired with D_i^{user} as the positive pairs. We obtain the visual embeddings from the image by feeding it to an image encoder, and the textual embeddings from the dialogue turn by feeding it to a text encoder. After these embeddings pass through a linear projection, we calculate the similarity using the dot product between the two resulting vectors. Utilizing the contrastive learning objective, on a batch of object-dialogue pairs, this cross-modal alignment architecture learns by maximizing the similarity of the positive pairs and minimizing the similarity of the negative pairs (Figure 3).

Object-Dialogue Similarity Learning Strategy

The original contrastive learning approaches the object-dialogue alignment task as a one-to-one function, where the positive sample of D_i is only O_i in Figure 3. This is different from the actual nature of multimodal object identification, where more than one object could be relevant to a dialogue turn. For this reason, in addition to the original contrastive learning, we explore two modifications of the learning objective, where: 1) the positive samples of D_i include O_i (image pair) and similar objects¹ to O_i ; and 2) the positive samples of D_i include O_i and other supposedly identified

¹We define similar objects to O_i as any other objects in the corresponding scene that use the same prefabricated design as O_i in the SIMMC 2.1 dataset.

objects in D_i . For simplicity, we refer to these methods as **CLIPPER (v1)** and **CLIPPER (v2)**.

3.4 Approach 3: Scene-Dialogue Alignment

For scene-dialogue alignment, we aim to combine the spatial understanding learned from object detection training with the image-text matching for multimodal similarity learning to solve multimodal object identification. For this approach, we utilize a pre-trained object detection model, i.e., DETR, and two pre-trained language models, i.e., BERT and GPT2. The resulting models are referred to as **DETR-BERT** and **DETR-GPT2**, respectively. We illustrate the overview of this approach in Figure 4.

In this approach, we first frame our dataset as an object detection task, where a data instance consists of a scene image I_i^{scene} and its object annotations $O^{scene} = \{(b_1, c_1), (b_2, c_2), \dots, (b_m, c_m)\}$, and train an object detection model (DETR) on it. The resulting model is then used to extract the visual representations of all objects in the scene image I^{scene} by matching the object queries with O^{scene} using Hungarian matching (Stewart et al., 2016).

For the next step, we frame our dataset as a binary classification task, where a data instance consists of a user dialogue turn D_i^{user} , an object O_j^{scene} in a corresponding scene I_i^{scene} , and a binary label (i.e., whether the object is identified by the user dialogue turn or not). We utilize a dialogue encoder to extract textual representation from a user dialogue turn D_i^{user} . The textual representation of D_i^{user} and the visual representation of O_j^{scene} are projected into a latent space. We compute the dot product of the two and use the resulting vector as the prediction logits for training and inference.




Figure 4: Scene-dialogue alignment. We pre-extract the visual embeddings from an object detection model trained on our dataset. The visual embeddings are used together with dialogue embeddings in the next training to perform multimodal object detection as a binary classification task.

4 Experiment

4.1 Dataset

For all of our experiments, we utilize the ambiguous candidate identification task from the SIMMC 2.1 dataset (Kottur et al., 2021). The dataset studies conversational scenarios where the system shares a co-observed vision (i.e., the same scene) with the user. The dataset focuses on improving the shopping experience in two domains: fashion and furniture. In the setting of SIMMC 2.1, the system is able to access the ground truth meta information of all objects (e.g., object price, size, material, brand, etc.) in the scene O^{scene} , while the user observes objects only through the scene viewpoints $\{I_1^{scene}, I_2^{scene}, \dots, I_n^{scene}\}$ to describe a request.

Each dialogue in the dataset can utilize different scene viewpoints at different dialogue turns throughout the session. This represents scenarios where the user navigates the scene during the interaction in a real physical store. Therefore, the multimodal dialogue system needs to understand user requests using both the dialogue history and the scene image as a unified multimodal context. The statistics of the ambiguous candidate identification of SIMMC 2.1 dataset is presented in Table 1.²

4.2 Baselines

We incorporate various baselines including simple heuristics and deep learning based multimodal

²We use the devtest split of SIMMC 2.1 dataset as the test set in our experiment.

Split	# Sample	# Dialogue	$\frac{O_{match}}{O_{scene}}$
Train	4239	3983	28.74%
Validation	414	371	24.72%
Test	940	905	30.78%

Table 1: Statistics of the ambiguous candidates identification of the SIMMC 2.1 dataset.

matching methods from SIMMC 2.1.³ For the heuristic methods, we incorporate uniform random prediction (**Random**), empty prediction (**No object**), and all objects prediction (**All objects**) as our baselines. For the deep learning approaches (**ResNet50-BERT** and **ResNet50-GPT2**), we apply cosine similarity between the feature extracted from ResNet-50 (He et al., 2016)⁴ and two widely-used pre-trained LMs, i.e., BERT (Devlin et al., 2019)⁵ and GPT2 (Radford et al., 2019)⁶.

In addition to these baselines, we incorporate several additional baselines: 1) pre-trained CLIP (Radford et al., 2021)⁷, which serves as a baseline for the object-dialogue alignment approach and 2) pre-trained MDETR (Kamath et al., 2021)⁸, which represents a text-conditioned object detection baseline trained with an explicit align-

³SIMMC 2.1 repository: <https://github.com/facebookresearch/simmc2>.

⁴We use the pre-extracted visual feature provided in the SIMMC 2.1 repository.

⁵<https://huggingface.co/bert-base-uncased>.

⁶<https://huggingface.co/gpt2>.

⁷We use the checkpoint from <https://huggingface.co/openai/clip-vit-base-patch32>.

⁸We use the EfficientNet B5 (ENB5) backbone checkpoint from <https://github.com/ashkamath/mdetr>.

ment between phrases and objects. For CLIP, we report both zero-shot (**CLIP (zero-shot)**) and direct fine-tuning (**CLIP**) performances, while for MDETR, we only use the zero-shot performance (**MDETR (zero-shot)**) due to the unavailability of the explicit alignment between objects and dialogues in the dataset.

4.3 Models

We propose three different approaches to solve the multimodal object identification task §3. For the dialogue-contextualized object detection approach, we incorporate one model, namely **SitCoM-DETR** which will be compared to the MDETR baseline. For the object-dialogue alignment approach, we incorporate two model variants, i.e., **CLIPPER (v1)** and **CLIPPER (v2)**. For the scene-object alignment approach, we incorporate two model variants, i.e., **DETR-BERT** and **DETR-GPT2**.

4.4 Evaluation

Given a label set L and a prediction set P , we define the number of true positive $N^{correct}$ as the objects that appear in both the prediction and the label sets. Using this definition, we evaluate the models’ performance on the multimodal object identification task using three evaluation metrics, i.e., recall, precision, and F1-score. The definition of each metric is defined as:

$$Recall = \frac{N^{correct}}{\|L\|} \quad (1)$$

$$Precision = \frac{N^{correct}}{\|P\|} \quad (2)$$

$$F1 = \frac{2 * Precision * Recall}{Precision + Recall} \quad (3)$$

4.5 Implementation Details

Dialogue Preprocessing In all of our experiments, following prior works in end-to-end task-oriented dialogue system, we encode the last three utterances from the dialogue into a single text. For example a user dialogue turn $D_i^{user} = \{u_1, s_1, u_2, s_2, \dots, u_i\}$ is encoded into a text "U: <u_{i-1}> S: <s_{i-1}> U: <u_i>" to be further processed by the dialogue encoder.

Inference strategy for object-dialogue alignment

For the proposed CLIPPER model in the object-dialogue alignment approach, we simply apply sigmoid to the logits and use a threshold value of 0.5 (denoted as *Sigmoid*), since it has a built-in capability to perform multi-label classification. While

for the CLIP model, which serves as a baseline, does not have the same capability, hence we use the mean value of the logits as the threshold (denoted as *Mean*). Additionally, we also evaluate the performance of the model if the top- k objects with the highest logits are considered valid predictions, where k denotes the correct amount of objects in the ground-truth label (denoted as *Oracle*).

Inference strategy for dialogue-contextualized object detection

For the dialogue-contextualized object detection, since the model is originally for the object detection task, we develop our own inference strategy to allow it to perform multi-label classification for object identification. This is done through several steps: 1) we perform Hungarian matching using all objects, 2) we compute intersection over union (IoU) of all pairs of matched prediction and ground-truth bounding boxes⁹, and 3) we take all objects having IoU score $\geq 10\%$ ¹⁰.

Hyperparameter Details

For the dialogue-contextualized object detection, we fine-tune the SitCoM-DETR model for a maximum of 200 epochs with AdamW optimizer using a linear learning rate decay, a learning rate between [1e-4..1e-5], and an early stopping of 10 epochs. For the object-dialogue alignment, we fine-tune the CLIP and CLIPPER models for a maximum of 200 epochs with AdamW optimizer using a linear learning rate decay, a learning rate between [1e-4..1e-5], and an early stopping of 10 epochs. For the scene-dialogue alignment, we fine-tune the DETR-BERT and DETR-GPT2 models for a maximum of 200 epochs with AdamW optimizer using a linear learning rate decay, a learning rate between [1e-4..1e-5], and an early stopping of 10 epochs.

5 Result and Analysis

5.1 Result Overview

The results of our experiments are shown in Table 2. The best baseline performance is achieved by **CLIP (fine-tuned)** with 45.09% F1-score outperforming the baselines provided by the SIMMC 2.1 (i.e., **ResNet50-GPT2** and **ResNet50-BERT**), showing the superiority of image-text alignment pre-training over separate unimodal pre-trainings for multimodal object identification. For the

⁹We do not consider the class label in the scoring to have a fairer comparison with the zero-shot MDETR approach.

¹⁰We align this with MDETR’s class probability setting during inference.

Method Type	Approach	Recall	Precision	F1-score
<i>Baselines</i>				
<i>Heuristic</i>	No object	0.00%	0.00%	0.00%
	Random	49.90%	22.43%	30.95%
	All objects	100.00%	22.34%	36.52%
<i>SIMMC 2.1</i>	ResNet50-GPT2	36.40%	42.26%	39.11%
	ResNet50-BERT	36.70%	43.39%	39.76%
<i>Dialogue-Contextualized Object Detection</i>	MDETR (zero-shot)	<u>16.33%</u>	<u>29.70%</u>	<u>21.07%</u>
<i>Object-Dialogue Alignment</i>	CLIP (zero-shot)	55.70%	26.39%	35.81%
	CLIP (fine-tuned)	<u>73.00%</u>	<u>32.62%</u>	45.09%
<i>Proposed Methods</i>				
<i>Dialogue-Contextualized Object Detection</i>	SitCoM-DETR (aug)	47.82%	25.69%	33.42%
	SitCoM-DETR (no aug)	<u>49.51%</u>	<u>25.81%</u>	<u>33.93%</u>
<i>Object-Dialogue Alignment</i>	CLIPPER (v1)	73.41%	<u>33.00%</u>	<u>45.53%</u>
	CLIPPER (v2)	59.95%	25.60%	35.88%
<i>Scene-Dialogue Alignment</i>	DETR-BERT	<u>65.47%</u>	51.48%	57.64%
	DETR-GPT2	63.81%	56.79%	60.10%

Table 2: Experimental results of multimodal object identification on the SIMMC 2.1 dataset (Kottur et al., 2021). **Bold** denotes the best performances of baselines and proposed methods. Underline denotes the best performances within a method type.

dialogue-contextualized object detection methods, the proposed **SitCoM-DETR** outperforms **MDETR (zero-shot)**. Nevertheless, its performance for multimodal object identification is low despite having an acceptable object detection quality. We conjecture that a better method for adapting an object detection model for multimodal object identification is required, which is also shown by our *scene-dialogue alignment* approach in §3.4.

For the object-dialogue alignment, our **CLIPPER (v1)** marginally outperforms the **CLIP (fine-tuned)** baseline. This shows the effectiveness of modifying the CLIP objective which is explained in more detail in §5.3. For the scene-dialogue alignment (i.e., **DETR-BERT** and **DETR-GPT2**), where we combine the object detection and the image-text contrastive objective, we show a significant improvement over **CLIP (fine-tuned)**, which is the highest-performing baseline, by $\sim 10\text{-}15\%$ F1-score. This suggests the importance of combining object detection representation and image-text contrastive learning to fulfill the need for both visual and spatial matching to solve multimodal object identification.

5.2 Pitfalls of the Best Performing Models

We manually analyze the incorrect predictions made by our scene-dialogue alignment approaches, i.e., **DETR-BERT** and **DETR-GPT2**. Based on our analysis in Table 5, our models encounter two main issues. First, our models have difficulties in identifying objects when faced with a sudden object shift in the dialogue, e.g., the sudden shift from beds to a chair in this user dialogue turn U: “I need a new bed too. Any suggestions?”, S: “Both of these grey beds are in stock.”, U: “What’s the rating on that chair?”.

The second issue is the ineffectiveness of handling textual coreferences. For instance, in the user dialogue turn U: “How about a hat, but cheap and in a small?”, S: “I have the black hat third from the front, the white hat at the front, and the black hat between them.”, U: “What’s the brand and reviews for the black hat?”, the models fail to recognize that “the black hat” in the user utterance is anaphoric to either “the black hat third from the front” or “the black hat between them” in the system utterance, which leads to the system’s failure to identify both black hats as O^{match} . This shortcoming also be-




Figure 5: Frequency of error types of 100 misclassified samples from **DETR-BERT** and **DETR-GPT2**.

comes more pronounced if the coreference chains are longer.

These issues show the limitation of pre-trained LMs for discourse understanding and analysis, especially in terms of coreference and entity linking (Jurafsky and Martin, 2019; Pandia et al., 2021; Koto et al., 2021). Additionally, some other cases require the model to process long-term dialogue history dependency which existing LMs are not able to handle because of the quadratic cost bottleneck of the attention mechanism of the transformer architecture (Vaswani et al., 2017). Adapting an efficient attention mechanism with linear complexity might be beneficial to mitigate this problem.

5.3 Impact of Changing CLIP Objective

As shown in Table 3, the CLIPPER models with binary cross-entropy objective have a built-in capability for multi-label classification with **Sigmoid** which consistently performs better compared to the **Mean** thresholding. In addition, **CLIPPER (v1)** outperforms the original CLIP model which is trained with the cross-entropy loss. These facts suggest that changing the CLIP objective is beneficial for performing multi-label classification tasks such as multimodal object identification.

When using **Oracle**, we can observe a significant improvement in F1-score score, which mainly comes from the improvement in the precision with only a minor degradation on recall. This suggests that there is a very sensitive range of logits which consists of many negative samples with a few positive samples. To better segregate these few positive samples from the negative ones, hard negative mining techniques such as focal loss (Lin et al., 2020) might be beneficial to alleviate this problem.

Approach	Rec.	Prec.	F1
CLIP — Cross-Entropy			
Mean	73.00%	32.62%	45.09%
Oracle	74.99%	74.96%	74.98%
CLIPPER (v1) — Binary Cross-Entropy			
Sigmoid	73.41%	33.00%	45.53%
Mean	73.08%	31.97%	44.48%
Oracle	73.37%	73.34%	73.36%
CLIPPER (v2) — Binary Cross-Entropy			
Sigmoid	59.95%	25.60%	35.88%
Mean	53.90%	23.42%	32.65%
Oracle	54.92%	54.89%	54.91%

Table 3: Results for object-dialogue alignment models with different thresholding strategies.

6 Discussion

Based on the results and analysis, we show that the *scene-object alignment* approach is the best performing approach, achieving $\sim 55\text{-}60\%$ F1-score in the multimodal object identification task of SIMMC 2.1. We analyze the behavior of the model and conjecture that existing LMs have a limitation on understanding discourse. Additionally, we show the potential benefit of modeling the long-term dependency of dialogue history to further improve the quality of multimodal object identification task (§5.2). Lastly, we analyze the limitation of the existing image-text contrastive approaches for multimodal object identification and propose an alternative objective to alleviate this limitation (§5.3).

For future work, we aim to focus on the scene-dialogue alignment methods to further improve the model performance on the multimodal object identification capability. We note five potential points of improvement that can be further explored to improve the model performance in multimodal object identification: 1) the incorporation of cross-object attention in the modality fusion phase to enable a better relative position understanding between objects, 2) the incorporation of linear attention mechanism to handle the long-term dependency of dialogue history, 3) the exploration on better contrastive objectives for multimodal object identification, 4) the exploration on improving discourse understanding for LMs to better handle coreference and sudden object shift, and 5) the synthetic scene-dialogue data augmentation through the utilization of other publicly available object detection datasets to handle the in-domain data scarcity problem.

7 Conclusion

In this paper, we explore three methods to tackle multimodal object identification and evaluate them on SIMMC 2.1. Our best method, scene-dialogue alignment, improves the performance by $\sim 20\%$ F1-score compared to the SIMMC 2.1 baselines. We provide an analysis of incorrect predictions by our best approach and the impact of changing the CLIP learning objective. We further provide discussion regarding the limitation of our methods and the potential directions for future works.

Acknowledgement

We appreciate the guidance that Prof. Dan Xu has provided for this research. This work has been supported by the School of Engineering PhD Fellowship Award, the Hong Kong University of Science and Technology and PF20-43679 Hong Kong PhD Fellowship Scheme, Research Grant Council, Hong Kong.

References

- Huda Alamri, Vincent Cartillier, Abhishek Das, Jue Wang, Anoop Cherian, Irfan Essa, Dhruv Batra, Tim K Marks, Chiori Hori, Peter Anderson, et al. 2019. Audio visual scene-aware dialog. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7558–7567.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433.
- Jake Brawer, Olivier Mangin, Alessandro Roncone, Sarah Widder, and Brian Scassellati. 2018. Situated human–robot collaboration: predicting intent from grounded natural language. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 827–833. IEEE.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *Computer Vision – ECCV 2020*, pages 213–229, Cham. Springer International Publishing.
- Wenliang Dai, Samuel Cahyawijaya, Zihan Liu, and Pascale Fung. 2021. **Multimodal end-to-end sparse model for emotion recognition**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5305–5316, Online. Association for Computational Linguistics.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. 2017. Visual dialog. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 326–335.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dimitra Gkatzia, Verena Rieser, Phil Bartie, and William Mackaness. 2015. **From the virtual to the RealWorld: Referring to objects in real-world spatial scenes**. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1936–1942, Lisbon, Portugal. Association for Computational Linguistics.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. **Deep residual learning for image recognition**. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- MD Zakir Hossain, Ferdous Sohel, Mohd Fairuz Shiratuddin, and Hamid Laga. 2019. A comprehensive survey of deep learning for image captioning. *ACM Computing Surveys (CSUR)*, 51(6):1–36.
- Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, and Trevor Darrell. 2016. Natural language object retrieval. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4555–4564.
- Nikolai Ilinykh, Sina Zarrieß, and David Schlangen. 2019. **Tell me more: A dataset of visual scene description sequences**. In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 152–157, Tokyo, Japan. Association for Computational Linguistics.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR.
- Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2901–2910.
- Dan Jurafsky and James H Martin. 2019. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics,*

- and *Speech Recognition (3rd draft ed.)*. Stanford Univ.
- Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. 2021. **MDETR - modulated detection for end-to-end multi-modal understanding**. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE.
- Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. 2014. **ReferItGame: Referring to objects in photographs of natural scenes**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 787–798, Doha, Qatar. Association for Computational Linguistics.
- Fajri Koto, Jey Han Lau, and Timothy Baldwin. 2021. **Discourse probing of pretrained language models**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3849–3864, Online. Association for Computational Linguistics.
- Satwik Kottur, Seungwhan Moon, Alborz Geramifard, and Babak Damavandi. 2021. **SIMMC 2.0: A task-oriented dialog dataset for immersive multimodal conversations**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4903–4912, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Satwik Kottur, José M. F. Moura, Devi Parikh, Dhruv Batra, and Marcus Rohrbach. 2019. **CLEVR-dialog: A diagnostic dataset for multi-round reasoning in visual dialog**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 582–595, Minneapolis, Minnesota. Association for Computational Linguistics.
- Satwik Kottur, José MF Moura, Devi Parikh, Dhruv Batra, and Marcus Rohrbach. 2018. Visual coreference resolution in visual dialog using neural module networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 153–169.
- Chia-Wen Kuo and Zsolt Kira. 2022. Beyond a pretrained object detector: Cross-modal textual and visual context for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17969–17979.
- Frédéric Landragin. 2006. Visual perception, language and gesture: A model for their understanding in multimodal dialogue systems. *Signal Processing*, 86(12):3578–3595.
- Zhihui Li, Lina Yao, Xiaoqin Zhang, Xianzhi Wang, Salil Kanhere, and Huaxiang Zhang. 2019. Zero-shot object detection with textual descriptions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8690–8697.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2020. **Focal loss for dense object detection**. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2):318–327.
- Runtao Liu, Chenxi Liu, Yutong Bai, and Alan L Yuille. 2019. Clevr-ref+: Diagnosing visual reasoning with referring expressions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4185–4194.
- Sharid Loáiciga, Simon Dobnik, and David Schlangen. 2021a. **Annotating anaphoric phenomena in situated dialogue**. In *Proceedings of the 1st Workshop on Multimodal Semantic Representations (MMSR)*, pages 78–88, Groningen, Netherlands (Online). Association for Computational Linguistics.
- Sharid Loáiciga, Simon Dobnik, and David Schlangen. 2021b. **Reference and coreference in situated dialogue**. In *Proceedings of the Second Workshop on Advances in Language and Vision Research*, pages 39–44, Online. Association for Computational Linguistics.
- Holy Lovenia, Bryan Wilie, Romain Barraud, Samuel Cahyawijaya, Willy Chung, and Pascale Fung. 2022. **Every picture tells a story: Image-grounded controllable stylistic story generation**. In *Proceedings of the 6th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 40–52, Gyeongju, Republic of Korea. International Conference on Computational Linguistics.
- Stephanie M. Lukin, Felix Gervits, Cory J. Hayes, Pooja Moolchandani, Anton Leuski, John G. Rogers III, Carlos Sanchez Amaro, Matthew Marge, Clare R. Voss, and David Traum. 2018. **ScoutBot: A dialogue system for collaborative navigation**. In *Proceedings of ACL 2018, System Demonstrations*, pages 93–98, Melbourne, Australia. Association for Computational Linguistics.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. 2016. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20.
- Kyungbok Min, Minh Dang, and Hyeonjoon Moon. 2021. Deep learning-based short story generation for an image using the encoder-decoder structure. *IEEE Access*, 9:113550–113557.
- Margaret Mitchell, Kees van Deemter, and Ehud Reiter. 2010. Natural reference to objects in a visual domain. In *Proceedings of the 6th international natural language generation conference*.

- Seungwhan Moon, Satwik Kottur, Paul Crook, Ankita De, Shivani Poddar, Theodore Levin, David Whitney, Daniel Difranco, Ahmad Beirami, Eunjoon Cho, Rajen Subba, and Alborz Geramifard. 2020. [Situating and interactive multimodal conversations](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1103–1121, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Nasrin Mostafazadeh, Chris Brockett, Bill Dolan, Michel Galley, Jianfeng Gao, Georgios Spithourakis, and Lucy Vanderwende. 2017. [Image-grounded conversations: Multimodal context for natural question and response generation](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 462–472, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Lalchand Pandia, Yan Cong, and Allyson Ettinger. 2021. [Pragmatic competence of pre-trained language models through the lens of discourse connectives](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 367–379, Online. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Amrita Saha, Mitesh Khapra, and Karthik Sankaranarayanan. 2018. Towards building large scale multimodal domain-aware conversation systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Paul Hongsuck Seo, Andreas Lehrmann, Bohyung Han, and Leonid Sigal. 2017. Visual reference resolution using attention memory for visual dialog. *Advances in neural information processing systems*, 30.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565.
- Kurt Shuster, Samuel Humeau, Antoine Bordes, and Jason Weston. 2020. [Image-chat: Engaging grounded conversations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2414–2429, Online. Association for Computational Linguistics.
- Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. 2022. [Flava: A foundational language and vision alignment model](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15638–15650.
- Russell Stewart, Mykhaylo Andriluka, and Andrew Y. Ng. 2016. [End-to-end people detection in crowded scenes](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2325–2333.
- Qingfeng Sun, Yujing Wang, Can Xu, Kai Zheng, Yaming Yang, Huang Hu, Fei Xu, Jessica Zhang, Xiubo Geng, and Daxin Jiang. 2022. [Multimodal dialogue response generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2854–2866, Dublin, Ireland. Association for Computational Linguistics.
- Anirudh Sundar and Larry Heck. 2022. [Multimodal conversational AI: A survey of datasets and approaches](#). In *Proceedings of the 4th Workshop on NLP for Conversational AI*, pages 131–147, Dublin, Ireland. Association for Computational Linguistics.
- Ronja Utescher and Sina Zarrieß. 2021. What did this castle look like before? exploring referential relations in naturally occurring multimodal texts. In *Proceedings of the Third Workshop on Beyond Vision and Language: Integrating Real-world Knowledge (LANTERN)*, pages 53–60.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Zhaoqing Wang, Yu Lu, Qiang Li, Xunqiang Tao, Yandong Guo, Mingming Gong, and Tongliang Liu. 2022. [Cris: Clip-driven referring image segmentation](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11686–11695.
- Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. 2016. Modeling context in referring expressions. In *European Conference on Computer Vision*, pages 69–85. Springer.

A MDETR Architecture

We provide Figure 6 for illustrative comparison with our proposed SitCoM-DETR in §3.2.




Figure 6: MDETR architecture.