

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

7-2022

Declaration-based prompt tuning for visual question answering

Yuhang LIU

Wei WEI

Feida ZHU

Feida ZHU

Singapore Management University, fdzhu@smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

LIU, Yuhang; WEI, Wei; ZHU, Feida; and ZHU, Feida. Declaration-based prompt tuning for visual question answering. (2022). *Proceedings of the 2022 International Joint Conference on Artificial Intelligence, Vienna, Austria, July 23-29*. 3264-3270.

Available at: https://ink.library.smu.edu.sg/sis_research/7752

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Declaration-based Prompt Tuning for Visual Question Answering

Yuhang Liu^{1,2}, Wei Wei^{1,2*}, Daowan Peng^{1,2} and Feida Zhu³

¹Cognitive Computing and Intelligent Information Processing (CCIIP) Laboratory, School of Computer Science and Technology, Huazhong University of Science and Technology, China

²Joint Laboratory of HUST and Pingan Property & Casualty Research (HPL), China

³School of Computing and Information Systems, Singapore Management University, Singapore
 {lyuhang, weiw, pengdw}@hust.edu.cn, fdzhu@smu.edu.sg

Abstract

In recent years, the *pre-training-then-fine-tuning paradigm* has yielded immense success on a wide spectrum of cross-modal tasks, such as visual question answering (VQA), in which a visual-language (VL) model is first optimized via self-supervised task objectives, *e.g.*, masked language modeling (MLM) and image-text matching (ITM), and then fine-tuned to adapt to downstream task (*e.g.*, VQA) via a brand-new objective function, *e.g.*, answer prediction. However, the inconsistency of the objective forms not only severely limits the generalization of pre-trained VL models to downstream tasks, but also requires a large amount of labeled data for fine-tuning. To alleviate the problem, we propose an innovative VL fine-tuning paradigm (named **Declaration-based Prompt Tuning**, abbreviated as **DPT**), which fine-tunes the model for downstream VQA using the pre-training objectives, boosting the effective adaptation of pre-trained models to the downstream task. Specifically, DPT reformulates the VQA task via (1) *textual adaptation*, which converts the given questions into declarative sentence form for prompt-tuning, and (2) *task adaptation*, which optimizes the objective function of VQA problem in the manner of pre-training phase. Experimental results on GQA dataset show that DPT outperforms the fine-tuned counterpart by a large margin regarding accuracy in both fully-supervised (2.68%) and zero-shot/few-shot (over 31%) settings. The data and codes are available at <https://github.com/CCIPLab/DPT>.

1 Introduction

Recently, large-scale vision-language pre-training has been an emerging topic in the multi-modal community, and delivered strong performance in numerous vision-language tasks [Yao *et al.*, 2021; Li *et al.*, 2020; Zhang *et al.*, 2021; Chen *et al.*, 2020; Lu *et al.*, 2019; Su *et al.*, 2020]. Typically, a commonly-used practice is to follow the *pre-training-then-fine-tuning paradigm* [Liu *et al.*, 2021b], in which a generic

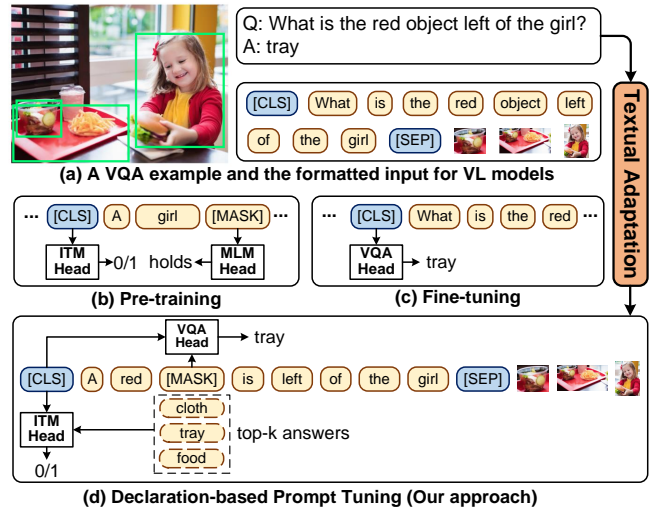


Figure 1: Illustration of (a) a VQA example and the formatted input for VL models, (b) pre-training VL models with masked language model (MLM) and image-text matching (ITM) tasks, (c) vanilla fine-tuning for VQA with a new classification head, and (d) our proposed declaration-based prompt tuning (DPT) framework that reformulates VQA task into fill-in-the-blank and image-text matching problems via textual and task adaptation. Only parts of the relevant image regions are shown for illustration.

Transformer [Vaswani *et al.*, 2017] is pre-trained on large-scale image-text datasets in a self-supervised manner, and then adapted to different downstream tasks by introducing additional parameters and fine-tuning using task-specific objectives, *e.g.*, auxiliary fully-connected layer for answer classification in visual question answering. This paradigm has greatly pushed forward the state-of-the-art of VQA task.

Despite the promising performance achieved, it's worth noting that there exists a natural gap in objective forms between pre-training and fine-tuning stages. As illustrated by Figure 1(b-c), most VL models are pre-trained via masked language modeling and image-text matching objectives, *i.e.*, recovering the masked token on the cross-modal contexts and predicting the matching scores of image-text pairs. However, in the fine-tuning stage, VQA problem is usually conducted and optimized using a brand-new task objective, *i.e.*, classifying [CLS] token into the semantic labels (*i.e.*, answers),

*Corresponding author.

where additional parameters are typically introduced. As a result, there exist great disparities in the task forms between pre-training and fine-tuning. This gap hinders the generalization of pre-trained VL models to downstream VQA task, which leads to suboptimal performance and a demand for large amount of labeled data for fine-tuning.

Inspired by the recent progress of vision-language pre-trained models (VL-PTM) [Li *et al.*, 2020; Zhang *et al.*, 2021] and prompt tuning paradigms in cross-modal domain [Yao *et al.*, 2021; Tsimpoukelli *et al.*, 2021; Radford *et al.*, 2021], in this paper we propose **D**eclaration-based **P**rompt **T**uning (DPT), a novel paradigm of fine-tuning VL-PTM for VQA problem. Our core insight is to reformulate the objective form of downstream VQA task into the format of pre-training phase, maximally mitigating the gap between two stages. To achieve this goal, we reformulate the VQA task from two aspects (refer to Figure 1(d)): (1) *textual adaptation* that converts the textual input (*i.e.*, questions) into declarative sentence form, and (2) *task adaptation* that solves VQA by recovering the masked token from the declarative sentences, and selecting the one that best matches the image. In this way, answer prediction can be achieved via cloze-filling and image-text matching, imitating the behavior of MLM and ITM tasks in the pre-training phase.

By mitigating the gap between pre-training and fine-tuning, DPT enables strong performance over various VL models and VQA datasets in both fully-supervised and zero/few-shot settings. For example, with respect to the accuracy, our method achieves 2.68% absolute improvement in the fully-supervised setting, and 31.8%~37.4% absolute improvement in the zero-shot/few-shot settings in GQA evaluation. Furthermore, the generalization experiment on VQA v2.0 equipped with recently proposed VL models shows 0.45%~1.01% absolute improvement compared to the vanilla fine-tuning approach.

In summary, the main contributions are the following,

- We introduce **D**eclaration-based **P**rompt **T**uning (DPT), a novel fine-tuning paradigm that solves VQA via adapting downstream problem to pre-training task format. To the best of our knowledge, this is the first attempt in the prompt tuning using declaration sentences for visual question answering.
- We propose novel textual and task adaptation approaches to reformulate VQA into cloze-filling and image-text matching problems, *i.e.*, MLM and ITM. The adapted tasks significantly outperform the fine-tuning counterparts in fully-supervised and few-shot settings.
- We conduct comprehensive experiments over various VL-PTMs and VQA datasets, which demonstrates the effectiveness and generalizability of DPT.

2 Related Work

2.1 Pre-trained Vision-language Models

Recently, there exists numerous work on training generic models for various downstream cross-modal tasks [Liu *et al.*, 2021a], such as visual question answering (VQA) or image caption [Cho *et al.*, 2021; Radford *et al.*, 2021; Kim *et al.*,

2021; Zhang *et al.*, 2021; Li *et al.*, 2020; Cheng *et al.*, 2020; Tan and Bansal, 2019]. Typically, a commonly-used practice is to follow a paradigm from model pre-training to model fine-tuning. In specific, in pre-training stage, a BERT-like architecture [Devlin *et al.*, 2019] is first built for pre-training in learning multi-modal representations via a variety of self-supervised tasks, for instance, a mask language model (MLM) task of recovering the masked textual tokens in the multi-modal context [Tan and Bansal, 2019; Li *et al.*, 2020], or an image-text matching (ITM) task to verify the alignment of an image to a given text [Tan and Bansal, 2019; Zhang *et al.*, 2021]. Next, in the fine-tuning stage, the pre-trained model is then fine-tuned to adapt to downstream tasks using totally different task-specific objectives, such as predicting the answer for the VQA task. In this work, instead of optimizing brand-new task objectives in the *fine-tune* stage, we attempt to reformulate VQA into the pre-training format, boosting the effective generalization of pre-trained VL models to the downstream task.

2.2 Cross-modal Prompt Tuning

Recently, prompt tuning has increasingly received attentions due to its powerful capability in keeping the optimization objectives of the pre-trained model and the downstream task consistent [Liu *et al.*, 2021b; Radford *et al.*, 2021; Yao *et al.*, 2021; Tsimpoukelli *et al.*, 2021], which enables pre-trained models generalize to downstream tasks with few/zero samples for fine-tuning. Indeed, there already exist many attempts on this topic, for example, [Radford *et al.*, 2021; Zhou *et al.*, 2021] make use of crafted templates and learnable continuous representations to reformulate the objective forms of downstream tasks. [Cho *et al.*, 2021; Jin *et al.*, 2021; Tsimpoukelli *et al.*, 2021] take account of utilizing an unified text generation framework to uniformly optimize with autoregressive objective. However, the fixed templates or pre-defined unified generation paradigm may be inadequacy in designing a suitable prompt model owing to the complex semantics in the given questions. To overcome the problem, in this paper we propose an innovative declaration-based prompt model, which exploits question-adaptive declarative sentence as prompt template so that the textual format for VQA task is more consistent with the pre-training phase, diminishing the textual gap between *pre-train* and *fine-tune* stages.

3 Methodology

In the following sections, we first present the problem statement of the VQA task (Section 3.1). Then, we describe our proposed DPT method (Section 3.2). The overall framework is depicted in Figure 2. Specifically, the image and question are converted into the input form and fed to the pre-trained VL model for multi-modal fusion, in which the declaration is typically introduced for prompt tuning. After that, the outputs of the model are exploited to perform the adapted MLM and ITM tasks for model fine-tuning and deciding the answer.

3.1 Preliminary

In this paper, we follow the problem definition in [Agrawal *et al.*, 2015], and thus the VQA problem is formulated as a

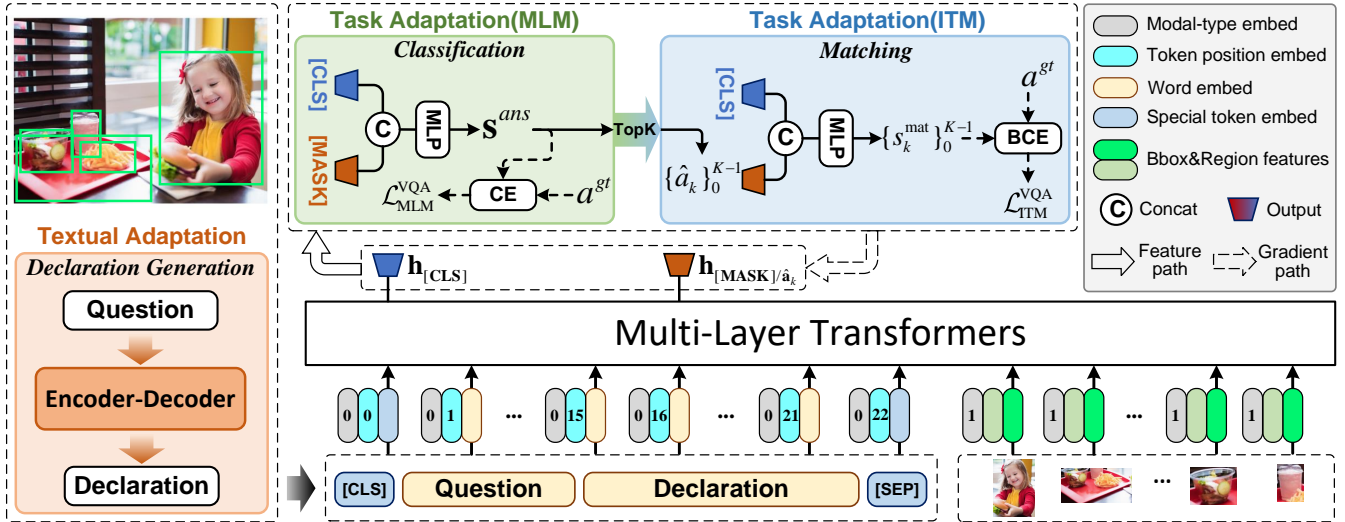


Figure 2: The framework of our proposed DPT method. Questions are converted into declarations which are concatenated to form as textual input and fed to the pre-trained VL model along with the regional features. The output [MASK] and [CLS] representations will prompt the model to predict the answer scores or (image-text) matching scores.

multi-class classification problem. Formally, the VQA task aims to select a correct answer a from a candidate answer set when given an image I and a question Q . To this end, we present the classical paradigm for VQA, namely, *pre-training-then-fine-tuning paradigm*.

Pre-training-then-fine-tuning paradigm. Given a generic architecture, *e.g.*, Transformer, the model is first pre-trained on large-scale image-text corpus via manually designed self-supervised tasks, *e.g.*, MLM and ITM. To this end, a set of region proposals extracted from the image I , $\{\mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_n\}$ and word embeddings of the question Q , $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_m\}$ are converted to the input format, *i.e.*, $\{\mathbf{e}_{[CLS]}, \mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_m, \mathbf{e}_{[SEP]}, \mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_n\}$, which is fed to the model and fused to produce the hidden representations $\{\mathbf{h}_i\}_{i=0}^{m+n+2}$, where $\mathbf{e}_{[CLS]}, \mathbf{e}_{[SEP]}$ are embeddings of special tokens. The model is further optimized using self-supervised objectives. Then, in the fine-tuning stage for VQA task, the output [CLS] is exploited to perform multi-class classification and optimized via cross-entropy loss. This paradigm introduces a brand-new task for fine-tuning, which requires a large amount of labeled data to generalize in downstream task.

3.2 Declaration-based Prompt Tuning

To facilitate the generalization of pre-trained VL models to downstream VQA tasks, we propose a declaration-based prompt tuning (DPT) paradigm that reformulates VQA into pre-training task format. As illustrated in Figure 1(b-d), there exist two challenges, *i.e.*, different forms of textual input (question *vs.* declaration) and different task objectives (MLM&ITM *vs.* answer classification). To address these issues, we present (1) **Textual Adaptation** module to convert questions into their corresponding declarative sentences, and (2) **Task Adaptation** module to reformulate answer prediction into MLM and ITM tasks. The two adapted tasks are combined to decide the final answer.

Textual Adaptation via Declaration Generation

Textual adaptation aims to convert the textual input (*i.e.*, questions) into the pre-training form (*i.e.*, declarative sentences), *e.g.*, the declaration form of “What is the red object left of the girl?” is “A red [MASK] is left of the girl.”. To this end, we introduce declaration generation which formulates this procedure as a translation problem, where the source and target texts are *question* and corresponding *declaration*, respectively. Formally, we first construct a declaration dataset using the annotations from GQA dataset [Hudson and Manning, 2019a], where the “fullAnswer” is regarded as the declaration and the short answer word/phrase in “fullAnswer” is replaced with a [MASK] token. Then, an encoder-decoder network (T5 [Raffel *et al.*, 2020]) is trained on this dataset and optimized using the standard auto-regressive cross-entropy loss. Finally, the model can be used to convert questions into declarative sentences for various VQA datasets, *e.g.*, GQA [Hudson and Manning, 2019a] and VQA [Agrawal *et al.*, 2015]. More details are provided in Section 4.1 and Appendix.

Task Adaptation

Equipped with declarative sentences, VQA can be reformulated into pre-training task format, *i.e.*, MLM and ITM. The adaptations mainly involve two aspects: textual input format and task objectives. Specifically, MLM reserves a [MASK] token in the textual input, and predicts the answer via multi-class classification. ITM replaces [MASK] with the top-k candidate answers predicted from MLM, and predicts the matching scores using binary classification.

Adaptation to MLM task. To reformulate VQA into MLM task, the question and declaration sentence are concatenated to form as the textual input:

$$\mathcal{T}^{MLM}(Q) = [CLS] Q \text{ Answer: } D [SEP] \quad (1)$$

where \mathcal{T}^{MLM} represents the conversion function that con-

verts the question Q to the input format. D denotes the declaration sentence. In Equation (1), we reserve the question in the textual input, because we find declaration sentence alone drops performance due to the lack of reasoning contexts (refer to Appendix for details). It’s worth noting that D reserves a [MASK] token, *e.g.*, *a red [MASK] is left of the girl*. In this way, the model is prompted to decide the token to fill in the mask, which exactly indicates the answer word/phrase.

On the basis of the adapted textual input, a pre-trained VL model is exploited to fuse the text and image features, producing a set of hidden representations. The outputs from [CLS] and [MASK] tokens (*i.e.*, $\mathbf{h}_{[CLS]}$ and $\mathbf{h}_{[MASK]}$) are concatenated to predict the answer:

$$\mathbf{s}^{ans} = \text{MLP}_{MLM}([\mathbf{h}_{[CLS]}; \mathbf{h}_{[MASK]}]), \quad (2)$$

$$p_1(a = a_i | Q, D, I) = \frac{\exp(\mathbf{s}_i^{ans})}{\sum_{j=0}^{|\mathcal{C}|} \exp(\mathbf{s}_j^{ans})}, \quad (3)$$

where $\mathbf{s}^{ans} \in \mathbb{R}^{|\mathcal{C}|}$ denotes the scores over the answer set \mathcal{C} . The model is optimized using cross-entropy loss, defined as:

$$\mathcal{L}_{MLM}^{VQA} = -\mathbb{E}_{\mathcal{D}}[\log p_1(a^{gt} | Q, D, I)], \quad (4)$$

where a^{gt} is the ground-truth answer. \mathcal{D} denotes the VQA dataset.

Adaptation to ITM task. To reformulate VQA into ITM task, the [MASK] token in the declaration sentence D is replaced by the top- k answers $\{\hat{a}_0, \hat{a}_1, \dots, \hat{a}_{K-1}\}$ predicted from Equation (2), resulting in K candidate declarations:

$$\{D_0^{ans}, D_1^{ans}, \dots, D_{K-1}^{ans}\}. \quad (5)$$

Based on the candidates, the textual input can be formed via concatenation of the question Q and the declaration sentence D_k^{ans} , defined as follows:

$$\mathcal{T}^{ITM}(Q) = [\text{CLS}] Q \text{ Answer: } D_k^{ans} [\text{SEP}] \quad (6)$$

where \mathcal{T}^{ITM} represents the conversion function. D_k^{ans} denotes the declaration sentence, in which the [MASK] token is replaced by the k -th candidate answer \hat{a}_k , *e.g.*, *a red tray/food/cloth is left of the girl*.

In this way, pre-trained VL models are prompted to determine whether the image-text is matched. To achieve this, the image and textual inputs are fed to the VL model, and the outputs from [CLS] and answer token (*i.e.*, $\mathbf{h}_{[CLS]}$ and $\mathbf{h}_{\hat{a}_k}$) are concatenated to predict the matching score:

$$s_k^{mat} = \text{MLP}_{ITM}([\mathbf{h}_{[CLS]}; \mathbf{h}_{\hat{a}_k}]), \quad (7)$$

$$p_2(a = \hat{a}_k | Q, D_k^{ans}, I) = \text{sigmoid}(s_k^{mat}), \quad (8)$$

where s_k^{mat} denotes the matching score of the image and the k -th candidate answer. Intuitively, the image-text pair with ground-truth answer should have higher matching score. Therefore, the model is optimized using binary cross-entropy loss, defined as follows:

$$y_k = \mathbb{I}[\hat{a}_k = a^{gt}], \quad (9)$$

$$\mathcal{L}_{ITM}^{VQA} = -\mathbb{E}_{\mathcal{D}} \frac{1}{K} \sum_{k=0}^{K-1} [y_k \log p_2(\hat{a}_k) + (1 - y_k) \log (1 - p_2(\hat{a}_k))], \quad (10)$$

where $\mathbb{I}[x] : X \rightarrow \{0, 1\}$ denotes the indicator function, which takes value 1 if x is positive and zero otherwise.

Training and inference. On the top of task adaptation, VQA has been reformulated into MLM and ITM problems. During training, we integrate the loss terms from Eq. (4) and (9) to fine-tune VL models. The total loss of DPT is defined as:

$$\mathcal{L}_{DPT} = \mathcal{L}_{MLM}^{VQA} + \mathcal{L}_{ITM}^{VQA}, \quad (11)$$

During inference, the normalized scores predicted by MLM and ITM are combined via simple summation, and the answer \hat{a} with the highest score is chosen as the final prediction result, defined as follows:

$$\hat{a} = \underset{\hat{a} \in \{\hat{a}_i\}_0^{K-1}}{\text{argmax}} (p_1(\hat{a}) + p_2(\hat{a})), \quad (12)$$

Zero-shot and few-shot learning. Equipped with DPT, previous pre-trained VL models can also be easily transformed for zero-shot or few-shot learning based VQA tasks, only if reformulating Equation (2) and (7) into the same form as the one in pre-trained phrase, and is initialized with the pre-trained weights, which can be rewritten as follows,

$$\mathbf{s}^{ans} = \text{MLP}_{MLM}^{pt}(\mathbf{h}_{[MASK]}), \quad (13)$$

$$s_k^{mat} = \text{MLP}_{ITM}^{pt}(\mathbf{h}_{[CLS]}), \quad (14)$$

where MLP_*^{pt} denotes the MLP layer initialized with pre-trained weights. Since the number of answers is less than that of vocabulary tokens, only the weights corresponding to answer words are taken to initialize MLP_{MLM}^{pt} .

4 Experiments

4.1 Implementation Details

Datasets. GQA [Hudson and Manning, 2019a] and VQA v2.0 [Agrawal *et al.*, 2015] are used to build declaration generation dataset and evaluate our proposed methods on VQA task. More details are provided in the Appendix.

Model training. T5-small [Raffel *et al.*, 2020] is chosen for declaration generation. As for VQA, VinVL [Zhang *et al.*, 2021] is selected as our base architecture. Our proposed DPT is applied to VinVL via textual and task adaptation. The model is fine-tuned using the adapted task objectives, resulting in two variants regarding the tasks for training, *i.e.*, DPT(MLM) and DPT(MLM&ITM). The number of answers used for ITM K is set to 8. For fair comparison, we follow the same training settings as reported in the previous works in the following experiments. The details of hyper-parameters are reported in Appendix.

4.2 Experimental Results

For online evaluation of GQA dataset, we compare our method with the state-of-the-art models, including non-pretrained models *i.e.*, MMN [Chen *et al.*, 2021], NSM [Hudson and Manning, 2019b], and pre-trained VL models *i.e.*, LXMERT [Tan and Bansal, 2019], VILLA [Gan *et al.*, 2020], OSCAR [Li *et al.*, 2020], VinVL [Zhang *et al.*, 2021], MDETR [Kamath *et al.*, 2021], VL-T5 [Cho *et al.*,

Method	Pre-trained	Accuracy (%)	
		Test-dev	Test-std
MMN [2021]		-	60.83
NSM [2019b]	✗	-	63.17
LXMERT [2019]		60.00	60.33
VILLA [2020]		60.98	61.12
OSCAR [2020]		61.58	61.62
VL-T5 [2021]	✓	-	60.80
MDETR [2021]		62.95	62.45
VinVL [†] _{bal} [2021]		60.76	60.89
VinVL [2021]		65.05	64.65
DPT _{bal}		63.55	63.57
DPT	✓	65.20	64.92

Table 1: Accuracy comparisons over the GQA dataset. ‘-’ and ‘†’ denote the numbers are not available and our implementation, respectively. *bal* denotes the model trained on the *balanced* split.

Prompt	Output	Task	Accuracy (%)	
			GQA	VQA v2
Baseline	[C]	Baseline	60.26	74.05
Mask	[C]&[M]	MLM	60.88	74.30
Dynamic	[C]&[M]	MLM	62.09	74.39
Declaration	[M]	MLM	60.03	73.90
	[C]&[M]	MLM	62.71	74.39
	[C]&[M]	MLM&ITM	63.13	74.50

Table 2: Effectiveness validation of declarative sentences for prompt tuning on GQA and VQA v2.0 datasets. *Output* and *Task* denote the outputs for prediction and the adapted tasks for fine-tuning, respectively. [C] and [M] are abbreviated as [CLS] and [MASK].

2021]. The results are reported in Table 1. When only exploiting *balanced* split for training, our method achieves 63.55% and 63.57% overall accuracy on test-dev and test-std, respectively, outperforming the state-of-the-art non-pretrained/pre-trained models. Specifically, our method (DPT_{bal}) surpasses the fine-tuned counterpart (VinVL_{bal}) by a significant margin of 2.68% on test-std. When using *all* split to bootstrap our model similar to [Chen *et al.*, 2021; Zhang *et al.*, 2021], our method (DPT) still ranks the top regarding overall accuracy, and outperforms the counterpart (VinVL) by 0.27% on test-std. Among the compared models, MMN and NSM also achieve competitive results even if no pre-training is performed, which is thanks to the usage of deliberately generated scene graphs or supervision of execution programs.

4.3 Ablation Study

For a deeper understanding of DPT, we further conduct the ablation studies on the local validation split of GQA and VQA v2.0 datasets (*textdev* on GQA and *val* on VQA v2.0).

Different prompts. To illustrate the effectiveness of declarative sentences for prompt tuning, several prompt variants are proposed for comparison in Table 2, defined as follows:

- **Baseline:** Vanilla fine-tuning VinVL [Zhang *et al.*,

Model	Task	Accuracy (%)	Δ (%)
VinVL[2021]	Baseline	74.05	0.00
	MLM	74.39	0.34
	MLM&ITM	74.50	0.45
ViLT[2021]	Baseline	70.71	0.00
	MLM	71.01	0.30
	MLM&ITM	71.17	0.46
UNITER[2020]	Baseline	67.72	0.00
	MLM	68.69	0.97
	MLM&ITM	68.73	1.01

Table 3: Effectiveness validation of DPT over different pre-trained VL models on VQA v2.0 datasets. Δ (%) denotes the absolute accuracy improvement margin compared with baseline.

2021] without prompt.

- **Mask:** “Answer: [MASK]”.
- **Dynamic:** “Answer: [V1] [V2] ... [V16] [MASK]”.
- **Declaration (Ours):** “Answer: *D*”.

where ‘[V1]’-‘[V16]’ denote the learnable tokens which are jointly trained during fine-tuning. As Table 2 shows, on GQA dataset, our proposed declaration-based prompt is more effective than manually designed templates (*i.e.*, Mask and Dynamic). For example, DPT with MLM task (row 5) surpasses the Mask and Dynamic with 1.83% and 0.62%, respectively. Equipped with both MLM and ITM tasks, our full model (row 6) surpasses Baseline by 2.87%. To measure the confidence of the results, we have performed additional 3 runs for our best-performing model on GQA and VQA v2.0 datasets, getting standard deviations of 0.10% and 0.06%, respectively.

Generalizability over different datasets. Table 2 shows the ablation results on VQA v2.0 with respect to different prompts. Consistent with the results on GQA, our proposed DPT surpasses the fine-tuning using fixed templates, *i.e.*, Mask or Dynamic. Specifically, our model with DPT outperforms Baseline by 0.45%. The difference in accuracy gain between GQA and VQA (2.87% vs. 0.45%) is mainly due to the question complexity and the quality of the generated declaration sentences (refer to Appendix for details).

Generalizability over different VL models. To illustrate the generalizability of our proposed method over different pre-trained VL models, we apply our DPT to the recently proposed VL models that have been pre-trained via MLM and ITM tasks, *e.g.*, UNITER [Chen *et al.*, 2020] and ViLT [Kim *et al.*, 2021]. As shown in Table 3, for all the three baselines, equipped with our DPT method, a consistent performance improvement (0.64% on average) can be observed. For example, ViLT+DPT and UNITER+DPT achieve absolute performance gains of 0.46% and 1.01% compared with the fine-tuning counterparts, respectively.

Accuracy over different question types. Figure 3 shows the accuracy breakdown on different question semantic types. It can be observed that the adapted MLM task achieves large accuracy improvement in attribute questions against Baseline




Figure 3: Accuracy breakdown over question semantic types on the GQA dataset.




Figure 4: Testdev accuracy in zero-shot and few-shot settings on GQA dataset. The mean and standard deviation are reported over 5 random splits. Δ indicates the absolute accuracy improvement compared to Baseline.

(70.46% vs. 64.87%). This shows the strength of declaration-based prompt in capturing the object attributes. Moreover, the adapted ITM task brings more performance improvement in global questions (59.24% vs. 56.69%), indicating its superior ability in the understanding of global semantics.

4.4 Zero-shot and Few-shot Results

Figure 4 shows the accuracy in zero-shot and few-shot settings on GQA dataset. We remove *yes/no* questions in the sampled splits in advance since the large proportion of *yes/no* questions (18.81% and 17.47% questions have *yes* and *no* answers, respectively) will cause large variance ($\sim 8\%$) in Baseline evaluation. As shown in Figure 4, it can be observed that DPT outperforms the vanilla fine-tuning counterparts and other prompt variants (*i.e.*, Mask and Dynamic) by a significant margin. For example, with no samples for training, our DPT achieves a strong accuracy of 36.6% while the fine-tuning counterpart can not predict correct answers due to random guessing. When provided 1~128 samples, our DPT method achieves 31.8%~37.4% absolute accuracy improvement compared to Baseline.

4.5 Case study

In Figure 5, we visualize two successful cases from our proposed DPT method. Regarding the first case, the baseline

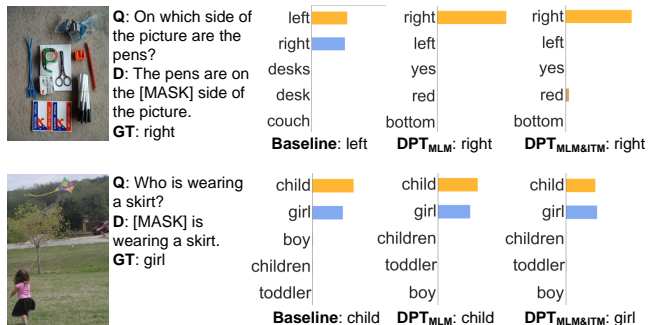


Figure 5: Visualization of the predictions from Baseline and our proposed DPT method. **Q**, **D** and **GT** denote question, generated declarative sentence and ground-truth answer, respectively.

yields almost the same probabilities for ‘left’ and ‘right’, indicates its weakness in solving such direction-related questions. In contrast, equipped with the ability of masked language model, our DPT confidently predicts the correct answer ‘right’. As for the second case, the baseline model and DPT_{MLM} both wrongly predict the answer ‘child’ mainly attribute to ‘child’ being a more frequent object that occurs in the train set. Besides, ‘child’ is a hypernym of ‘girl’ and ‘boy’, making it a universal answer to many questions. On the contrary, DPT with the adapted ITM task takes account of the semantics of answers, and gives a higher score to the answer ‘girl’, leading to the correct answer.

5 Conclusion

We propose to reformulate the VQA task into masked language model (MLM) and image-text matching (ITM) problems, maximally mitigating the gap between vision-language (VL) pre-training and fine-tuning stages. To achieve this, we first convert questions into declarative sentences with reserved [MASK] or candidate answers, mitigating the discrepancies regarding the textual input. Then, VQA problem is reformulated into pre-training format via task adaptation, which solves VQA in the manner of MLM and ITM tasks. Extensive experiments on two benchmarks validate the effectiveness and generalizability of our proposed DPT paradigm over different pre-trained VL models in both fully-supervised and zero-shot/few-shot settings.

6 Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant No.61602197, Grant No.L1924068, Grant No.61772076, in part by CCF-AFSG Research Fund under Grant No.RF20210005, in part by the fund of Joint Laboratory of HUST and Pingan Property & Casualty Research (HPL), and in part by the National Research Foundation (NRF) of Singapore under its AI Singapore Programme (AISG Award No: AISG-GC-2019-003). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of National Research Foundation, Singapore. The authors would also like to thank the anonymous reviewers for their comments on improving the quality of this paper.

References

- [Agrawal *et al.*, 2015] Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, C. Lawrence Zitnick, Devi Parikh, and Dhruv Batra. Vqa: Visual question answering. *International Journal of Computer Vision*, 123, 2015.
- [Chen *et al.*, 2020] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, 2020.
- [Chen *et al.*, 2021] Wenhui Chen, Zhe Gan, Linjie Li, Yu Cheng, William Yang Wang, and Jingjing Liu. Meta module network for compositional visual reasoning. *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2021.
- [Cheng *et al.*, 2020] Ling Cheng, Wei Wei, Xianling Mao, Yong Liu, and Chunyan Miao. Stack-vs: Stacked visual-semantic attention for image caption generation. *IEEE Access*, 8, 2020.
- [Cho *et al.*, 2021] Jaemin Cho, Jie Lei, Haochen Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *ICML*, 2021.
- [Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *ArXiv*, 2019.
- [Gan *et al.*, 2020] Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. Large-scale adversarial training for vision-and-language representation learning. *ArXiv*, 2020.
- [Hudson and Manning, 2019a] Drew A. Hudson and Christopher D. Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [Hudson and Manning, 2019b] Drew A. Hudson and Christopher D. Manning. Learning by abstraction: The neural state machine. In *NeurIPS*, 2019.
- [Jin *et al.*, 2021] Woojeong Jin, Yu Cheng, Yelong Shen, Weizhu Chen, and Xiang Ren. A good prompt is worth millions of parameters? low-resource prompt-based learning for vision-language models. *ArXiv*, 2021.
- [Kamath *et al.*, 2021] Aishwarya Kamath, Mannat Singh, Yann LeCun, Ishan Misra, Gabriel Synnaeve, and Nicolas Carion. Mdetr - modulated detection for end-to-end multi-modal understanding. *ArXiv*, 2021.
- [Kim *et al.*, 2021] Wonjae Kim, Bokyoung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, 2021.
- [Krishna *et al.*, 2016] R. Krishna, Yuke Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, Stephanie Chen, Yannic Kalantidis, L. Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123, 2016.
- [Li *et al.*, 2020] Xiujun Li, Xi Yin, Chunyuan Li, Xiaowei Hu, Pengchuan Zhang, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. Oscar: Object-semantic aligned pre-training for vision-language tasks. In *ECCV*, 2020.
- [Liu *et al.*, 2021a] Daizong Liu, Shuangjie Xu, Xiao-Yang Liu, Zichuan Xu, Wei Wei, and Pan Zhou. Spatiotemporal graph neural network based mask reconstruction for video object segmentation. In *AAAI*, 2021.
- [Liu *et al.*, 2021b] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ArXiv*, 2021.
- [Lu *et al.*, 2019] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vlbart: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019.
- [Radford *et al.*, 2021] 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. Learning transferable visual models from natural language supervision. In *ICML*, 2021.
- [Raffel *et al.*, 2020] Colin Raffel, Noam M. Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *ArXiv*, 2020.
- [Su *et al.*, 2020] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vl-bert: Pre-training of generic visual-linguistic representations. *ArXiv*, 2020.
- [Tan and Bansal, 2019] Hao Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *EMNLP*, 2019.
- [Tsimpoukelli *et al.*, 2021] Maria Tsimpoukelli, Jacob Menick, Serkan Cabi, S. M. Ali Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. *ArXiv*, 2021.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017.
- [Yao *et al.*, 2021] Yuan Yao, Ao Zhang, Zhengyan Zhang, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. Cpt: Colorful prompt tuning for pre-trained vision-language models. *ArXiv*, 2021.
- [Zhang *et al.*, 2021] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5575–5584, 2021.

[Zhou *et al.*, 2021] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *ArXiv*, 2021.

Appendix

A Datasets

VQA v2.0. VQA v2.0 [Agrawal *et al.*, 2015] is the most commonly used VQA benchmark. It contains real images and annotated question-answer pairs. Each image has an average of 5 questions. Each question has 10 answers annotated by different annotators, and the most frequent answer is regarded as the ground-truth answer. The dataset is split into *train*, *val*, and *test* sets, statistically detailed in Table 4. The evaluation metric (*i.e.*, accuracy) on this dataset is robust to inter-human variability, calculated as follows:

$$\text{Acc}(\mathbf{ans}) = \min \left\{ \frac{\#\text{humans that said } \mathbf{ans}}{3}, 1 \right\} \quad (15)$$

GQA. GQA [Hudson and Manning, 2019a] is a VQA dataset that characterizes in compositional question answering and visual reasoning about real-world images. With the help of the scene graph annotations from Visual Genome [Krishna *et al.*, 2016], GQA is able to maximally mitigate the language priors that exist widely in previous VQA datasets. Additionally, the questions are generated via the engine that operates over 524 patterns, spanning 117 question groups. Therefore, it requires more complicated reasoning skills to answer the questions. GQA consists of two splits, *i.e.*, *all* split that contains 22M QA pairs, and *balanced* split that consists of 1.7M QA pairs with resampled question-answer distribution. The dataset is split into 70% train, 10% validation, 10% test and 10% challenge. The statistics of *balanced* split are detailed in Table 5.

Declaration dataset. As introduced in Section 4.1, the declaration dataset is generated from the annotations from GQA dataset. Specifically, each annotation in GQA training set contains three keys, *i.e.*, *question*, *answer*, and *fullAnswer*, illustrated in Table 6. It can be observed that ‘answer’ is usually a word or phrase while the ‘fullAnswer’ is a complete declarative sentence. According to statistics, we find that most short ‘answer’s are included in the ‘fullAnswer’s, which inspires us that instead of choosing one from a candidate answer set, we can extract a word/phrase as the answer from the complete declarative sentence in the form of cloze-filling. Formally, we replace the short ‘answer’ included in the ‘fullAnswer’ with a [MASK] token, resulting in a declarative sentence that can be used for answer clozing, illustrated in Table 7.

In this way, we are able to convert the questions into declaration format from the *all* split of GQA dataset, resulting in 726k and 181k question-declaration pairs for training and validation, respectively. Several examples from our proposed declaration dataset are shown in Table 8.

Split	#Images	#Questions	#Answers
Train	82,783	443,757	4,437,570
Val	40,504	214,354	2,143,540
Test	81,434	447,793	-
All	204,721	1,105,904	-

Table 4: Statistics of samples in VQA-v2 dataset

Split	#Images	#Questions	#Vocab
Train	72,140	943,000	3,097
Val	10,234	132,062	
Test-dev	398	12,578	
Test	2,987	95,336	
All	85,759	1,182,976	3,097

Table 5: Statistics of balance-split in GQA dataset

questionId	201640614
question	Who is wearing the dress?
answer	woman
fullAnswer	The woman is wearing a dress.

Table 6: An annotation example of GQA dataset. Only part of the annotations is listed for illustration.

B Model Training and Evaluation

Declaration generation training. We use T5-small [Rafel *et al.*, 2020] as our declaration generation model due to its transferable abilities in unified text-to-text generation. Specifically, the generation process is regarded as a *translation* problem, in which *question* and *declaration* represent the source and target texts with the same vocabulary, respectively. The model is fine-tuned on the training set and evaluated on the validation set. The parameters are trained via Adam optimizer with learning rate of $5e^{-5}$ and batch size of 4 for 480k steps in total. The fine-tuned model can be used to convert questions for GQA and VQA v2.0 datasets.

GQA training and evaluation. Since GQA consists of two splits, *i.e.*, *balanced* and *all*, we evaluate our proposed methods on both splits. With regard to the *balanced* split, the model is trained on the concatenation of *train+val* and evaluated on the *testdev*. As for *all* split, we follow VinVL [Zhang *et al.*, 2021] to set up the fine-tuning procedure, which fine-tunes the model on *all* split for 5 epochs and then fine-tunes on the *balanced* split for 2 epochs. We evaluate our methods on both splits in the online evaluation. As for local ablation study, only *balanced* split is used.

VQA v2.0 training and evaluation. VQA v2.0 is typically designed to assess the generalizability of our proposed methods. Therefore, the model is trained on the *train* split and evaluated on local *valid* split. As for model training, we follow the same settings as the previous works (*i.e.*, VinVL [Zhang *et al.*, 2021], ViLT [Kim *et al.*, 2021] and UNITER [Chen *et al.*, 2020]), on which our DPT is implemented based.

C Additional Experimental Results

C.1 Declaration Generation

The commonly used metric in machine translation, *i.e.*, Bilingual Evaluation Understudy (BLEU) is adopted to evaluate the translation quality of the generated declarative sentences. Specifically, BLEU calculates scores based on the overlap of words in the reference text and the generated text. The higher score of BLEU denotes the better generation quality. We eval-

questionId	201640614
question	Who is wearing the dress?
declaration	The [MASK] is wearing a dress.

Table 7: An example of question and the corresponding declaration.

Type	Content
Q	Are there any black gloves or hats?
D	[MASK]
Q	What is the vehicle that is to the right of the man?
D	the vehicle is a [MASK].
Q	On which side of the image is the woman?
D	the woman is on the [MASK] of the image.
Q	Who is swinging the bat?
D	the [MASK] is swinging the bat.
Q	What color are the gloves?
D	the gloves are [MASK].

Table 8: Five samples from the declaration dataset. Q and D indicate question and declaration, respectively.

uate the best T5 model on the 181k validation set, getting BLEU score of 0.97, indicating that the fine-tuned T5 model is able to generate declarative sentences of very high quality.

C.2 Number of Candidate Answers

Table 9 shows the ablation results over the number of candidate answers for image-text matching. It can be observed that the best performance 63.13% is achieved when $K = 8$. Additionally, computational complexity will increase dramatically when K grows larger, while the accuracy does not change much. Therefore, we set $K = 8$ throughout the experiments.

C.3 Accuracy Breakdown Analysis

Figure 6 shows the accuracy breakdown over the question length on *testdev* of GQA dataset. Generally, the question length reflects the reasoning complexity to a certain extent. As Figure 6 shows, it can be observed that our DPT obtains more accuracy improvement on the question with length of 14~20, while achieves similar accuracy on the short questions (with length of 3~13), demonstrating the effectiveness of our proposed DPT over long questions.

D Case Study

D.1 Declaration Generation

To perform predictions on the test split of GQA or VQA v2.0 datasets, the fine-tuned T5 model is exploited to convert questions into declarative sentences. In Table 10, we visualize several declarative sentences from GQA dataset, which are generated by the fine-tuned T5 model. It can be observed that [MASK] token is placed at the proper position that can prompt the model to predict the answer. However, we find that some generated sentences may not contain complete semantic information of the question. For example, the sample-3 refers to *vehicle in front of the flag*, while only *vehicle* is

K	Accuracy (%)
0	62.79
1	62.79
2	62.90
4	63.06
8	63.13
16	63.12

Table 9: Ablation experiments on the GQA balanced dataset. K stands for the number of candidate answers chosen for image-text matching.

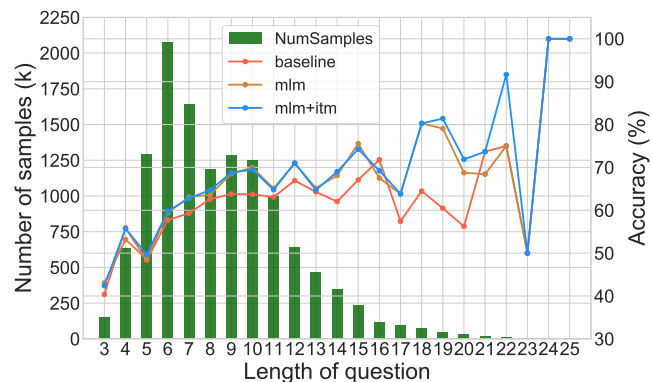


Figure 6: Accuracy breakdown over question length on GQA dataset.

generated in the declarative sentence. This may cause ambiguity in visual reasoning. Therefore, we preserve the original question in the textual input of VL models, maximally reserving the semantics of questions.

Table 11 shows several declarative sentences generated from VQA v2.0 dataset. Since the questions in VQA v2.0 are raised manually, there exists various question types, which requires broader capabilities beyond the image understanding. As Table 11 shows, T5 model is able to generalize to VQA v2.0 dataset, and generate appropriate declarative sentences for most cases. For example, the question pattern in sample-1 (*i.e.*, *what ... say?*) has not appeared in GQA dataset, but the T5 model still produces the proper declarative sentence. However, there also exist several question types that are difficult to convert, *e.g.*, *why*, *how*, etc. For example, the sample-2 asks the reason about *the sailboats have their sails lowered*, but the [MASK] token in the declarative sentence is unable to prompt the answer. This limitation can explain the differences of the absolute accuracy improvement on GQA and VQA v2.0 datasets shown in Table 2.

D.2 Visual Question Answering

In Figure 7, we visualize several prediction results from Baseline, DPT(MLM) and DPT(MLM&ITM) models. The samples are taken from the cases where Baseline predicts wrong answers while our proposed DPT gets the correct ones. From the first two rows, it can be observed that Baseline fails on the direction-related questions, especially *right/left* questions. The predicted scores of *right* and *left* from Baseline are al-

Id	Type	Content
0	Q D	Which color is the shirt? the shirt is [MASK].
1	Q D	What is beneath the microwave? the [MASK] is beneath the microwave.
2	Q D	Is this a bed or a cabinet? this is a [MASK].
3	Q D	Which kind of vehicle is in front of the flag? the vehicle is a [MASK].
4	Q D	Are there rivers or oceans that are not calm? [MASK].

Table 10: Examples from the declaration generation on GQA test split. Q and D denote the question and generated declarative sentence by T5 model, respectively.

Id	Type	Content
0	Q D	What is the man doing? the man is [MASK].
1	Q D	What does the sign say? the sign says the [MASK].
2	Q D	Why do the sailboats have their sails lowered? the sailboats have [MASK] lowered.
3	Q D	Could this roof be tiled? [MASK]
4	Q D	What pattern is the person’s shirt? the shirt is [MASK].

Table 11: Examples from the declaration generation on VQA v2.0 dataset. Q and D denote the question and generated declarative sentences by T5 model, respectively.

most equal, indicating that the model has difficulty in judging the directions. In contrary, our proposed DPT reserves a [MASK] token for answer cloze-filling. Benefit from the capability learned from the pre-training phase (*i.e.*, MLM), DPT is able to predict the correct answer with high confidence. The last two rows shows examples in which the [MASK] token indicates an object in the image. While the Baseline predicts wrong answers (*i.e.*, *door* and *man*), our proposed DPT can refer to the target objects in images, resulting in correct answers (*i.e.*, *chair* and *woman*).

Figure 8 visualizes the samples where Baseline and DPT(MLM) predict wrong answers while DPT(MLM&ITM) gets the correct ones. As shown in the examples, the Baseline and DPT(MLM) tend to predict general answers that appear more often in the training set, *e.g.*, *boy*, *man*, *truck* *etc.* Such bias is caused by the classification layer in the model. On the contrary, our proposed DPT with ITM task chooses the answer via the matching scores, instead of the answer probabilities. Therefore, our proposed method is able to mitigate the answer prediction bias to a extent, and predicts more granular answers due to the consideration of answer semantics, *e.g.*,

controller vs. *wii controller*, *truck* vs. *fire truck*.

D.3 Attention Visualization

In Figure 9, we visualize the attention maps corresponding to [CLS] and [MASK] tokens from the last Transformer layer. From these examples, it can be observed that in some cases, the [CLS] token learned from Baseline model fails to focus on the question-relevant regions, resulting in wrong answers. On the contrary, the learned [CLS] and [MASK] tokens show reasonable attentions on the relevant objects, thus producing the right answers.

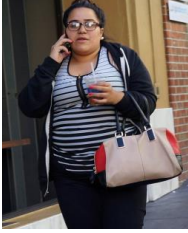



Sample	Baseline	DPT(MLM)	DPT(MLM&ITM)
 <p>Q: On which side of the picture is the leather bag? D: The leather bag is on the [MASK] of the image. GT: right</p>	<p>left ██████</p> <p>right ██████</p> <p>desks</p> <p>couch</p> <p>parking lot</p> <p>Baseline: left</p>	<p>right ██████</p> <p>left ██████</p> <p>no</p> <p>yes</p> <p>pink</p> <p>DPT_{MLM}: right</p>	<p>right ██████</p> <p>left ██████</p> <p>no</p> <p>yes</p> <p>pink</p> <p>DPT_{MLM&ITM}: right</p>
 <p>Q: On which side of the picture is the chair? D: The chair is on the [MASK] of the image. GT: right</p>	<p>left ██████</p> <p>right ██████</p> <p>desks</p> <p>couch</p> <p>parking lot</p> <p>Baseline: left</p>	<p>right ██████</p> <p>left ██████</p> <p>front</p> <p>yes</p> <p>bottom</p> <p>DPT_{MLM}: right</p>	<p>right ██████</p> <p>left ██████</p> <p>front</p> <p>yes</p> <p>bottom</p> <p>DPT_{MLM&ITM}: right</p>
 <p>Q: What is the picture hanging above? D: The picture is hanging above [MASK]. GT: chair</p>	<p>door ██████</p> <p>wall ██████</p> <p>hair ██████</p> <p>chair ██████</p> <p>table ██████</p> <p>Baseline: door</p>	<p>chair ██████</p> <p>seat</p> <p>chairs</p> <p>chair</p> <p>wall</p> <p>DPT_{MLM}: chair</p>	<p>chair ██████</p> <p>seat</p> <p>chairs</p> <p>chair</p> <p>wall</p> <p>DPT_{MLM&ITM}: chair</p>
 <p>Q: Who is staring at the giraffe? D: The [MASK] is staring at the giraffe. GT: woman</p>	<p>man ██████</p> <p>woman ██████</p> <p>people</p> <p>guy</p> <p>lady</p> <p>Baseline: man</p>	<p>woman ██████</p> <p>man ██████</p> <p>lady</p> <p>people</p> <p>couple</p> <p>DPT_{MLM}: woman</p>	<p>woman ██████</p> <p>man ██████</p> <p>lady</p> <p>people</p> <p>couple</p> <p>DPT_{MLM&ITM}: woman</p>

Figure 7: Visualization of *testdev* results on the GQA dataset. Top-5 answers predicted by Baseline, DPT(MLM) and DPT(MLM&ITM) models are displayed in the right three columns.


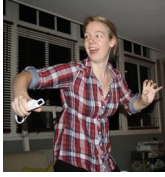


Sample	Baseline	DPT(MLM)	DPT(MLM&ITM)
 <p>Q: Who is wearing the jacket? D: The [MASK] is wearing a jacket. GT: skier</p>	boy ██████ man ██████ child ██████ skier ██████ guy ██████ Baseline: boy(0.70%)	man ██████ boy ██████ skier ██████ woman ██████ guy ██████ DPT_{MLM}: man(1.23%)	man ██████ boy ██████ skier ██████ woman ██████ guy ██████ DPT_{MLM&ITM}: skier(0.09%)
 <p>Q: What device is the woman holding? D: The woman is holding the [MASK] GT: wii controller</p>	controller ██████ wii controller ██████ remote control ██████ game controller ██████ wii ██████ Baseline: controller(0.07%)	controller ██████ wii controller ██████ remote control ██████ game controller ██████ wii ██████ DPT_{MLM}: controller(0.07%)	controller ██████ wii controller ██████ remote control ██████ game controller ██████ wii ██████ DPT_{MLM&ITM}: wii controller (0.05%)
 <p>Q: What is the light blue sky higher than? D: The sky is higher than the [MASK]. GT: mountains</p>	low ██████ tall ██████ high ██████ short ██████ hill ██████ Baseline: low(0.005%)	sky ██████ mountains ██████ hills ██████ mountain ██████ hill ██████ DPT_{MLM}: sky(0.05%)	sky ██████ mountains ██████ hills ██████ mountain ██████ hill ██████ DPT_{MLM&ITM}: mountains (0.02%)
 <p>Q: What type of vehicle is on the street? D: The vehicle is a [MASK]. GT: fire truck</p>	truck ██████ bus ██████ fire truck ██████ van ██████ car ██████ Baseline: truck(0.31%)	truck ██████ fire truck ██████ car ██████ van ██████ bus ██████ DPT_{MLM}: truck(0.31%)	truck ██████ fire truck ██████ car ██████ van ██████ bus ██████ DPT_{MLM&ITM}: fire truck (0.03%)

Figure 8: Visualization of *testdev* results on the GQA dataset. Top-5 answers predicted by Baseline, DPT(MLM) and DPT(MLM&ITM) models are displayed in the right three columns. The percentage in parentheses denotes the proportion of answers in the training set.







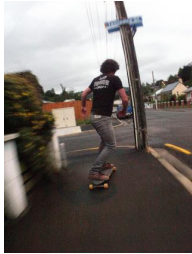
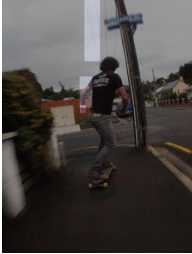
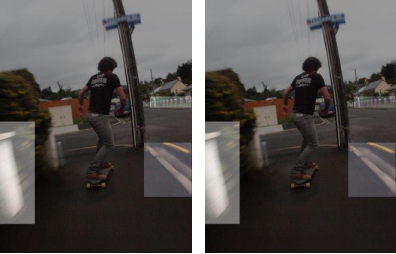


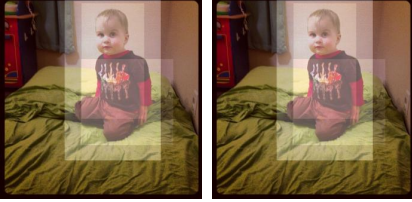
Sample	Baseline	DPT(MLM)
 <p>Q: Is the soccer player that is to the left of the ball female or male? D: The soccer player is [MASK]. GT: male</p>	 <p>Baseline: female</p>	 <p>DPT_{MLM}: male</p>
 <p>Q: Is the man to the left of a ball? D: [MASK] GT: yes</p>	 <p>Baseline: no</p>	 <p>DPT_{MLM}: yes</p>
 <p>Q: How clean is the concrete sidewalk? D: The sidewalk is [MASK]. GT: dirty</p>	 <p>Baseline: clean</p>	 <p>DPT_{MLM}: dirty</p>
 <p>Q: Is the boy to the right or to the left of the curtain? D: The boy is to the [MASK] of the curtain. GT: right</p>	 <p>Baseline: left</p>	 <p>DPT_{MLM}: right</p>

Figure 9: Visualization of attention maps corresponding to [CLS] and [MASK] tokens. All attention maps are obtained from the last Transformer layer. Top-3 objects with the maximum attention scores are visualized in the image. As for Baseline, the attention map from [CLS] is exploited for visualization. As for DPT(MLM), the two images show the attention maps of [CLS] (left) and [MASK] (right) tokens, respectively.