

UniFine: A Unified and Fine-grained Approach for Zero-shot Vision-Language Understanding

Rui Sun^{1*}, Zhecan Wang^{1*}, Haoxuan You^{1*}, Noel Codella²,
Kai-Wei Chang³, Shih-Fu Chang¹

¹ Columbia University ² Microsoft Research ³ University of California, Los Angeles
{rs4110, zw2627, hy2612, sc250}@columbia.edu
ncodella@microsoft.com, kwchang@cs.ucla.edu

Abstract

Vision-language tasks, such as VQA, SNLI-VE, and VCR are challenging because they require the model’s reasoning ability to understand the semantics of the visual world and natural language. Supervised methods working for vision-language tasks have been well-studied. However, solving these tasks in a zero-shot setting is less explored. Since Contrastive Language-Image Pre-training (CLIP) has shown remarkable zero-shot performance on image-text matching, previous works utilized its strong zero-shot ability by converting vision-language tasks into an image-text matching problem, and they mainly consider global-level matching (e.g., the whole image or sentence). However, we find visual and textual fine-grained information, e.g., keywords in the sentence and objects in the image, can be fairly informative for semantics understanding. Inspired by this, we propose a unified framework to take advantage of the fine-grained information for zero-shot vision-language learning, covering multiple tasks such as VQA, SNLI-VE, and VCR. Our experiments show that our framework outperforms former zero-shot methods on VQA and achieves substantial improvement on SNLI-VE and VCR. Furthermore, our ablation studies confirm the effectiveness and generalizability of our proposed method. Code is available at <https://github.com/ThreeSR/UniFine>.

1 Introduction

VQA (Antol et al., 2015), SNLI-VE (Xie et al., 2019), and VCR (Zellers et al., 2019) are vision-language tasks, which utilize the text and corresponding image to test a system’s cross-modal reasoning ability. These tasks are challenging as requiring models to obtain a joint understanding of visual and textual modality. Nevertheless, they are also meaningful since this capability plays an essential role in daily human-robot interaction, e.g.,

* Equal Contribution

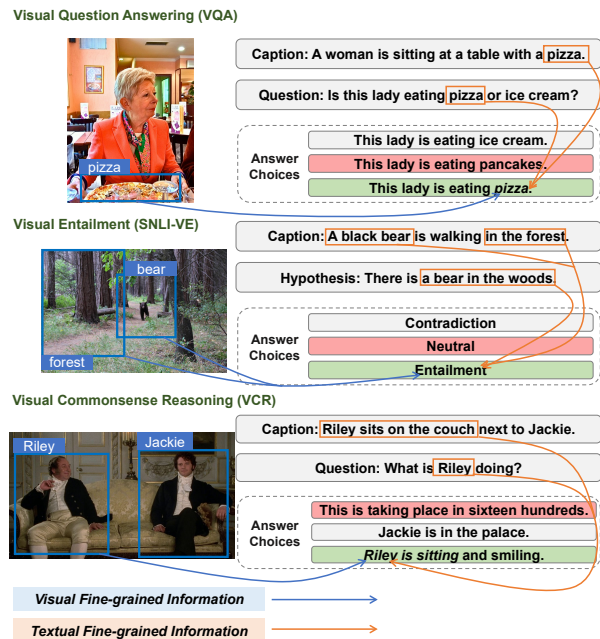


Figure 1: Examples of how fine-grained information is utilized to help CLIP from VQA, SNLI-VE, and VCR. Before extracting the fine-grained information, CLIP gives the wrong answer shown as the red box. With the assistance of visual and textual fine-grained information, CLIP can make the correct decision as the green box shows. (For visualization, only three answer choices are kept in VQA and VCR. And unisex names (Riley and Jackie) are added in VCR.)

asking a robot how many people are in the image. Despite the difficulty, a line of work (Tan and Bansal, 2019; Li et al., 2019; Lu et al., 2019; Chen et al., 2019; Su et al., 2019; Li et al., 2020) has been dedicated to resolving these vision-language tasks in a supervised setting and obtaining impressive progress. However, these methods all suffer from a significant problem of being costly as they require expert knowledge to collect well-annotated image-text data. On the other hand, zero-shot methods for vision-language tasks can successfully bypass this problem without costly annotations. Unfortunately, limited methods and relatively fewer works have been dedicated to exploring this direction.

Recently, CLIP (Radford et al., 2021) has been proposed to acquire visual concepts using natural language supervision. It jointly trains an image encoder and a text encoder on 400M noisy image-text pairs collected from the Internet by aligning images and texts through a contrastive loss.

Previous works (Song et al., 2022; Subramanian et al., 2022; Shen et al., 2021; Wang et al., 2022b) demonstrated that CLIP can achieve strong zero-shot performance of vision-language tasks by converting original tasks into the image-text matching format. However, they mainly consider matching on an instance or global level, *i.e.*, the whole image or sentence, ignoring the significance of fine-grained elements, *e.g.*, keywords in the sentence and objects in the image. Meanwhile, we find these fine-grained elements are important for specific downstream tasks, especially in zero-shot learning.

For instance, in Fig. 1, CLIP makes three incorrect predictions in three zero-shot vision-language tasks. For VQA, the model infers the wrong object "pancake" for the verb "eating", as it does not capture the details in the image (pizza on the table) and captions (pizza is mentioned). We posit that if we can find a proper solution to navigate the model to focus on these detailed pieces of textual and visual information, the model would likely have a better chance of selecting the correct answer label. This conjecture also seems true and generalizable across multiple zero-shot downstream tasks as shown by the three examples from different vision-language tasks, *i.e.*, VCR, VQA, and SNLI-VE in Fig. 1. Yet, we also recognize potential challenges also exist as those different tasks may differ from many perspectives including the distribution of image categories or scenes, the different semantic focus, and format of text premises between declarative statements and questions, and different task formats in terms of image-text matching or classification.

To overcome these challenges, we first identify two common fundamental steps required to utilize the fine-grained information across different vision-language tasks: 1) Extraction of the fine-grained information from context information, *e.g.*, the extraction of the word "pizza" from the caption in VQA as in Fig. 1. 2) The semantic matching between these extracted fine-grained information and answer choices or hypothesis. Based on these, we propose a unified approach leveraging these two common steps thus it can assist the model to generalize over different vision-language tasks. For

the extractor, we have two branches – 1) the vision branch and 2) the textual branch. In the vision branch, we employ Faster-RCNN (Ren et al., 2015) to extract object-level information. We select relevant object regions guided by the question in VQA and VCR or hypothesis in SNLI-VE. After that, we concatenate the whole image and its selected image regions and input them into the image encoder of CLIP. For textual information extraction, we exploit rich information from the image caption generated by a recently-developed captioning model OFA (Wang et al., 2022a) and question in VQA and VCR or hypothesis in SNLI-VE to boost the zero-shot performance.

It's noted that although we employ the image caption and question on a sentence level rather than a word level, we compute the cosine similarity between them and answer texts, which means if there are keywords in the answer texts which can be matched in the caption or question, then we will obtain high scores in zero-shot prediction.

Therefore, it is still a process of fine-grained information extraction. By using fine-grained information, our model outperforms previous methods on zero-shot VQA and we are the first to benchmark zero-shot VCR and SNLI-VE. The experiments confirm the effectiveness of our proposed method.

Our contributions can be summarized as follows:

- To the best of our knowledge, we are the first to propose a unified approach based on fine-grained information extraction for zero-shot learning of different vision-language tasks.
- Our approach outperforms previous CLIP-based methods for zero-shot learning of VQA and we are the first to study CLIP's zero-shot ability for SNLI-VE and VCR.
- The experiments and ablation studies confirm the generalizability of our proposed method and the significance of visual and textual fine-grained information for zero-shot learning of vision-language tasks.

2 Related Work

Vision-language understanding tasks. Unlike unimodal tasks, vision-language understanding tasks need joint understanding between vision and language, which require a deeper reasoning ability of the system. In VQA (Goyal et al., 2017), given a question, the model needs to understand the details of the corresponding image based on

the question to answer correctly. The real images in VQA come from MS COCO (Lin et al., 2014) and each of them is paired with a caption in COCO Captions (Chen et al., 2015). For another task VCR (Zellers et al., 2019), its semantic focus is different from VQA since it concentrates more on common-sense questions. The model needs to answer the recognitive questions (like VQA) at first, then it is also required to correctly answer the cognitive questions, which are rationales of the choice in the first question. The images in VCR are collected from movie clips. SNLI-VE originated from Stanford Natural Language Inference (SNLI) (Bowman et al., 2015), which is a text entailment (TE) task based on the Flickr30k (Young et al., 2014) image captions. It extends TE into the visual domain and it has a different task format from the VQA and VCR because the previous question-answering format is replaced with the hypothesis. Given the image and hypothesis, the model needs to predict whether the image semantically entails the text. The images in SNLI-VE are from Flickr30k with annotated captions.

Vision-language pre-trained models. Early vision-language pre-trained models (Tan and Bansal, 2019; Lu et al., 2019; Li et al., 2019; Chen et al., 2019; Su et al., 2019; Li et al., 2020) utilize cross-modal transformer (Vaswani et al., 2017) pre-trained on well-annotated image-text pairs. Different from these models, contrastive learning frameworks (Radford et al., 2021; Pham et al., 2021; Jia et al., 2021) are trained on noisy image-text pairs crawled from the Internet through contrastive loss, which employs the dot product between visual and textual modality. Due to the large-scale training data, these models acquire rich prior knowledge and show strong zero-shot ability on vision benchmarks like ImageNet (Deng et al., 2009).

Vision-language zero-shot learning. There is a line of work utilizing CLIP to do zero-shot learning for vision-language tasks. ReCLIP (Subramanian et al., 2022) utilizes CLIP to present a zero-shot method for referring expression comprehension (ReC), which outperforms prior zero-shot ReC approaches. CLIP-ViL (Shen et al., 2021) exploits CLIP to do zero-shot VQA by simply concatenating question and answer pair for each question and constructing "question: [question text] answer: [answer text]" as the prompt. Then, they feed the text and image into the text encoder and the image encoder of CLIP, which produces the near-chance

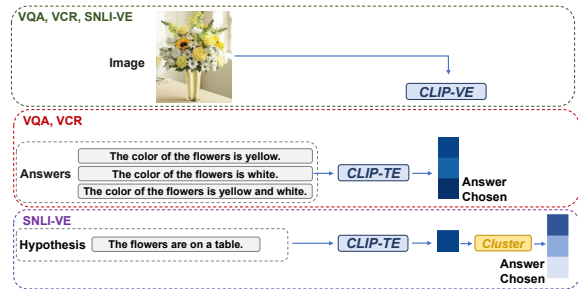


Figure 2: Baseline method of UniFine. (Note: For visualization, only three answer choices are kept in VQA and VCR. CLIP-VE denotes CLIP Visual Encoder, and CLIP-TE denotes CLIP Text Encoder.)

level performance. The most relevant work to ours is TAPC (Song et al., 2022), which manually designs the prompt and leverages T5 (Raffel et al., 2020), a large pre-trained text-to-text Transformer, to convert the question-answering problem into the image-text matching task. Then, it employs CLIP’s remarkable zero-shot image-text matching ability on VQA, whose results surpass CLIP-ViL by a large margin. However, these works handle different tasks on an instance level rather than fully utilizing the visual and textual fine-grained information (*i.e.*, keywords in the sentence and objects in the image) like ours. Moreover, we can tackle a diverse set of tasks but they just concentrate on one specific task.

3 Method

In this section, we introduce our method for visual and textual fine-grained information extraction to improve zero-shot learning of vision-language tasks including VQA, VCR, and SNLI-VE.

3.1 Baseline Method

In the baseline method shown in Fig. 2, we use CLIP to do zero-shot learning of vision-language tasks. CLIP consists of a visual encoder \mathbb{V} (*e.g.*, ResNet (He et al., 2016) and ViT (Dosovitskiy et al., 2020)) and a text encoder \mathbb{T} (*e.g.*, transformer (Vaswani et al., 2017)), where the image and text are processed independently.

Followed by the encoder, there is the dot product (*i.e.*, alignment score) between visual and textual features, *i.e.*, $\mathbb{V}(\text{image}) \cdot \mathbb{T}(\text{text})$. We input the image from VQA, VCR, and SNLI-VE into the CLIP visual encoder. Since there is a difference in task format, answer choices from VQA and VCR and the hypothesis from SNLI-VE are input into the CLIP text encoder. After encoding, we can obtain

the alignment score between the image and text. In VQA and VCR, we select the answer with the highest score. In SNLI-VE, there is a clustering process after the dot product, which is demonstrated in Algo. 1, and we select the answer with the lowest score.

3.2 Visual fine-grained information extraction

In visual fine-grained information extraction, we aim to find the related image regions to the question in VQA and VCR or the hypothesis in SNLI-VE since these regions can provide local visual clues to complement the global image. The objects and attributes are detected by Faster-RCNN (Ren et al., 2015), which is pre-trained on Visual Genome (Krishna et al., 2017) provided by Anderson et al. (2018). We select the top N relevant image regions (N is a hyperparameter, which will be analyzed in Sec. 4.3) by image region score (*i.e.*, cosine similarity) between the textual features of the question or hypothesis and the object class&attribute (*e.g.*, yellow flowers) encoded by RoBERTa (Liu et al., 2019):

$$\text{Top-}N\{\cos(\mathbb{R}(\text{Query}), \mathbb{R}(\{\text{Attr}(o_i), \text{Class}(o_i)\}))\} \quad (1)$$

where \mathbb{R} is RoBERTa, $\cos(\cdot)$ is cosine similarity, O is the set of objects detected by Faster-RCNN, $\text{Attr}()$ and $\text{Class}()$ are attribute and class of object respectively, and Query is the question in VQA and VCR or the hypothesis in SNLI-VE. After selection, the global image and selected image regions will be fed into CLIP visual encoder to obtain the encoded feature of each.

3.3 Textual fine-grained information extraction

Next, we present how textual fine-grained information is extracted and incorporated into our framework. To be more specific, two types of information are studied: image caption and question. Questions as a prior can narrow down the range of answer candidates and get rid of irrelevant answers. Image caption can transform the information inside the image into text so that it can be compared with answers in the same domain. Image captions are generated from the image, but their format is language. Thus, we arguably regard image captions as textual fine-grained information. Overcoming the challenge in different formats of vision-language tasks, we introduce a relatively unified way to extract and utilize textual fine-grained information in the zero-shot scenario.

Visual Question Answering: VQA

Following previous work, we experiment on the validation set of VQAv2 (Goyal et al., 2017). Typically, VQA is regarded as a classification problem and there are 3,129 most frequent answers used for classification. There are 65 types of questions (*e.g.*, *does this* type) and 3 types of answers including *Yes/No*, *Number*, and *Other* in VQAv2.

Although in VQA, each image is paired with a ground truth caption from MS COCO, we still choose to use OFA, a SOTA model of image captioning, to generate the caption given the image, because not every dataset is annotated with ground truth captions and we would like to make our method generalizable.

As Shen et al. (2021) shows, directly inputting the concatenation of the question and answer into CLIP will lead to near-chance level performance. In addition, there are more than 3,000 answer candidates in VQAv2, which will largely slow down the inference speed of zero-shot VQA with all answers input into CLIP. To bypass that, we utilize an answer-filtering method to downsize the number of answer choices inspired by Song et al. (2022).

Following Song et al. (2022), we first convert the question-answering format into declarative templates with the $\langle \text{extra_id_0} \rangle$ token by T5 low-shot demonstration. Then, templates with $\langle \text{extra_id_0} \rangle$ token are input into the T5 and we obtain the plausibility of each answer candidate based on T5 output probability. Next, we select the top K answers. More details can be found at Sec. A.2.

In this way, we can downsize the number of answers in VQA. There are three different answer types in VQA, which are processed differently in the answer filtering process. For *Yes/No* type, we treat it as a binary classification problem. For *Number* type, since its answers are highly related to numerical answers in the 3,129 most frequent answers set, we heuristically filter 285 numerical answers from 3,129 answers before answer filtering. As for *Other* type, we preserve the original answer candidates without filtering.

After obtaining top K filtered answers, on one hand, they will be sent to the CLIP text encoder and dot-product with image features will be calculated, denoted as CLIP alignment score S_{CLIP} . On the other hand, we will calculate the question prior score S_{Question} (*i.e.*, cosine similarity between textual features, encoded by RoBERTa, of the question and answers) and the caption prior score S_{Caption}

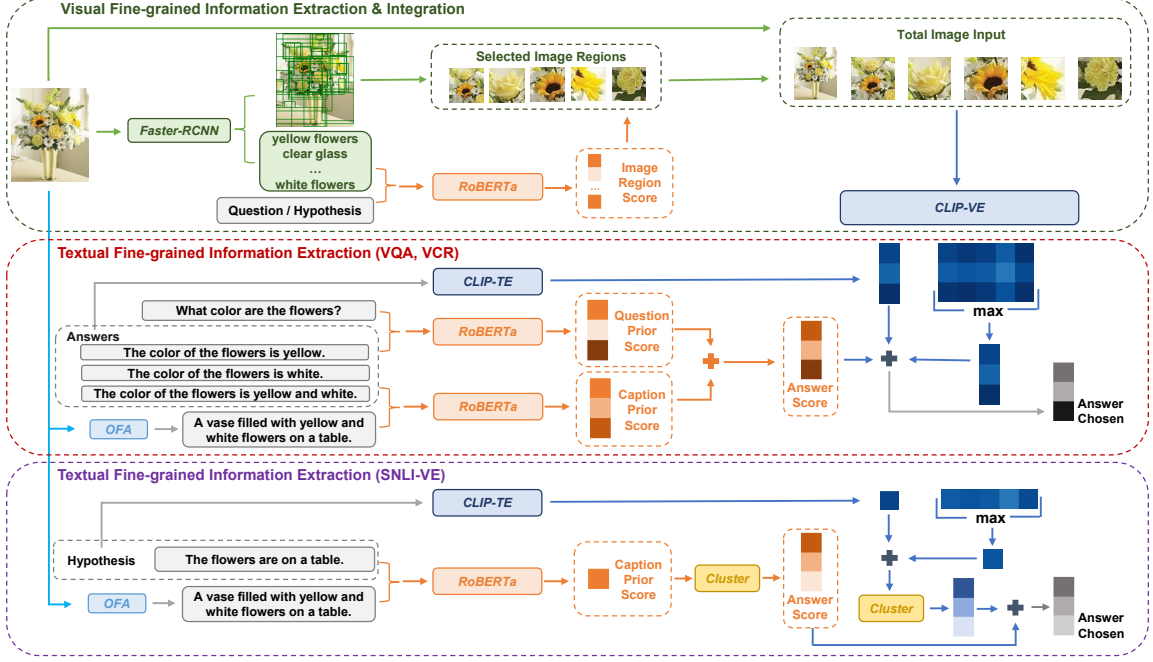


Figure 3: The overview of our proposed UniFine method. Visual fine-grained information is extracted under the guidance of the question (in VQA and VCR) or hypothesis (in SNLI-VE). In addition, textual fine-grained information is extracted by utilizing the question (in VQA and VCR) or hypothesis (in SNLI-VE) and the image caption. (Note: CLIP-VE denotes CLIP Visual Encoder and CLIP-TE denotes CLIP Text Encoder.)

(i.e., cosine similarity between textual features, encoded by RoBERTa, of image caption generated by OFA and answers). The whole process can be summarized as the following equations:

$$\begin{aligned}
 S_{\text{CLIP}} &= \mathbb{T}(A) \cdot \mathbb{V}(I) \\
 S_{\text{Question}} &= \cos(\mathbb{R}(Q), \mathbb{R}(A)) \\
 S_{\text{Caption}} &= \cos(\mathbb{R}(\mathbb{O}(I)), \mathbb{R}(A))
 \end{aligned} \quad (2)$$

where \mathbb{V} and \mathbb{T} are image and text encoders of CLIP, \mathbb{R} is RoBERTa, \mathbb{O} is OFA, and $\cos(\cdot, \cdot)$ means cosine similarity. I denotes images including one global image I_g and N selected image regions $\{I_l \in \text{Reg}\}$. Q and A correspond to the question and its top K filtered answers. $\mathbb{O}(I)$ means image caption generated by OFA.

In the end, all scores are ensembled. We select the answer with the highest score as zero-shot prediction result:

$$\begin{aligned}
 \arg \max_i \{ & S_{\text{CLIP}}(A_i, I_g) + k_1 \cdot \max_{I_l \in \text{Reg}} \{ S_{\text{CLIP}}(A_i, I_l) \} \\
 & + k_2 \cdot S_{\text{Question}}(Q, A_i) + k_3 \cdot S_{\text{Caption}}(I_g, A_i) \}, \quad (3)
 \end{aligned}$$

where k_1 , k_2 , and k_3 are hyperparameters.

Visual Commonsense Reasoning: VCR

VCR is similar to VQA since both of them are in question-answering formats. However, there are only four answer choices per question, which means we don't need to do answer filtering. $Q2A$

and $QA2R$ are two subtasks of VCR. $Q2A$ is similar to VQA in that there is only one question per sample. So the process of $Q2A$ is the same as VQA except for omitting answer filtering. $QA2R$ aims to dig out the rationale why one correct answer is chosen in $Q2A$ question. Since there is no question text in $QA2R$ and the correct answer is provided, we directly utilize the correct answer as the question text. Other procedures in $QA2R$ are the same as $Q2A$.

Visual Entailment: SNLI-VE

The task format of SNLI-VE is different from VQA and VCR. For each sample, only one image premise I and one hypothesis H are given, without answer candidates. It is a three-way classification problem, aiming to predict the relation between the image premise and hypothesis text into one of three classes: *Entailment*, *Contradiction*, and *Neutral*.

Since there are no answer candidates, we cannot directly compare CLIP alignment scores of answers to select the best answer, as in VQA and VCR. To tackle that, we compute the CLIP alignment scores between image and hypothesis of each sample in whole evaluation set, and cluster those scores into three clusters with three centroids. We rank the centroids from high to low and sequentially treat them as entailment centroid C_{CLIP}^e , neutral centroid C_{CLIP}^n and contradictory centroid C_{CLIP}^c . The detail of

clustering can be found in Algo. 1. It’s noted that, to make cluster centroids meaningful, an assumption is required: three relationships are uniformly distributed in the evaluation dataset. That assumption is true in SNLI-VE but not guaranteed in other less-calibrated datasets. We can measure how close S_{CLIP} of each sample is to each centroid:

$$\text{Dis}(C_{\text{CLIP}}^i, S_{\text{CLIP}}) = \|C_{\text{CLIP}}^i - (S_{\text{CLIP}}(H, I_g) + k_1 \cdot \max_{I_i \in \text{Reg}} \{S_{\text{CLIP}}(H, I_i)\})\| \quad (4)$$

where centroid $C_{\text{CLIP}}^i \in \{C_{\text{CLIP}}^e, C_{\text{CLIP}}^n, C_{\text{CLIP}}^c\}$.

Besides the CLIP alignment score comparison, we can obtain the caption prior score $S_{\text{Caption}}(I, H)$ using the image caption generated by OFA. Same as above, we also use the clustering method in Algo. 1, with only changing CLIP score to caption score, to get three centroids $\{C_{\text{Caption}}^e, C_{\text{Caption}}^n, C_{\text{Caption}}^c\}$. And we measure how close S_{Caption} of each sample is to each centroid:

$$\text{Dis}(C_{\text{Caption}}^i, S_{\text{Caption}}) = \|C_{\text{Caption}}^i - S_{\text{Caption}}(I, H)\| \quad (5)$$

It’s noted that due to the lack of answer candidates, we can’t get the question prior score S_{Question} . In the end, we ensemble two distances and predict the relationship by picking the closest centroid:

$$\arg \min_i \{ \text{Dis}(C_{\text{CLIP}}^i, S_{\text{CLIP}}) + k_2 \cdot \text{Dis}(C_{\text{Caption}}^i, S_{\text{Caption}}) \} \quad (6)$$

4 Experiments

In this section, we will talk about benchmark comparison first to show our strong performance. Then, we conduct extensive ablation studies to confirm the effectiveness of fine-grained information.

4.1 Experimental setup

Datasets. We analyze three vision-language tasks in our paper. For each of them, we utilize the validation set of VQAv2 (Goyal et al., 2017), VCR (Zellers et al., 2019), and SNLI-VE (Xie et al., 2019). More details about the validation set can be found in Sec. A.1. In VQAv2, we employ *vqa scores* to evaluate the model. In VCR and SNLI-VE, we use the accuracy of the validation set for evaluation.

Models. The core component of our method is CLIP¹. There are different variants of CLIP since we can use different models to act as the image or text encoder. Following previous work, we leverage CLIP Res50x16 and CLIP ViT-B/16 in VQA

¹<https://github.com/openai/CLIP>

Methods	VQA Answer Types			
	Yes/No	Number	Other	All
CLIP-ViL				
w/ CLIP _{Res50x16}	56.16	9.76	1.39	23.07
w/ CLIP _{ViT-B/16}	53.89	7.67	0.70	21.40
TAP-C* (Baseline)				
w/ CLIP _{Res50x16}	68.9(71.7)	25.9(18.7)	16.7(18.2)	37.5(38.4)
w/ CLIP _{ViT-B/16}	68.6(71.4)	25.4(21.0)	16.7(18.6)	37.3(38.7)
UniFine-Base (Ours)				
w/ CLIP _{Res50x16}	69.69	29.61	19.85	39.87
w/ CLIP _{ViT-B/16}	69.49	29.34	20.15	39.91
UniFine-Large (Ours)				
w/ CLIP _{ViT-L/14@336px}	70.36	29.95	20.19	40.33
Random	50.00	-	-	18.80

Table 1: Zero-shot VQAv2 results on the validation set. * denotes our reimplement. Reported results from TAP-C are in the bracket.

Methods	SNLI-VE Answer Types				VCR Tasks	
	C	N	E	All	Q2A	QA2R
Baseline						
w/ CLIP _{ViT-B/16}	67.59	18.66	55.92	47.37	53.24	46.51
UniFine-Base						
w/ CLIP _{ViT-B/16}	68.08	28.55	51.67	49.41	54.97	50.72
UniFine-Large						
w/ CLIP _{ViT-L/14@336px}	68.29	29.57	52.68	50.16	58.48	51.88
Supervised						
w/ EVE-Image	71.04	70.55	73.10	71.56	-	-
w/ R2C			-		63.8	67.2
Random	-	-	-	33.33	25.00	25.00

Table 2: Zero-shot SNLI-VE and VCR results on the validation set. (C: Contradiction, N: Neutral, E: Entailment)

for comparison. Since we are the first to evaluate CLIP’s zero-shot ability in SNLI-VE and VCR, there is no need for us to compare them with prior work. So we just exploit CLIP ViT-B/16 in VCR and SNLI-VE. We believe the scale of the model will have a big impact on the result, so we also utilize CLIP ViT-L/14@336px in VQA, VCR, and SNLI-VE to see how much improvement can be obtained by using a larger model. In addition to CLIP, we also use T5-large² for task format conversion, OFA-base³ for image captioning, RoBERTa-large⁴ for the following calculation of cosine similarity, and Faster-RCNN⁵ for object detection.

4.2 Benchmark comparison

VQA. Results of zero-shot VQA are reported in Tab. 1. For a fair comparison, we compare our

²<https://huggingface.co/models>

³<https://github.com/OFA-Sys/OFA>

⁴<https://github.com/UKPLab/sentence-transformers>

⁵<https://github.com/peteanderson80/bottom-up-attention>

method with two CLIP-based methods. We choose TAP-C (Song et al., 2022) as our baseline. Since the author didn't release the code, we reimplement it from scratch. After reimplementation, we obtain a lower score than TAP-C. There might be some reasons like differences in specific prompt design and answer filtering process making our result different from the original one. Although our reimplemented results are lower than the reported ones, we can surpass TAP-C after extracting and exploiting visual and textual fine-grained information. Compared to our reimplemented results, our method can elevate the performance of all answer types. By using a larger CLIP model, we can achieve better performance. Our best performance can surpass the reimplemented and TAP-C result by 2.83% and 1.63%. Currently, our method outperforms previous CLIP-based methods for zero-shot VQA.

SNLI-VE. We report the results of SNLI-VE in Tab. 2. By using the baseline method, we can get an accuracy of 47.37% in all categories, which is 14.04% higher than random performance. This result reveals that our baseline method is strong and it confirms CLIP's zero-shot ability in SNLI-VE. By extracting fine-grained information and upscaling the model, we can increase accuracy by 2.79% at most. For each answer type, *Neutral* type increases the most (+10.91%) and *Entailment* type decreases by 3.24%. We need to note that *Neutral* type is more complex than *Entailment* and *Contradiction* since this type is not as clear as the other two types requiring a model's deeper reasoning ability. The improvement in *Neutral* type shows the significance of fine-grained information. As for the decrement in *Entailment* type, it is likely due to the deficiency of our clustering method, which should be improved in the future. Since there is no CLIP-based zero-shot method for SNLI-VE before, we choose the supervised method EVE-Image from SNLI-VE paper (Xie et al., 2019) for comparison. Although the overall performance is still not comparable to the supervised method, our result of *Contradiction* type is approaching EVE-Image.

VCR. The results of VCR are reported in Tab. 2. We carry out experiments in two VCR sub-tasks, namely *Q2A* and *QA2R*. Compared to the random performance of *Q2A* and *QA2R*, our baseline method can improve by 28.24% and 21.51% respectively. The improvement confirms CLIP's strong zero-shot ability for VCR. By extracting fine-grained information and using a larger model,

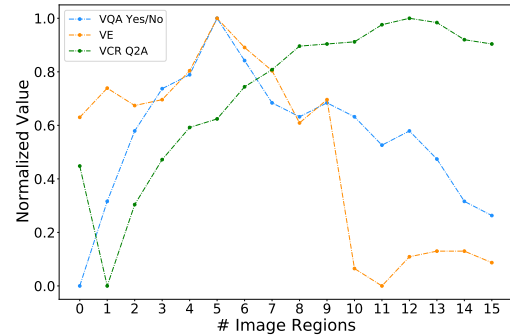


Figure 4: Normalized value to the number of image regions.

we can improve the baseline by 5.24% and 5.37% at most, which proves the effectiveness of our proposed method. There is no prior CLIP-based method for zero-shot VCR so we select the supervised model R2C, proposed in VCR paper (Zellers et al., 2019), for comparison. Although we cannot surpass the supervised model, the result of *Q2A* is approaching R2C and our results are competitive.

4.3 Ablation studies

In this section, we will analyze every important component of our proposed method. In Tab. 3, we can see all of the fine-grained (FG) information can help zero-shot learning and all fine-grained (FG) information combined together can bring more improvement.

Textual FG Information - Question: By adding the question prior information, we can see it can help VCR the most. We think the first reason is the question and answer in VCR are longer and more complex than the other two datasets. Consequently, the question and answers can provide more useful and richer information in zero-shot inference. Secondly, the correct answer is likely to have more overlap with the question. Plus, we can observe that question doesn't help a lot in VQA *Yes/No* answer type since this is a binary classification problem and a large number of questions are like "Is this A or B?" type, which cannot provide more useful information to zero-shot prediction.

Visual FG Information - Image Region: We can observe that the image region can largely improve the performance of *Other* answer type in VQA because the questions of this type tend to query the details of the image. And image regions can provide finer details to zero-shot inference. At the same time, we also find that the image region cannot help SNLI-VE much. We think SNLI-VE concentrates more on the global image thus image

Methods	VLU Tasks	Baseline	w/ Question	w/ Image Region (Faster-RCNN)	w/ Image Caption (OFA)	w/ All (Generation)	w/ Image Caption (GT)	w/ Image Region (GT)	w/ All (GT)	
	VQA	Yes/No	68.61	68.67 [+0.06]	68.80 [+0.19]	69.40 [+0.79]	69.49 [+0.88]	69.32 [+0.71]	-	69.54 [+0.93]
		Number	25.36	26.22 [+0.86]	26.21 [+0.85]	28.99 [+3.63]	29.34 [+3.98]	28.52 [+3.16]	-	28.88 [+3.52]
		Other	16.65	17.61 [+0.96]	19.00 [+2.35]	18.26 [+1.61]	20.15 [+3.50]	18.03 [+1.38]	-	20.04 [+3.39]
		All	37.33	37.94 [+0.61]	38.67 [+1.34]	38.90 [+1.57]	39.91 [+2.58]	38.69 [+1.36]	-	39.81 [+2.48]
UniFine-Base w/ CLIP _{Base}	SNLI-VE	C	67.59	-	68.59 [+1.00]	66.87 [-0.72]	68.08 [+0.49]	69.41 [+1.82]	-	69.67 [+2.08]
		N	18.66	-	20.94 [+2.28]	19.84 [+1.18]	28.55 [+9.89]	24.87 [+6.21]	-	25.50 [+6.84]
		E	55.92	-	53.13 [-2.79]	56.43 [+0.51]	51.67 [-4.25]	63.03 [+7.11]	-	62.90 [+6.98]
		All	47.37	-	47.54 [+0.17]	47.68 [+0.31]	49.41 [+2.04]	52.42 [+5.05]	-	52.66 [+5.29]
VCR	Q2A	53.24	54.60 [+1.36]	53.93 [+0.69]	53.35 [+0.11]	54.97 [+1.73]	53.42 [+0.18]	53.79 [+0.55]	54.72 [+1.48]	
	QA2R	46.51	50.10 [+3.59]	47.66 [+1.15]	46.79 [+0.28]	50.72 [+4.21]	46.60 [+0.09]	47.02 [+0.51]	50.16 [+3.65]	
	VQA	Yes/No	69.38	69.42 [+0.04]	69.75 [+0.37]	70.04 [+0.66]	70.36 [+0.98]	70.16 [+0.78]	-	70.41 [+1.03]
		Number	28.44	28.67 [+0.23]	28.64 [+0.20]	29.48 [+1.04]	29.95 [+1.51]	29.32 [+0.88]	-	29.72 [+1.28]
		Other	16.74	17.35 [+0.61]	19.03 [+2.29]	17.95 [+1.21]	20.19 [+3.45]	17.74 [+1.00]	-	20.06 [+3.32]
		All	38.07	38.41 [+0.34]	39.36 [+1.29]	39.05 [+0.98]	40.33 [+2.26]	38.97 [+0.90]	-	40.26 [+2.19]
UniFine-Large w/ CLIP _{Large}	SNLI-VE	C	67.57	-	68.44 [+0.87]	66.69 [-0.88]	68.29 [+0.72]	69.37 [+1.80]	-	70.30 [+2.73]
		N	25.17	-	20.97 [-4.20]	29.05 [+3.88]	29.57 [+4.40]	25.71 [+0.54]	-	25.44 [+0.27]
		E	51.57	-	55.14 [+3.57]	52.70 [+1.13]	52.68 [+1.11]	62.84 [+11.27]	-	62.84 [+11.27]
		All	48.05	-	48.15 [+0.10]	49.46 [+1.41]	50.16 [+2.11]	52.62 [+4.57]	-	52.84 [+4.79]
VCR	Q2A	56.92	58.12 [+1.20]	57.26 [+0.34]	57.11 [+0.19]	58.48 [+1.56]	57.01 [+0.09]	57.37 [+0.45]	58.29 [+1.37]	
	QA2R	48.06	51.20 [+3.14]	49.31 [+1.25]	48.44 [+0.38]	51.88 [+3.82]	48.16 [+0.10]	48.44 [+0.38]	51.30 [+3.24]	

Table 3: Ablation studies of zero-shot VQA, SNLI-VE, and VCR. (C: Contradiction, N: Neutral, E: Entailment, GT: Ground Truth, CLIP_{Base}: CLIP_{ViT-B/16}, CLIP_{Large}: CLIP_{ViT-L/14@336px})

regions can’t help a lot.

Textual FG Information - Image Caption: In Tab. 3, we can observe that the image caption can better assist the *Number* and *Other* answer type in VQA. For *Number* type, we think the image caption may contain numerical information which aids zero-shot prediction of *Number* type. Since there are a large number of questions in *Other* type, they will cover diverse question types, some of which may focus on information on an instance level. Normally, the image caption captures the instance-level information, so it can help VQA *Other* answer type. We can also notice that using image captions may hurt some categories of SNLI-VE, we think this result may suffer from the quality of the generated caption.

Generation vs. Ground Truth: Since not every dataset is well human-annotated, we employ these two settings to test the generalizability of our proposed method. In the generation setting, we generate image captions by OFA and detect objects by Faster-RCNN. In the ground truth setting, as mentioned above, there are ground truth captions paired with images in VQA and SNLI-VE. For VCR, images are not paired with human-annotated caption texts. However, 68% images of VCR validation set are the same as images in VisualCOMET (Park et al., 2020) and VisualCOMET is paired with the

ground truth caption. Thus, we directly leverage captions from VisualCOMET in VCR. Although images in VCR are not paired with captions, they are annotated with ground truth bounding boxes, so we have a ground truth image region experiment for VCR. However, VQA and SNLI-VE are not annotated with ground truth bounding boxes. As Tab. 3 shows, we can conclude that our method can work well in a situation without many annotations because we achieve similar performance in generation and ground truth scenarios, which confirms the generalizability of our proposed method.

Model Scale: We believe that the model scale will affect the final result since larger models are able to better process visual and textual information. In our experiments, we mainly focus on two variants of CLIP, namely CLIP ViT-B/16 and CLIP ViT-L/14@336px. We also carry out experiments on CLIP Res50x16 in VQA task, which can be found in Tab. 6. We can observe that larger models can elevate the performance and all of our best results are achieved by using CLIP ViT-L/14@336px.

Number of Image Regions: In this subsection, we would like to see how selected N image regions affect the zero-shot performance of different vision-language tasks. For convenience, we select *Yes/No* answer type of VQA, SNLI-VE, and *Q2A* task of VCR to carry out experiments. Full results

are reported in Tab. 8. For better visualization, we normalize the results. In Fig. 4, we can observe that with the increment of the image regions, the performance of all three tasks increases and then decreases. Moreover, selecting 5 image regions is optimal for VQA and VE. For VCR, 12 image regions are optimal. Visual fine-grained information can help CLIP and play an important role in the zero-shot prediction since it provides fine details of the image but more image regions after a certain point will result in a decrement. Too many image regions will introduce irrelevant visual information. In our experiments, we select 5 regions for VQA and SNLI-VE, and 12 regions for VCR.

5 Conclusion

In this work, we propose a unified and fine-grained approach for vision-language tasks including VQA, SNLI-VE, and VCR. We outperform previous CLIP-based methods for zero-shot VQA. Plus, we are the first to empirically study CLIP’s zero-shot ability for SNLI-VE and VCR, which achieves strong zero-shot performance. In addition to the benchmark comparison, we conduct extensive ablation studies confirming the significance of visual and textual fine-grained information and the generalizability of our proposed method.

Limitations

Although our proposed method is effective in three vision-language tasks, we still have some limitations. Firstly, we utilize T5 to convert the question-answering format into the declarative sentence in VQA and it works well in most cases, but it still faces out-of-coverage problems, which will affect the following zero-shot prediction of CLIP. We need to design more rules for these special cases for better conversion. Secondly, our clustering algorithm for SNLI-VE can achieve strong zero-shot performance, but the clustering centroids are close to each other and the algorithm is sensitive to these centroids. The robustness of this algorithm should be improved. What’s more, we leverage Faster-RCNN in visual fine-grained information extraction, so the detectable object attributes and classes are constrained in a relatively limited object set of Faster-RCNN, which may hinder further improvement from visual fine-grained information. The Faster-RCNN can be replaced with a better vision module. Besides, since we only utilize CLIP in our paper, we can explore the zero-shot ability of other

contrastive pre-training models in future work.

Ethics Statement

There are many large-scale pre-trained models used in our paper like OFA, T5, RoBERTa, and CLIP. Our method relies heavily on CLIP, which is pre-trained on approximately 400M image-text pairs crawled from the Internet. Since the pre-training dataset is noisy, CLIP is likely to have potential racial and gender bias. Therefore, if someone finds our work interesting and would like to use it in a specific environment, we suggest the user check the potential bias before application. We think one advantage of our work is we only utilize existing pre-trained models and we don’t need to train any new models. Compared to the energy-consuming model training, our method can be more environmentally friendly.

Acknowledgements

We thank anonymous reviewers for their comments. This work is supported by the DARPA MCS program under Cooperative Agreement N66001-19-2-4032.

References

- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*.
- 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.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2019. Uniter: Learning universal image-text representations.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference*

- on computer vision and pattern recognition, pages 248–255. Ieee.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- 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 *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 4904–4916. PMLR.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. 2020. Oscar: Object-semantic aligned pre-training for vision-language tasks. In *European Conference on Computer Vision*, pages 121–137. Springer.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Doll’ar, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32.
- Jae Sung Park, Chandra Bhagavatula, Roozbeh Mottaghi, Ali Farhadi, and Yejin Choi. 2020. Visualcomet: Reasoning about the dynamic context of a still image. In *European Conference on Computer Vision*, pages 508–524. Springer.
- Hieu Pham, Zihang Dai, Golnaz Ghiasi, Kenji Kawaguchi, Hanxiao Liu, Adams Wei Yu, Jiahui Yu, Yi-Ting Chen, Minh-Thang Luong, Yonghui Wu, et al. 2021. Combined scaling for open-vocabulary image classification. *arXiv preprint arXiv: 2111.10050*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28.
- Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, and Kurt Keutzer. 2021. How much can clip benefit vision-and-language tasks? *arXiv preprint arXiv:2107.06383*.
- Haoyu Song, Li Dong, Wei-Nan Zhang, Ting Liu, and Furu Wei. 2022. Clip models are few-shot learners: Empirical studies on vqa and visual entailment. *arXiv preprint arXiv:2203.07190*.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2019. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*.
- Sanjay Subramanian, Will Merrill, Trevor Darrell, Matt Gardner, Sameer Singh, and Anna Rohrbach. 2022. Reclip: A strong zero-shot baseline for referring expression comprehension. *arXiv preprint arXiv:2204.05991*.
- Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*.
- 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.

Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022a. [OFA: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 23318–23340. PMLR.

Zhecan Wang, Noel Codella, Yen-Chun Chen, Luwei Zhou, Xiyang Dai, Bin Xiao, Jianwei Yang, Haoxuan You, Kai-Wei Chang, Shih-fu Chang, et al. 2022b. Multimodal adaptive distillation for leveraging unimodal encoders for vision-language tasks. *arXiv preprint arXiv:2204.10496*.

Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. 2019. Visual entailment: A novel task for fine-grained image understanding. *arXiv preprint arXiv:1901.06706*.

Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.

Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6720–6731.

A Appendix

A.1 Data Statistics

VLU Tasks		Number of Questions	Number of Images
VQAv2	Yes/No	80541	-
	Number	28134	-
	Other	105679	-
	All	214354	40504
SNLI-VE	C	5939	-
	N	5960	-
	E	5959	-
	All	17858	1000
VCR	-	26534	9929

Table 4: Statistics of VQAv2, SNLI-VE, and VCR (C: Contradiction, N: Neutral, E: Entailment, VLU: Vision-Language Understanding)

Following previous work, we use the val2014 split of VQAv2. In zero-shot SNLI-VE and VCR, we use the validation set.

A.2 Answer filtering for VQA

Answer filtering. As in TAP-C (Song et al., 2022), we first manually design the demonstrations and employ T5 to convert the question-answering format of VQA into the declarative template with the $\langle extra_id_0 \rangle$ token. Then, we input the concatenation of demonstrations and declarative statements converted from question-answering format with the $\langle extra_id_0 \rangle$ token into the T5 encoder. Next, encoded features from the T5 encoder and answer candidates are input into the T5 decoder. At the end of the T5 decoder, it will calculate the probability of each answer candidate. We select Top K answers to replace $\langle extra_id_0 \rangle$ token in the template to generate K prompts, which will be fed into the CLIP text encoder.

Setting of hyperparameter K . Since we employ answer filtering to select top K answers, K is a significant hyperparameter. In Tab. 5, we show how the zero-shot performance of VQA *Number* and *Other* type varies with the increment of selected top K answers. We carry out six and seven experiments on these two types. We can observe that with the increment of K , the performance first increases and then decreases. When K is small, many correct answers are directly removed by T5, which makes it impossible for CLIP to choose the right answer. Conversely, if K is very big, there are too many answers, which are likely to disturb CLIP’s zero-shot prediction. In our experiments, we select the top 10 answers in VQA *Other* type and the top 4 answers in VQA *Number* type.

Algorithm 1 Pseudocode of clustering algorithm

Input: \mathbb{V} : CLIP image encoder, \mathbb{T} : CLIP text encoder, I : all images in SNLI-VE val split, H : all hypotheses in SNLI-VE val split, N : the number of samples in SNLI-VE val split;

Output: *centroid*.

```
1: dictionary centroid initialized to 0
2: array scores initialized to 0
3: // use CLIP to calculate dot product
4: for  $i = 0; i < N; i++$  do
5:    $scores[i] = \mathbb{T}(H[i]) \cdot \mathbb{V}(I[i])$ 
6: end for
7: // obtain three centroids
8:  $scores = \text{sort}(scores)$  // ascending order
9: // C: Contradiction, N: Neutral, E: Entailment
10:  $centroid[C] = \text{sum}(scores[:N/3]) / (N/3)$ 
11:  $centroid[N] = \text{sum}(scores[N/3:2N/3]) / (N/3)$ 
12:  $centroid[E] = \text{sum}(scores[2N/3:]) / (N/3)$ 
```

A.3 More ablation studies of VQAv2

Since previous work uses CLIP RN50x16 to do zero-shot VQA, we also conduct ablation studies on it. Results can be found in Tab. 6

A.4 Clustering algorithm and centroids of SNLI-VE

Algo. 1 is utilized in zero-shot SNLI-VE. After running the Algo. 1, we can obtain three clustering centroids. In fact, we can cache the centroids in advance. In order to achieve better performance, we tune the centroids, which are reported in Tab. 7. The effectiveness of Algo. 1 is based on the relatively even data distribution. K-Means⁶ can also be utilized here but it also requires a relatively even data distribution. In the validation split of SNLI-VE, we have 17858 samples, which are not divisible by 3. However, we can assume there are 5952, 5953, and 5953 samples in entailment, neutral, and contradiction category respectively.

A.5 How # image regions affect performance

The full results are reported in Tab. 8. They are values before normalization in Fig. 4. Through the table and figure, we can see how selected N images affects the zero-shot performance.

A.6 Zero-shot learning by only using textual fine-grained information

We think it is interesting to investigate the zero-shot performance if we only use textual fine-grained

⁶https://en.wikipedia.org/wiki/K-means_clustering

information. We only exploit the language model to accomplish zero-shot prediction in all three vision-language tasks. All results are shown in Tab. 9. In VQA, we use T5-large (for answer filtering) and RoBERTa-large. In SNLI-VE and VCR, we only utilize RoBERTa-large. Visual information is not considered and textual fine-grained information includes the image caption and question in this experimental setting. All results show that only using textual fine-grained information can achieve fair performance. (Note: We can notice that only using ground truth textual fine-grained information in SNLI-VE can surpass baseline performance. It is because the relation between the ground truth caption and hypothesis is well annotated in SNLI (Bowman et al., 2015))

Answer Type	2	3	4	5	9	10	11	20	40	200
Number	25.14	25.27	25.36	24.41	-	21.60	-	18.36	-	-
Other	-	-	-	14.98	16.60	16.65	16.61	15.54	14.93	10.87

Table 5: How # selected answers affect VQA

Methods	VLU Tasks	Baseline	w/ Question	w/ Image Region (Faster-RCNN)	w/ Image Caption (OFA)	w/ All (Generation)	w/ Image Caption (GT)	w/ Image Region (GT)	w/ All (GT)						
UniFine-Base w/ CLIP _{Res50x16}	Yes/No	68.85	68.90	[+0.05]	68.90	[+0.05]	69.63	[+0.78]	69.69	[+0.84]	69.59	[+0.74]	-	69.66	[+0.81]
	Number	25.85	26.20	[+0.35]	26.34	[+0.49]	29.26	[+3.41]	29.61	[+3.76]	28.91	[+3.06]	-	29.33	[+3.48]
	Other	16.74	17.36	[+0.62]	18.60	[+1.86]	17.86	[+1.12]	19.85	[+3.11]	17.64	[+0.90]	-	19.85	[+3.11]
	All	37.53	37.90	[+0.37]	38.53	[+1.00]	38.82	[+1.29]	39.87	[+2.34]	38.65	[+1.12]	-	39.82	[+2.29]

Table 6: Ablation Studies of zero-shot VQA. (GT: Ground Truth)

Answer Types	Methods					
	CLIP ViT-B/16		CLIP ViT-L/14@336px		RoBERTa-large	
	Baseline	w/ Image Region	Baseline	w/ Image Region	Only Image Caption (Generation)	Only Image Caption (GT)
Contradiction	0.23	0.47	0.17	0.37	0.22	0.29
Neutral	0.26	0.54	0.22	0.45	0.34	0.48
Entailment	0.27	0.55	0.23	0.46	0.43	0.60

Table 7: Clustering centroids of SNLI-VE (GT: Ground Truth).

Task	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
VQA Yes/No	68.61	68.67	68.72	68.75	68.76	68.80	68.77	68.74	68.73	68.74	68.73	68.71	68.72	68.70	68.67	68.66
VE	47.37	47.42	47.39	47.40	47.45	47.54	47.49	47.45	47.36	47.40	47.11	47.08	47.13	47.14	47.14	47.12
VCR Q2A	53.24	52.68	53.06	53.27	53.42	53.46	53.61	53.69	53.80	53.81	53.82	53.90	53.93	53.91	53.83	53.81

Table 8: How # image regions affect the performance of different tasks

Methods	VQA Answer Types				SNLI-VE Answer Types				VCR Tasks	
	Yes/No	Number	Other	All	C	N	E	All	Q2A	QA2R
Baseline w/ CLIP _{ViT-B/16}	68.61	25.36	16.65	37.33	67.59	18.66	55.92	47.37	53.24	46.51
Only textual fine-grained information (Generation)	67.13	28.77	13.44	35.64	64.83	17.32	56.19	46.09	35.13	41.93
Only textual fine-grained information (GT)	67.58	28.22	13.36	35.69	68.72	24.03	63.00	51.89	32.69	40.21

Table 9: Zero-shot performance when only using textual fine-grained information. (C: Contradiction, N: Neutral, E: Entailment, GT: Ground Truth)

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
6
- A2. Did you discuss any potential risks of your work?
7
- A3. Do the abstract and introduction summarize the paper’s main claims?
1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

4

- B1. Did you cite the creators of artifacts you used?
1, 2, 3, 4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
We will release the code and artifacts after we get the decision on the paper.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
We didn’t discuss it but our use is consistent with their intended use since they are open-sourced. Plus, we will specify the intended use when we get the decision on paper. We will release the code and provide a description of it.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
We don’t discuss it since we don’t focus on data collection and model pre-training. We conduct zero-shot learning experiments. We will consider it in the future.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
We will not release our code and documentation until we get the decision on the paper.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
I report the statistics of the validation set of datasets I use in the appendix section of the paper.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

C Did you run computational experiments?

4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

Our paper focuses on zero-shot learning. We don't need too many computational resources so we don't report it.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

4

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

4

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

4

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.