
Supplementary Material

Anonymous Author(s)

Affiliation

Address

email

1 A More details of Multimodal Pipeline

2 A.1 Image Encoder

3 In this work, we focus on patch features and apply vision transformer based models(ViTs) by [15] as
4 our visual encoder backbones. We split input image into a sequence of patches and adopt the linear
5 projection embedding of patch features v_q , which simplifies the step for fusing with text embedding.
6 Following the vision-language models, we train our model with multiple popular ViTs to examine the
7 influence of image encoder in OOD detection backpropagation process.

8 A.2 Knowledge Retrieval Module

9 Given the caption S , we parse it into triplets in the form of $T^{id} = \langle o(c), r(c), o'(c) \rangle$, where $o(c)$
10 and $o'(c)$ are concepts $\in C^{id}$, and $r(c)$ is the relation(s) between them, i.e., $\langle \text{man, riding, bicycle} \rangle$.
11 In our example, the seed triplets(ID triplets) parsed from the caption are $\langle \text{man, riding, bicycle} \rangle$ and
12 $\langle \text{bicycle, down, street} \rangle$. Then we construct knowledge graph by bridging these triplets with external
13 open knowledge including domain and commonsense knowledge graphs, e.g., ConceptNet [13].
14 ConceptNet provides a large scale commonsense knowledge with over 21 million edges by 36 type
15 of relations connecting 8 million nodes, i.e., IsA, UsedFor, AtLocation. In this study, to complete
16 our knowledge graph, we collect concepts by querying from ConceptNet using $o(c)$, $o'(c)$ and rel_i
17 where $i \in [0, 36]$ and integrate extracted triplets to seed triplets. For example, given "street" as $o(c)$
18 and "AtLocation" as rel_i , we will extract the related concepts are located at street to form triple t_i .
19 Specifically, we query explicit knowledge triplets of $o(c)$ and $o'(c)$ from ConceptNet to form T^{cn} ,
20 i.e., $\langle \text{bicycle, used for, transport} \rangle$. Finally, these knowledge triplets $\in T = T^{cn} \cup T^{id}$ are encoded as
21 language features, such as l_j using a language encoder(e.g., BERT[3]).

22 A.3 Multimodal Fusion Encoder

23 We consider our proposed OOD detection layer as a plug module to different vision-language
24 architectures. In ViLT and CLIP, we sum the visual and textual embeddings with incorporating $s(\cdot)$
25 as in Eq 4 of the main paper, and pass them to a standard L-depth transformer. Considering the BLIP
26 backbones, we perform a two-stream transformer pipeline consisting of stacked multiple layers to
27 joint vision and knowledge textual representations. For each layer, we have self-attention unit and
28 merged cross-attention unit which integrates vision and knowledge semantic information and the
29 alignments across them, and a position-wise feed-forward network.

30 Moreover, we update the image and language embedding outputs of themselves previous layer as
31 queries and concatenate them together as keys and values. To further improve the performance of
32 attention function, we use a multi-head attention which is composed by multiple paralleled attention
33 function in each head. The feed-forward layer transform the outputs of multi-head attention through
34 two fully-connected layers with GeLU activation.

35 B More details of Training objectives

36 We mainly introduce our training objectives in our pipeline in this section, including image text
37 matching (ITM) and masked language modeling (MLM).

38 B.1 Image Text Matching

39 To incorporate both the vision and the language representations, we adopt ITM which is widely used
40 in previous VL studies. Given an image and text of triple pair $\langle v_q, l_j \rangle$, ITM predicts whether they are
41 matched as positive examples or not, and it is a binary classification problem with the loss function in
42 Equation 5 of the main paper. We assume that each image and ID triple pair $\langle v_q, l_j \rangle$, as a positive
43 example. The negative pairs are constructed through batch-sampling.

44 B.2 Masked Language Modeling

45 MLM utilizes vision features and text features of ID concepts and relations to predict the masked
46 tokens in the caption sentence S . Following most VL models, we randomly masked some tokens
47 in S replacing as y^{msk} and predict them with their visual and textual features. Since some tokens
48 are replaced with "[mask]", the OOD score $s(\cdot)$ is changed based on the random masks. Thus, we
49 incorporate $s(\cdot)$ to calculate the predicted probability for a masked token similar to Eq 4 of the main
50 paper, and MLM loss is written as

$$\mathcal{L}_{\text{mlm}} = \mathbb{E}_{(v, \hat{l}) \sim D} \mathcal{H}(y_{\text{mlm}}, p(v, \hat{l})) \quad (1)$$

51 where \mathcal{H} denotes the cross-entropy, y_{mlm} is a one-hot vector where the ground truth tokens are with
52 probabilities of 1, \hat{l} denotes the masked text.

53 C More Experiments

54 In this section, we show implement details, more experiments on ablation studies, and more qualitative
55 analysis of our proposed VK-OOD models.

56 C.1 Implement details

57 **Datasets.** Following the practical settings, we adopt the training strategies of pre-training on more
58 data and fine-tuning on downstream tasks. We pre-train on three datasets, including COCO [5],
59 Visual Genome [4], and SBU Captions [9] with total of 1M images and 6.8M image-caption pairs, as
60 approximate 30% less than the baseline (ViLT). Each caption is parsed to 1 - 3 triplets and augmented
61 with 5 external knowledge triplets. For downstream datasets, we use Flickr30k [10] and COCO
62 for image-text retrieval, VQA_{v2} [1] and OKVQA [8] for visual question answering and ablation
63 studies, and NLVR2 [14] for visual reasoning. We resize each image to the size of 224×224 by
64 center-cropping. In the merged attention module, each multimodal encoder layer consists of one
65 multi-head self-attention block and one feedforward block, and total number of identical layers is 12.

66 **Encoder backbones.** First, we retrieve explicit knowledge triplets in pre-processing, by using
67 ConceptNet Numberbatch¹. Next, for ViLT, we use RoBERTa [6] as text encoder and ViT-B/32 by
68 [11] as visual encoder. To scale with CLIP, we use CLIP-ViT-B/32 [11] as both backbones. Then, we
69 follow the BLIP design with BERT-base as text encoder and ViT-B/16 as visual encoder.

70 **Network training.** We pre-train the model for 10 epochs, and use AdamW optimizer designed by [7]
71 with the learning rate of $1e-4$ and weight decay of $1e-2$. The warm-up ratio of learning rate is 10% of
72 the total training steps, and the learning rate was decayed linearly to 0 in the rest steps. Then, we
73 finetune our model for 5 epochs with learning rate of $2e-4$ for all downstream tasks. In addition, we
74 apply RandAugment [2] as augmentation strategy in finetuning steps. We pre-train and fine-tune both
75 on 8 NVIDIA RTX 2080Ti GPUs, and inference on 1 NVIDIA RTX 2080Ti GPU.

¹<https://github.com/commonsense/conceptnet-numberbatch>

Model	COCO						F30k					
	R@1	TR	R@10	R@1	IR	R@10	TR	R@5	R@10	R@1	IR	R@10
ViLT	61.8	86.2	92.6	41.3	72.0	82.5	81.4	95.6	97.6	61.9	86.8	92.8
UNITER	64.4	87.4	93.1	50.3	78.5	87.2	85.9	97.1	98.8	72.5	92.4	96.1
ALBEF	73.1	91.4	96.0	56.8	81.5	89.2	94.3	99.4	<u>99.8</u>	82.8	96.7	98.4
VinVL	74.6	92.6	96.3	58.1	83.2	90.1	-	-	-	-	-	-
BLIP	80.6	95.2	97.6	63.1	85.3	91.1	96.6	99.8	100	87.2	97.5	98.8
Ours(ViLT)	73.8	91.4	96	52.4	81.3	90.1	85.9	97.1	97.6	80.1	94.6	96.7
Ours(CLIP)	69.8	87.5	93.6	48.8	78.5	82.5	92.3	98.4	99.5	79.8	92.1	96.4
Ours(BLIP)	80.7	95.1	96.8	62.9	84.8	92.8	96.4	99.6	99.8	86.3	97.1	98.8

Table 1: Detailed results of image-text retrieval tasks on COCO and Flickr30k datasets. Our model with different backbones outperforms other models and achieve the best and second-best results.

Model	Objectives	VQA	Flickr30k	
		test-dev	TR@1	IR@1
ViLT	ITM	70.6	82.1	65.6
ViLT	MLM	72.8	-	-
ViLT	ITM+MLM	74.2	88.1	74.1
VK-OOD	ITM	72.1	84.5	69.8
VK-OOD	MLM	73.4	-	-
VK-OOD	ITM+MLM	74.8	89.0	77.2

Table 2: Ablation study experiment results of VK-OOD model. ViLT is our implementation without explicit knowledge and OOD detection layer. ITM is image-text matching, and MLM is masked language modeling. Results on VQA are on test-dev set. Both downstream results are in zero-shot settings. The bold values mean the best model in the table. Comparing with the baselines, our model with OOD detection layer outperforms on all objectives with two datasets. Training on combinations of objectives improves model performance.

76 C.2 More experimental results

77 **Detailed results of image-text retrieval tasks.** We provided detailed results on COCO and F30K
78 datasets, as shown in Table 1. The model of OOD detection layer with ViLT has significant improve-
79 ments in image retrieval and text retrieval tasks. Overall, our model achieved the best and second-best
80 results on both datasets comparing to other SOTA models.

81 **Training Objectives with OOD detection Layer.** To evaluate our proposed model, we perform
82 more ablations with the default training settings of the baseline and our model mentioned in Section
83 3.3 of the main paper. We consider different combinations of train objectives and evaluate in zero-
84 shot settings. We observe our model performance on training objectives. Note that, ViLT is our
85 implementation with the same subset of training datasets. Our raw results are presented in Table
86 2. We train on pre-train datasets with \mathcal{L}_{itm} in Equation 7, \mathcal{L}_{mlm} in Equation 8 and \mathcal{L} in Equation 9.
87 The results in Table 2 show that training on image-text matching and masked language modeling is
88 beneficial for both downstream tasks comparing to the baseline model, especially, there is promising
89 improvements in image retrieval and text retrieval tasks. Thus, it is beneficial to train on both ITM
90 and MLM for filtering outlier concepts and improve performance on downstream tasks.

91 C.3 More Qualitative Analysis

92 As the feature maps shown in Fig 2 of the main paper, our model result demonstrates more clusters
93 can be identified over the multimodal features extracted by VK-OOD. Thus, it illustrates that our
94 model is able to detect outliers and cluster images closest to the corresponding μ_k with image and
95 explicit knowledge triplets. Figure 1a and Figure 1b are examples that the nearest images in each
96 cluster.

97 Figure 2 and Figure 3 show more qualitative examples of our multimodal alignment results of our
98 models. We visualize the multimodal attention maps on images corresponding to concept triplets
99 using Grad-cam designed by [12]. Following our model architecture, the caption is parsed and
100 integrated with knowledge triplets. The right bottom subfigure in our model in Figure 2 and Figure 3



Figure 1: Example images in the clusters on COCO val set.

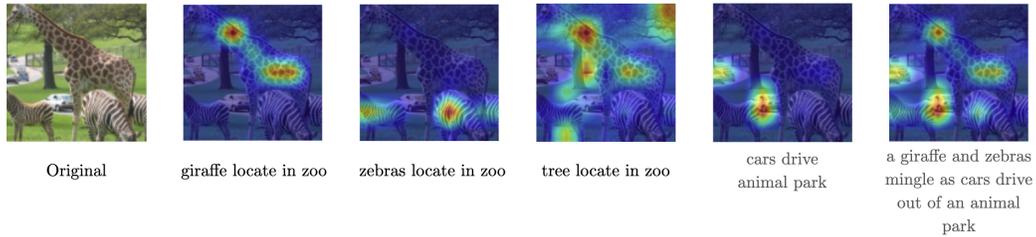


Figure 2: Visualization of the attention maps of image feature v_q and language features l_j of external knowledge alignment. The results are from our VKOOD-ViLT model. The original sample caption is “a giraffe and zebras mingle as cars drive out of an animal park”. We highlight areas in the example image corresponding to different knowledge triplets. Comparing with the attention maps of the baseline model, our model learns object shapes such as zebras and localize those objects correctly.

101 are the multimodal alignment of original captions from MSCOCO [5] dataset. Other subfigures show
 102 the alignments of extracted triplets on the image.

103 Interestingly, we find that our model is able to capture concept “plug” as a part of “refrigerator” or
 104 “microwave” in Figure 2. The heatmap area of “plug” and “microwave” in Figure 2 clearly suggest
 105 that our model has the capability to exploit different relevance between visual and corresponding
 106 conceptual text features. By contrast, the baseline results have not shown the relation between
 107 plug and microwave. In Figure 3, it shows that we detect three zebras comparing with baseline,
 108 but counting cars is not performing well as we expected – since the size (or scale) of cars is not
 109 sufficiently high, and moreover some parts of them are occluded.

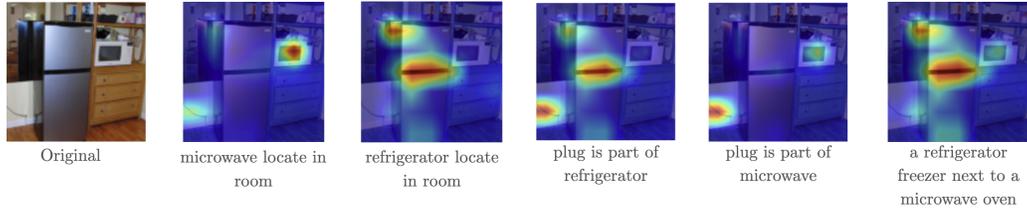


Figure 3: Visualization of the attention maps of image and knowledge concept triplets alignment. The results are from our VKOOD-ViLT model. The original sample caption is “a metallic refrigerator freezer next to a microwave oven”. We highlight areas in the example image corresponding to different knowledge triplets. Comparing with the attention maps of the baseline model, our model learns the relations between the parts (i.e., plug) of the objects correctly.

110 References

- 111 [1] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. Vqa: Visual
 112 question answering. In *Proceedings of the IEEE international conference on computer vision*,
 113 pages 2425–2433, 2015.
- 114 [2] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le. Randaugment: Practical automated data
 115 augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on*
 116 *Computer Vision and Pattern Recognition Workshops*, pages 702–703, 2020.
- 117 [3] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional
 118 transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- 119 [4] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li,
 120 D. A. Shamma, et al. Visual genome: Connecting language and vision using crowdsourced
 121 dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017.
- 122 [5] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick.
 123 Microsoft coco: Common objects in context. In *European conference on computer vision*, pages
 124 740–755. Springer, 2014.
- 125 [6] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer,
 126 and V. Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint*
 127 *arXiv:1907.11692*, 2019.
- 128 [7] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference*
 129 *on Learning Representations*, 2018.
- 130 [8] K. Marino, M. Rastegari, A. Farhadi, and R. Mottaghi. Ok-vqa: A visual question answering
 131 benchmark requiring external knowledge. In *Conference on Computer Vision and Pattern*
 132 *Recognition (CVPR)*, 2019.
- 133 [9] V. Ordonez, G. Kulkarni, and T. Berg. Im2text: Describing images using 1 million captioned
 134 photographs. *Advances in neural information processing systems*, 24, 2011.
- 135 [10] B. A. Plummer, L. Wang, C. M. Cervantes, J. C. Caicedo, J. Hockenmaier, and S. Lazebnik.
 136 Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence
 137 models. In *Proceedings of the IEEE international conference on computer vision*, pages
 138 2641–2649, 2015.
- 139 [11] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell,
 140 P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision.
 141 In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- 142 [12] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual
 143 explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE*
 144 *international conference on computer vision*, pages 618–626, 2017.

- 145 [13] R. Speer, J. Chin, and C. Havasi. Conceptnet 5.5: An open multilingual graph of general
146 knowledge. In *Thirty-first AAAI conference on artificial intelligence*, 2017.
- 147 [14] A. Suhr, S. Zhou, A. Zhang, I. Zhang, H. Bai, and Y. Artzi. A corpus for reasoning about natural
148 language grounded in photographs. *arXiv preprint arXiv:1811.00491*, 2018.
- 149 [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and
150 I. Polosukhin. Attention is all you need. In *Advances in neural information processing systems*,
151 pages 5998–6008, 2017.