

Less Is More: Linear Layers on CLIP Features as Powerful VizWiz Model

Fabian Deuser, Konrad Habel, Philipp J. Rösch, Norbert Oswald
University of the Bundeswehr Munich
ETTI - Institute for Distributed Intelligent Systems (VIS)

{fabian.deuser, konrad.habel, philipp.roesch, norbert.oswald}@unibw.de

Abstract

Current architectures for multi-modality tasks such as visual question answering suffer from their high complexity. As a result, these architectures are difficult to train and require high computational resources. To address these problems we present a CLIP-based architecture that does not require any fine-tuning of the feature extractors. A simple linear classifier is used on the concatenated features of the image and text encoder. During training an auxiliary loss is added which operates on the answer types. The resulting classification is then used as an attention gate on the answer class selection. On the VizWiz 2022 Visual Question Answering Challenge we achieve **60.15 %** accuracy on Task 1: Predict Answer to a Visual Question and AP score of **83.78 %** on Task 2: Predict Answerability of a Visual Question.

1. Introduction

Many new architectures were developed in recent years and applied to data sets like VQA_{v2}, GQA or VizWiz-VQA [2]. The VizWiz data set differs from other VQA data sets, because it has several issues in the data. Questions may not be answerable due to missing information in the images or the quality of the images may be extremely poor. Additionally the questions in the data set are not developed with a rigid set of rules, but are often colloquially. Last year’s winning team used an extension of OSCAR. They added an optical character recognition (OCR) module to the model and introduced reference image matching. Their final system is an ensemble of 96 models. While ensembles are important to achieve competitive results, they are extremely costly to train.

Our approach focus on simplicity and usability. We use pre-trained image and text encoders from CLIP [4] and train only a simple classification head. CLIP is based on CNN [3] respectively Vision Transformer [1] for image encoding and a Transformer [7] for text encoding. The CLIP model is pre-trained on 400 million image-text pairs with a contrastive




Figure 1. Our architecture for the VizWiz Challenge 2022.

objective to bring both modalities into the same embedding space. Since CLIP is trained with many samples, it also has OCR capabilities [4].

2. Methodology

The contribution of this paper is divided into (i) creating a suitable vocabulary for the classification task, (ii) using CLIP features with linear layers for VQA, and (iii) introducing an answer type gate to create a learnable masking.

Answer Vocabulary. The selection of appropriate answers has a major impact on the accuracy that can be achieved. Therefore, in this approach, the most common answer that returns the highest score per image-question pair is greedily selected. If this selection yields in several answers, the answer which appears most often in the whole training set is used. In case of a tie, the pairwise Levenshtein distance is used to find the answer that is most representative to all others. With this selection process the remaining number of answer candidates for training decreases to 5726.

CLIP-based Model. Previous CLIP-based models for VQA use the image encoder only [5] or generate prompts to match answers to questions [6]. Our approach utilises both

Model	VQA [Acc.]		Answerability [AP]	
	test-dev	test-std	test-dev	test-std
[R] RN50x64	60.73 %	59.40 %	82.74 %	82.54 %
[V] ViT-L/14@336px	60.66 %	59.01 %	83.50 %	82.86 %
[E] Ensemble	61.64 %	60.15 %	84.13 %	83.78 %

Table 1. Results in the VizWiz 2022 challenge.

image and text encoder. The resulting features are concatenated and passed to linear layers with layer normalisation and a high dropout value (0.5). As shown in Figure 1 answer types as well as the answers are predicted using an additional linear layer. Image size of the visual encoder is 448x448 for RN50x64 and 336x336 for ViT-L/14@336px. In both cases the linear classifier is trained using cross entropy loss with rotation as image augmentation. We train only the additional linear classifier and use the pre-trained CLIP model as image and text encoder. The CLIP part remains frozen and is not trained on the VizWiz data set, which allows fast and efficient training without large computational resources.

Answer Type Gate. We also introduce an auxiliary loss for answer type prediction. This loss helps to learn an answer masking for the eight answer types “other”, “numbers”, “yes”, “no”, “color”, “unsuitable” and “unanswerable”. The answer types are retrieved by regular expression matching from the best selected answer per image-question pair. The learned predictions for the answer types are linearly projected to a vector with the same dimension (5726) as the number of possible answer classes. After a sigmoid layer this vector is multiplied with the logits of the answer vocabulary. This enables to mask answers that do not correspond to the current answer type during inference. Both cross entropy losses, of the intermediate answer type prediction and the final answer classification, are weighted equally.

3. Conclusion

Our approach focuses on lightweight training by keeping the pre-trained CLIP backbone frozen, while still maintaining good accuracy. The OCR capabilities of CLIP, the large amount of pre-training data, and the multi-modality make CLIP an excellent feature extractor for this task. Unlike previous publications, the text Transformer is also used from CLIP. Although it was trained on alt-texts, it could be shown that meaningful representations of the questions are extracted without any fine-tuning. On the VizWiz VQA task we reach **59.40 %** with a single model and **60.15 %** with an ensemble of the RN50x64 and ViT-L/14@336px. On the answerability task we achieve **82.86 %** with a single model

and **83.78 %** with an ensemble.

Acknowledgement

The authors gratefully acknowledge the computing time granted by the Institute for Distributed Intelligent Systems and provided on the GPU cluster Monacum One at the University of the Bundeswehr Munich.

References

- [1] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 1
- [2] Danna Gurari, Qing Li, Abigale Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. Vizwiz grand challenge: Answering visual questions from blind people. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3608–3617, 2018. 1
- [3] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 1
- [4] 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, 2021. 1
- [5] Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, and Kurt Keutzer. How much can CLIP benefit vision-and-language tasks? In *International Conference on Learning Representations*, 2022. 1
- [6] Haoyu Song, Li Dong, Wei-Nan Zhang, Ting Liu, and Furu Wei. Clip models are few-shot learners: Empirical studies on vqa and visual entailment. *arXiv preprint arXiv:2203.07190*, 2022. 1
- [7] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 1