

VLATTACK: Multimodal Adversarial Attacks on Vision-Language Tasks via Pre-trained Models (Supplementary Materials)

Anonymous Author(s)

Affiliation

Address

email

1 A. Details of VL Models

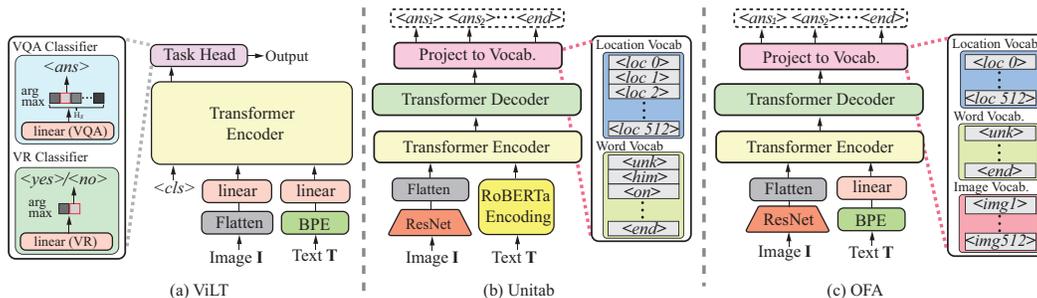


Figure 1: An illustration of ViLT, Unitab, and OFA model structures.

This section gives the details of ViLT, Unitab, and OFA models, and their structures are illustrated in Figure 1

- ViLT.** We select ViLT [1] as the **encoder-only** VL model because of its succinct structure and prominent performance on multiple downstream tasks. Given an input image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ and a sentence \mathbf{T} , ViLT yields M image tokens using a linear transformation on the flattened image patches, where each token is a 1D vector and $M = \frac{HW}{P^2}$ for a given patch resolution (P, P) . Word tokens are encoded through a Byte-Pair Encoder (BPE) [2] and a word-vector linear projection. Then tokens of two modalities and a special learnable token $\langle cls \rangle$ are concatenated. By attending visual and text tokens and a special token $\langle cls \rangle$ in a Transformer encoder with twelve layers, the output feature from the $\langle cls \rangle$ token is fed into a task-specific classification head for the final output. Taking the VQA task as an example, the VQA classifier adopts a linear layer to output a vector with H_s elements, where H_s is the number of all possible choices in the closed answer set of the VQA task. The final output is obtained through the element with the highest response in the vector.
- Unitab.** Unitab adopts an **encoder-decoder** framework. It first embeds text \mathbf{T} via RoBERTa [3] and flats features after encoding image \mathbf{I} through ResNet [4]. The attached visual and text token features are then fed into a standard Transformer network [5] with six encoder layers and six decoder layers. Finally, the sequence predictions $[\langle ans_1 \rangle, \langle ans_2 \rangle, \dots, \langle end \rangle]$ are obtained auto-regressively through a projection head. The network stops regressing when an end token $\langle end \rangle$ appears. For different tasks, the output tokens may come from different pre-defined vocabularies. Given the REC task as an example, four tokens $[\langle loc x_1 \rangle, \langle loc x_2 \rangle], [\langle loc x_3 \rangle, \langle loc x_4 \rangle]$ will be selected from the location vocabulary, which forms the coordinate of a bounding box. As a result, these models can handle both text and grounding tasks.

Table 1: An illustration of all datasets and tasks evaluated in our paper.

Datasets	Task	Task description	Attack Model			Attack Modality	
			OFA	Unitab	ViLT	Image	Text
VQAv2	VQA	Scene Understanding QA	✓	✓	✓	✓	✓
SNLI-VE	VE	VL Entailment	✓			✓	✓
RefCOCO	REC	Bounding Box Localization	✓	✓		✓	✓
RefCOCOg	REC	Bounding Box Localization	✓	✓		✓	✓
RefCOCO+	REC	Bounding Box Localization	✓	✓		✓	✓
NLVR2	VR	Image-Text Pairs Matching			✓	✓	✓
MSCOCO	Captioning	Image Captioning	✓			✓	
ImageNet-1K	Classification	Object Classification	✓			✓	

- 25 • **OFA**. As shown in Figure 1 (c), OFA also adopts an **encoder-decoder** structure. Different from
 26 Unitab, it adopts the BPE to encode text and extends the linguistic vocabulary by adding image
 27 quantization tokens [6] $\langle img \rangle$ for synthesis tasks. *Note that the main difference between OFA*
 28 *and Unitab lies in their pre-training and fine-tuning strategies rather than the model struc-*
 29 *ture*. For example, in the pre-training process, Unitab focuses on learning alignments between
 30 predicted words and boxes through grounding tasks, while OFA captures more general represen-
 31 tations through multi-task joint training that includes both single-modal and multi-modal tasks.
 32 Overall, OFA outperforms Unitab in terms of performance improvement.

33 B. Dataset and Implementation

34 B.1 Tasks and Datasets

35 To verify the generalization ability of our proposed VLATTACK, we evaluate a wide array of popular
 36 vision language tasks summarized in Table 1. Specifically, the selected tasks span from text under-
 37 standing (visual reasoning, visual entailment, visual question answering) to image understanding
 38 (image classification, captioning) and localization (referring expression comprehension).

39 For each dataset, we randomly select **5K correctly predicted samples** in the corresponding valida-
 40 tion dataset to evaluate the ASR performance. All validation datasets follow the same split settings
 41 as adopted in the respective attack models. Because VQA is a multiclass classification task, we
 42 select a correct prediction only if the prediction result is the same as the label with the highest VQA
 43 score¹, and regard the label as the ground truth in Eq. (1). In the REC task, a correct prediction is
 44 considered when the Intersection-over-Union (IoU) score between the predicted and ground truth
 45 bounding box is larger than 0.5. We adopt the same IoU threshold as in Unitab [7] and OFA [8].

46 B.2 Implementation Details

47 For the perturbation parameters of images, we follow the setting in the common transferable image
 48 attacks [9, 10] and set the maximum perturbation σ_i of each pixel to 16/255 on all tasks except
 49 REC. Considering that even a single coordinate change can affect the final grounding results to a
 50 great extent, the σ_i of the REC task is 4/255 to better highlight the ASR differences among distinct
 51 methods. The total iteration number N and step size are set to 40 and 0.01 by following the projected
 52 gradient decent method [11], and N_s is 20. For the perturbation on the text, the semantic similarity
 53 constraint σ_s is set to 0.95, and the number of maximum modified words is set to 1 by following
 54 the previous text-attack work [12, 13] to ensure the semantic consistency and imperceptibility. All
 55 experiments are conducted on a single GTX A6000 GPU. The analysis of parameter selection can
 56 be found in Section B.3.

57 B.3 Parameter Sensitivity Analysis

58 We discuss the effect of different iteration numbers of N and N_s in VLATTACK. All experiments
 59 are conducted on the VQAv2 dataset and the ViLT model. The total iteration number N is set from

¹The VQA score calculates the percentage of the predicted answer that appears in 10 reference ground truth answers. More details can be found via <https://visualqa.org/evaluation.html>

60 10 to 80, N_s is set to $\frac{N}{2}$. As depicted in Figure 2(a), the ASR performance is dramatically improved
 61 by increasing N from 10 to 20 steps and then achieves the best result when $N = 40$.

62 We next investigate the impact of different initial iteration
 63 numbers N_s . We test N_s from 5 to 40, but the total iteration
 64 number N is fixed to 40. As shown in Figure 2(b), the ASR
 65 score reaches the summit when N_s is 5, and it is smoothly
 66 decreased by continually enlarging N_s . Considering that the
 67 smaller initial iteration number N_s increases the ratio of text
 68 perturbations, we set N_s as 20 to obtain the best trade-off be-
 69 tween attack performance and the naturalness of generated ad-
 70 versarial samples in our experiments.

71 C. More Ablation Results

72 In Section 5.4, we conduct an ablation study to show the effec-
 73 tiveness of each module in our model design on VQA, VE, and
 74 REC tasks. Here, we conduct additional ablation experiments
 75 for the remaining tasks, including visual reasoning, image cap-
 76 tioning, and image classification.

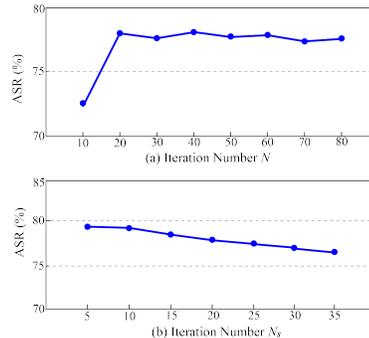


Figure 2: Investigation of iteration number N and N_s . (a) Various total iteration number N , where N_s is set to $\frac{N}{2}$. (b) Various initial iteration numbers N_s , where N is set to 40.

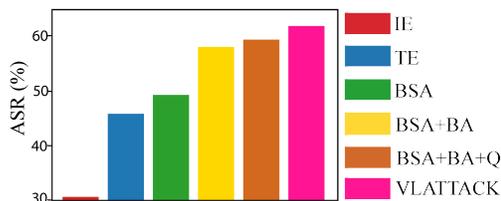


Figure 3: ViLT-VR.

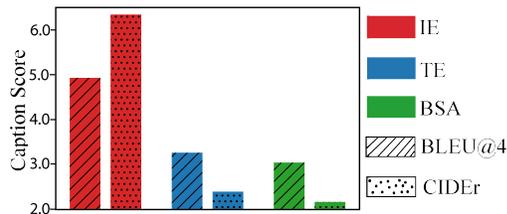


Figure 4: OFA-captioning.

77 Figure 3 shows the results of the ablation study on the VR task using the ViLT model. We can
 78 observe that only using the image encoder results in significantly low ASR. However, by combining
 79 it with the Transformer encoder (TE), BSA can achieve a high ASR. These results show the reason-
 80 ableness of considering two encoders simultaneously when attacking the image modality. The result
 81 of BSA+BA demonstrates the usefulness of attacking the text modality. Although BSA+BA+Q out-
 82 performs other approaches, its performance is still lower than that of the proposed VLATTACK. This
 83 comparison proves that the proposed iterative cross-search attack (ICSA) strategy is effective for the
 84 multimodal attack again.

85 Figure 4 shows the results of the image captioning task using the OFA model. Because the image
 86 captioning task only accepts a fixed text prompt for prediction, we only perturb the image and report
 87 the results on IE, TE, and BSA. For this task, we report BLEU@4 and CIDEr scores. **The lower,**
 88 **the better.** We can observe that the proposed BSA outperforms IE and TE, indicating our model
 89 design’s effectiveness.

90 Figure 5 shows the results of the image classification task using the
 91 OFA model. Similar to the image captioning task, we only attack
 92 images. The evaluation metric for this task is ASR. **The higher,**
 93 **the better.** We can have the same observations with other abla-
 94 tion studies, where attacking both encoders outperforms attacking a
 95 single encoder.

96 D. Different Optimization Methods

97 VLATTACK can be easily adapted to various optimization methods in image attacks. To demon-
 98 strate the generalizability of our method, we replace the projected gradient decent [11] in VLAT-
 99 TACK with Momentum Iterative method (MI) [14] and Diverse Input attack (DI) [15] since they

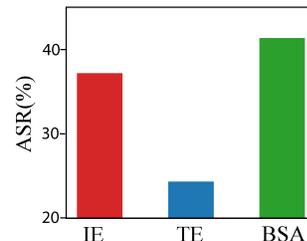


Figure 5: OFA-classification.

Table 2: Combining VLATTACK with different gradient-based image attack schemes.

Method	ViLT		Unitab			
	VQAv2	NLVR2	VQAv2	RefCOCO	RefCOCO+	RefCOCOg
BSA_{MI}	65.40	52.32	50.38	86.00	89.20	87.39
$VLATTACK_{MI}$	78.77	67.16	63.02	92.46	93.10	94.34
BSA_{DI}	65.94	52.30	42.74	90.30	91.56	91.00
$VLATTACK_{DI}$	78.07	67.50	61.22	93.98	94.04	94.76

100 have shown better performance than traditional iterative attacks [11, 16]. The replaced methods are
 101 denoted by BSA_{MI} , and $VLATTACK_{MI}$ using MI, BSA_{DI} and $VLATTACK_{DI}$ using DI, respec-
 102 tively. Experiments are developed on ViLT and Unitab. Results are shown in Table 2. Using MI
 103 and DI optimizations, BSA_{MI} and BSA_{DI} still outperform all baselines displayed in Table 1 in the
 104 main manuscript. Also, $VLATTACK_{MI}$ and $VLATTACK_{DI}$ outperform the image attack method
 105 BSA_{MI} and BSA_{DI} with an average ASR improvement of 9.70% and 9.29% among all datasets.
 106 The gain of performance demonstrates that the proposed VLATTACK can be further improved by
 107 combining with stronger gradient-based optimization schemes.

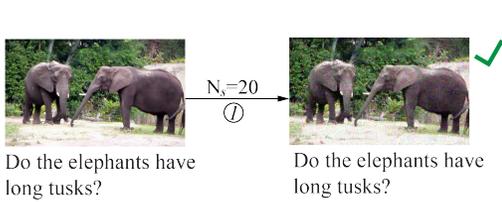


Figure 6: An adversarial image from BSA.

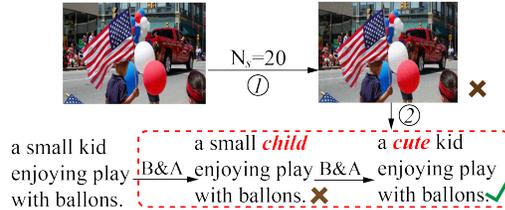


Figure 7: An adversarial sentence from text attack.

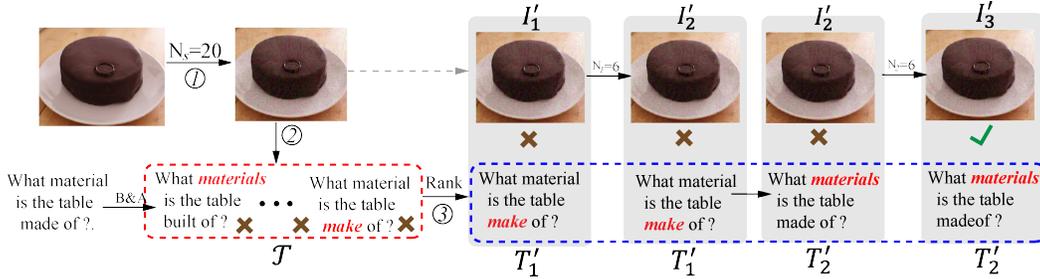


Figure 8: An adversarial image-text pair from multi-modal attack.

108 E. Case Study

109 E.1 How does VLATTACK generate adversarial samples?

110 The proposed VLATTACK aims to attack multimodal VL tasks starting by attacking single modalities.
 111 If they are failed, VLATTACK uses the proposed interactive cross-search attack (ICSA) strategy
 112 to generate adversarial samples. In this experiment, we display the generated adversarial cases from
 113 different attack steps, including the image modality in Figure 6, the text modality in Figure 7, and
 114 the multimodal attack in Figure 8.

115 **Single-modal Attacks (Section 4.1).** VLATTACK first perturbs the image modality using the pro-
 116 posed BSA and only outputs the adversarial image if the attack is successful (Algorithm 1 lines
 117 1-5). As shown in Figure 6, only attacking the image modality, VLATTACK can generate a suc-
 118 cessful adversarial sample to fool the downstream task. Then, VLATTACK will stop. Otherwise, it
 119 will perturb the text through BERT-Attack (B&A) and use the clean image as the input, which is
 120 illustrated in Figure 7 (Algorithm 1, lines 6-15). During the text attack, B&A will generate multiple
 121 candidates by replacing the synonyms of a word. Since the length of text sentences is very short in

122 the VL datasets, we only replace one word each time. From Figure 7, we can observe that B&A first
 123 replaces “kid” with its synonym “child”, but this is not an adversarial sample. B&A then moves to
 124 the next word “small” and uses its synonym “cute” as the perturbation. By querying the black-box
 125 downstream task model, VLATTACK succeeds, and the algorithm will stop.

126 **Multimodal Attack (Section 4.2).** If the single-modal attack fails, VLATTACK moves to the multi-
 127 modal attack by iteratively cross-updating image and text perturbations, where image perturbations
 128 are added through BSA, and text perturbations are added according to the semantic similarity. The
 129 cross-updating process is repeated until an adversarial image-text pair is found (Algorithm 1, lines
 130 16-24). Figure 8 shows an example. In Step ①, VLATTACK fails to attack the image modality
 131 and outputs a perturbed image denoted as I'_1 . In Step ②, VLATTACK also fails to attack the text
 132 modality and outputs a list of text perturbations \mathcal{T} . VLATTACK has to use the multimodal attack to
 133 generate adversarial samples in Step ③. It first ranks the text perturbations in \mathcal{T} according to the
 134 semantic similarity between the original text and each perturbation. The ranked list is denoted as
 135 $\{\hat{T}'_1, \dots, \hat{T}'_K\}$. Then it equally allocates the iteration number of the image attack to generate the
 136 image perturbations iteratively. In Figure 8, this number is 6, which means we run BSA with the
 137 budget 6 to generate a new image perturbation.

138 VLATTACK takes the pair (I'_1, \hat{T}'_1) as the input to query the black-box downstream model, where
 139 $\hat{T}'_1 =$ “What material is the table made of?”. If this pair is not an adversarial sample, then the pro-
 140 posed ICSA will adopt BSA to generate the new image perturbation I'_2 . The new pair (I'_2, \hat{T}'_1)
 141 will be checked again. If it is still not an adversarial sample, VLATTACK will use the next text pertur-
 142 bation $\hat{T}'_2 =$ “What materials is the table made of?” and the newly generated image perturbation I'_2
 143 as the input and repeat the previous steps until finding a successful adversarial sample or using up all
 144 K text perturbations in \mathcal{T} . VLATTACK employs a systematic strategy for adversarial attacks on VL
 145 models, sequentially targeting single-modal and multimodal perturbations to achieve successful ad-
 146 versarial attacks. Note that we miss one line “if $S(I'_{k+1}, T'_k) \neq y$ then return (I'_{k+1}, T'_k) ” between
 147 Lines 22 and 23 in Algorithm 1 of the main manuscript.

148 E.2 Case Study on Different Tasks

149 We also provide additional qualitative results from Figure 9 to Figure 14 for experiments on all
 150 six tasks. For better visualization, we display the adversarial and clean samples side by side in a
 151 single column. By adding pixel and word perturbations, the fidelity of all samples is still preserved,
 152 but predictions are dramatically changed. For instance, in the image captioning task of Figure 13,
 153 all generated captions show no correlation with the input images. Some texts may even include
 154 replacement Unicode characters, such as “\ufffd”, resulting in incomplete sentence structures.

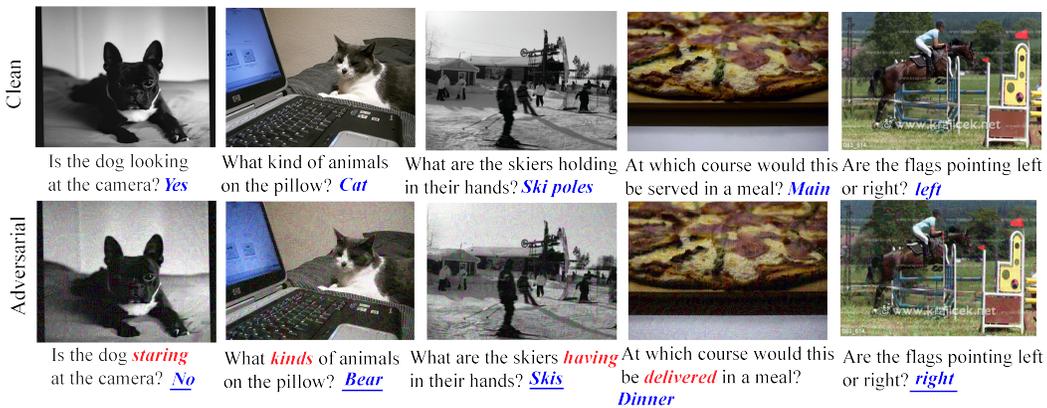


Figure 9: Additional quantitative results on visual question answering (VQA).

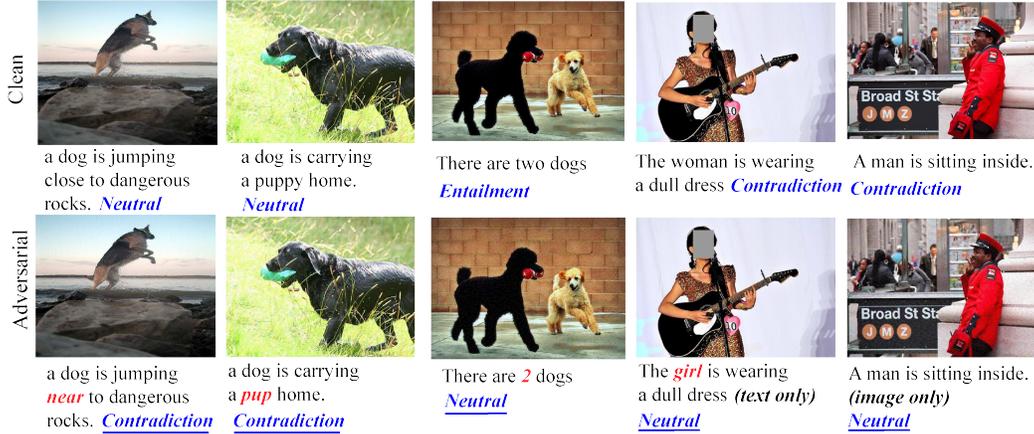


Figure 10: Additional quantitative results on visual entailment (VE).



Figure 11: Additional quantitative results on visual reasoning (VR).

155 F. Limitations

156 The limitations of our work can be summarized from the following two aspects. On the one hand,
 157 in our current model design, for the text modality, we directly apply the existing model instead of
 158 developing a new one. Therefore, there is no performance improvement on tasks that only accepts
 159 texts as input, such as text-to-image synthesis. On the other hand, our research problem is formul-
 160 ated by assuming the pre-trained and downstream models share similar structures. The adversarial
 161 transferability between different pre-trained and fine-tuned models is worth exploring, which we left
 162 to our future work.

163 G. Broad Impacts

164 Our research reveals substantial vulnerabilities in vision-language (VL) pretrained models, underlin-
 165 ing the importance of adversarial robustness cross pre-trained and fine-tuned models. By exposing
 166 these vulnerabilities through the VLATTACK strategy, we offer inspiration for developing more ro-
 167 bust models. Furthermore, our findings underscore the ethical considerations of using VL models
 168 in real-world applications, especially those dealing with sensitive information and big data. Overall,
 169 our work emphasizes the necessity of balancing performance and robustness in VL models, with
 170 implications extending across computer vision, natural language processing, and broader artificial
 171 intelligence applications.

172 References

173 [1] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without
 174 convolution or region supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning*

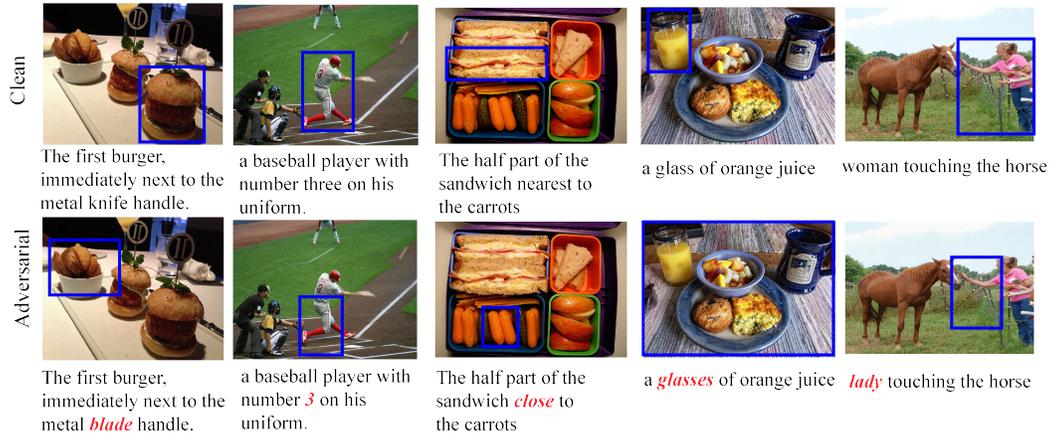


Figure 12: Additional quantitative results on referring expression comprehension (REC).

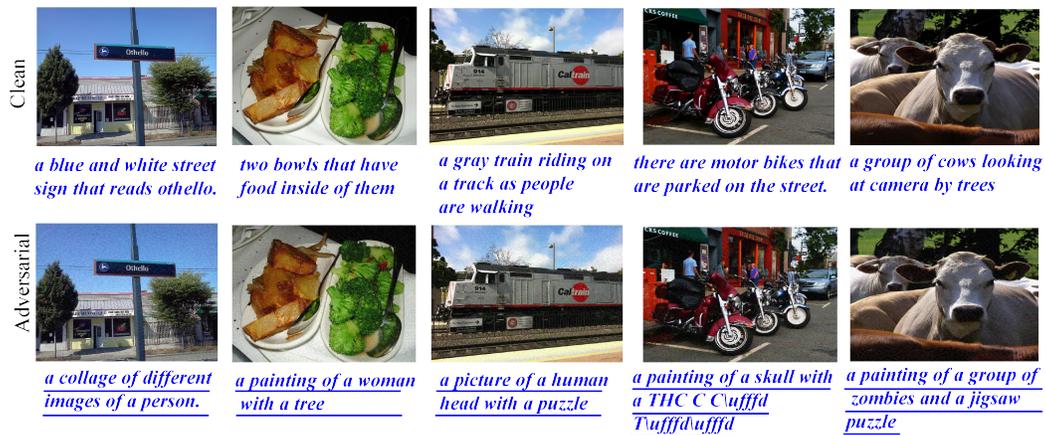


Figure 13: Additional quantitative results on the image captioning task.

- 175 *Research*, pages 5583–5594. PMLR, 2021.
- 176 [2] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words
177 with subword units. In *ACL (1)*. The Association for Computer Linguistics, 2016.
- 178 [3] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy,
179 Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT
180 pretraining approach. *CoRR*, abs/1907.11692, 2019.
- 181 [4] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
182 recognition. In *CVPR*, pages 770–778. IEEE Computer Society, 2016.
- 183 [5] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,
184 Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *CoRR*, abs/1706.03762, 2017.
- 185 [6] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution
186 image synthesis. In *CVPR*, pages 12873–12883. Computer Vision Foundation / IEEE, 2021.
- 187 [7] Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Faisal Ahmed, Zicheng Liu, Yumao
188 Lu, and Lijuan Wang. Unitab: Unifying text and box outputs for grounded vision-language
189 modeling. In *ECCV (36)*, volume 13696 of *Lecture Notes in Computer Science*, pages 521–
190 539. Springer, 2022.
- 191 [8] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou,
192 Jingren Zhou, and Hongxia Yang. OFA: unifying architectures, tasks, and modalities through
193 a simple sequence-to-sequence learning framework. In *ICML*, volume 162 of *Proceedings of
194 Machine Learning Research*, pages 23318–23340. PMLR, 2022.

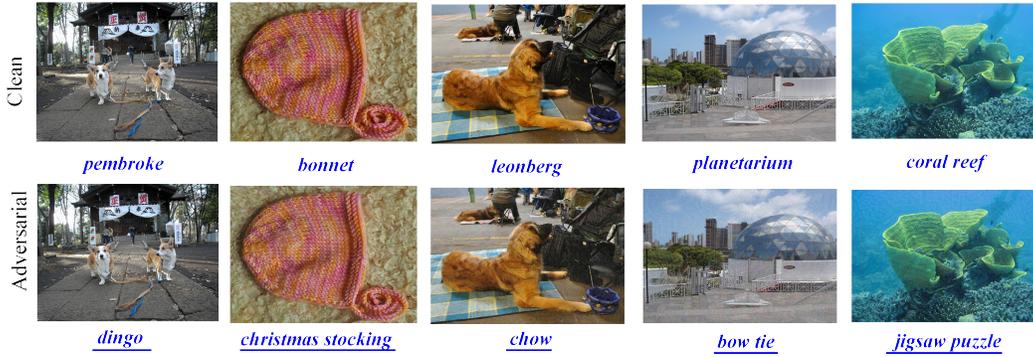


Figure 14: Additional quantitative results on the image classification task.

- 195 [9] Dongxian Wu, Yisen Wang, Shu-Tao Xia, James Bailey, and Xingjun Ma. Skip connections
 196 matter: On the transferability of adversarial examples generated with resnets. In *ICLR*. Open-
 197 Review.net, 2020.
- 198 [10] Zhipeng Wei, Jingjing Chen, Micah Goldblum, Zuxuan Wu, Tom Goldstein, and Yu-Gang
 199 Jiang. Towards transferable adversarial attacks on vision transformers. In *AAAI*, pages 2668–
 200 2676. AAAI Press, 2022.
- 201 [11] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian
 202 Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR (Poster)*. Open-
 203 Review.net, 2018.
- 204 [12] Lei Xu, Alfredo Cuesta-Infante, Laure Berti-Équille, and Kalyan Veeramachaneni. R&R:
 205 Metric-guided adversarial sentence generation. In *ACL/IJCNLP (Findings)*, pages 438–452.
 206 Association for Computational Linguistics, 2022.
- 207 [13] Jiaming Zhang, Qi Yi, and Jitao Sang. Towards adversarial attack on vision-language pre-
 208 training models. In *ACM Multimedia*, pages 5005–5013. ACM, 2022.
- 209 [14] Jiafeng Wang, Zhaoyu Chen, Kaixun Jiang, Dingkang Yang, Lingyi Hong, Yan Wang, and
 210 Wenqiang Zhang. Boosting the transferability of adversarial attacks with global momentum
 211 initialization. *CoRR*, abs/2211.11236, 2022.
- 212 [15] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L.
 213 Yuille. Improving transferability of adversarial examples with input diversity. In *CVPR*, pages
 214 2730–2739. Computer Vision Foundation / IEEE, 2019.
- 215 [16] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial machine learning at scale.
 216 In *ICLR (Poster)*. OpenReview.net, 2017.