

RoCOCO: Robustness Benchmark of MS-COCO to Stress-test Image-Text Matching Models

Seulki Park¹  Daeho Um²  Hajung Yoon² 
Sanghyuk Chun³  Sangdoo Yun³ 

¹University of Michigan ²Seoul National University ³NAVER AI Lab

Abstract. With the extensive use of vision-language models in various downstream tasks, evaluating their robustness is crucial. In this paper, we propose a benchmark for assessing the robustness of vision-language models. We believe that a robust model should properly understand both linguistic and visual semantics and be resilient to explicit variations. In pursuit of this goal, we create new variants of texts and images in the MS-COCO test set and re-evaluate the state-of-the-art (SOTA) models with the new data. Specifically, we alter the meaning of text by replacing a word, and generate visually altered images that maintain some visual context while introducing noticeable pixel changes through image mixing techniques. Our evaluations on the proposed benchmark reveal substantial performance degradation in many SOTA models (e.g., Image-to-Text Recall@1: 81.9% \rightarrow 48.4% in BLIP, 66.1% \rightarrow 37.6% in VSE ∞), with the models often favoring the altered texts/images over the original ones. This indicates the current vision-language models struggle with subtle changes and often fail to understand the overall context of texts and images. Based on these findings, we propose semantic contrastive loss and visual contrastive loss to learn more robust embedding. Datasets and code are available at <https://github.com/pseulki/rococo>.

Keywords: Image-Text Matching Models · Robustness · Stress-test

1 Introduction

Recently, vision-language (VL) models [38, 63] have generated significant research interest, with their learned visual-semantic features being actively employed across various domains. For instance, CLIP [63] latent features are used in image generation models [24, 60, 64], and CLIP scores are used to evaluate the alignment between generated images and texts [23, 67]. Similarly, BLIP-Diffusion [46] uses a pre-trained module from BLIP-2 [48] for controllable text-to-image generation. Additionally, pretrained CLIP models are used for semantic understanding in robotic manipulation [70]. Given the extensive use of VL models in various downstream tasks, evaluating the robustness of the models is imperative.

One of the primary methods for evaluating the VL models is through image-text matching (retrieval) tasks, which aim to align embedding features of images and texts effectively. It is worth noting that the above-mentioned applications

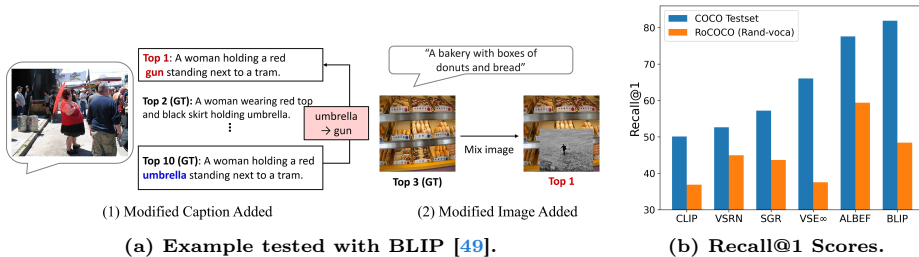


Fig. 1: Motivating Example. (a) When we add a new caption with only one word changed from “umbrella” to “gun”, this new caption is preferred by the model (Image-to-text). Likewise, when we add a new image created by inserting an unrelated image to the original one, this new image is ranked as top 1 (Text-to-image). (b) Recall@1 scores of existing SOTAs decrease when tested with our RoCOCO benchmark.

rely on image-text matching, which involves assessing the similarity between image features and text features. Many vision-language models assess image-text matching accuracy (*i.e.*, recall@1) on MS-COCO [56] or Flickr30K [61] benchmark datasets, with recent advancements showcasing notably strong performance [10, 16, 41, 49, 51, 75, 83]. However, there are still important questions remaining: How much can we trust these numbers? How robust is it when used in the wild?

While existing evaluation datasets are meticulously curated, they often lack the diversity and variability inherent in real-world data. Thus, it is challenging to assess whether a model robustly understands semantic and visual details using existing datasets. Figure 1a shows a motivating example of the retrieval results, when adding variants of texts and images to MS-COCO [56] test set, tested with the state-of-the-art (SOTA), BLIP [49]. When we retrieve the top-ranked text from a given image, the caption “A woman holding a red **gun** standing next to a tram” is preferred the most over the original caption “A woman holding a red **umbrella** standing next to a tram” by the model. Likewise, given the text, the image partially mixed with a completely unrelated image (*i.e.*, skiing on the snow) is retrieved as the highest similarity (top 1), instead of the correct clean image, which is ranked as top 3.

Although these variants are easily distinguishable by humans, we are surprised to find that SOTA models can be misled by such simple and explicit alterations. Considering the widespread use of image-text feature similarity across many downstream tasks, the lack of robustness and the neglect of such details can lead to various issues. If the highest text-image similarities are inaccurate, CLIP-based image generation evaluations [23, 67] might lose credibility, and there is a risk of incorrect robotic execution [70] due to a misunderstanding of commands. Lastly, these misleading results can greatly damage the trust in AI systems.

From this motivation, we propose a Robustness benchmark of MS-COCO (RoCOCO) to stress-test image-text matching models. **Our initial belief is that a robust model should properly understand both linguistic and visual semantics and be resilient to explicit variations in texts or images.**

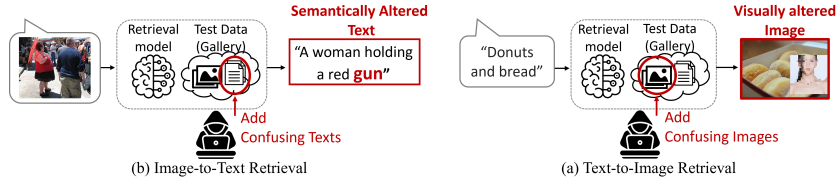


Fig. 2: Our Benchmark. By adding newly created confusing texts and images into the existing test set, our proposed benchmarks stress-test the model’s robustness in understanding visual and semantic details.

To explore this, we generate new sentences with minimal structural changes (*i.e.*, altering only one word) while changing its meaning (*e.g.*, “umbrella” → “gun”), as our language can generate an infinite number of different sentences by changing just one word. Similarly, for images, we contaminate the original images by mixing them with unrelated ones, thereby creating new images that retain some visual context but have disrupted visual coherence. This synthetic approach allows us to create controlled variations that can stress-test the model’s ability to handle diverse inputs.

These synthesized captions and images are added to the existing test set to evaluate the model’s capability to accurately retrieve the original ground-truth captions and images without being confused by the additional contents (See Figure 2). Despite the simplicity of the alteration, many state-of-the-art models show considerable performance degradation on the proposed benchmarks (*e.g.*, 81.9% → 48.4% in BLIP (ViT-B) [49], 66.1% → 37.6% in VSE ∞ (WSL Grid) [10] for Image-to-Text retrieval), as can be seen in Figure 1b. Our discovery reveals **that models trained on even hundreds of millions of data points can be susceptible to errors which are easily discernible by human observers**. In addition, our further analysis highlights the tendency of current image-text matching models to overlook subtle details and show more attention to specific words or image parts.

Finally, to address the vulnerabilities in understanding subtle details of image-text matching models, we propose Semantic Contrastive (SC) loss and Visual Contrastive (VC) loss. The SC loss encourages the model to learn semantic details by increasing the embedding distance between an original image or caption and a randomly corrupted caption. Similarly, the VC loss facilitates the effective learning of comprehensive visual information by enlarging the embedding distance between a randomly corrupted image and an original image or caption.

Our key contributions can be summarized as follows:

- We propose a robustness-evaluation benchmark for the image-text matching task, and discover significant performance drops across all models regardless of the extent of large-scale pre-training.
- We study vulnerabilities of image-text matching models and observe that these models often tend to focus on specific words or image components rather than comprehending the overall context.
- We propose Semantic and Visual Contrastive Losses to learn details in image and captions, enhancing the robust alignment in the embedding space.

2 Related Work

2.1 Image-Text Matching

Methods. Most image-text matching (ITM) methods [10, 16, 19, 22, 36, 44, 71, 76, 77] aim to learn joint visual-semantic embedding (VSE) such that paired image and text representation in the embedding space are close. In recent years, large-scale pre-training models [1, 7, 12, 13, 38, 42, 49, 50, 55, 57, 78, 83] have shown strong achievement in both zero-shot and fine-tuned performances. Most of these models adopt transformer architecture and can learn cross-modal representations benefiting from large-scale image-text pairs.

Datasets. New ITM benchmark datasets, such as Crisscrossed Captions (CxC) [58] and ECCV caption [15], have been proposed to extend MS-COCO. These datasets aim to improve associations and address false negatives in MS-COCO. However, our main focus differs as we aim to assess the vulnerability of models rather than providing improved benchmark datasets.

2.2 Robustness Test

Unimodal. Robustness in deep learning methods has been actively studied in both computer vision and natural language processing (NLP) areas. In computer vision, one research direction is data poisoning [5, 11, 29, 34, 72], which attacks the robustness of models during training by adding images with small perturbations. Meanwhile, adversarial attack studies [9, 17, 27, 30, 43] inject imperceptible noises to test images so that a model can make wrong predictions. For image retrieval task, Li et al. [47] showed that adding invisible noise to a query image can make the model return incorrect images. Another line of research [32, 33, 35] has proposed new ImageNet benchmarks for common robustness evaluation. In NLP, research on data poisoning [74] and adversarial attacks [2, 6, 20, 21, 25, 39, 45] has also been actively studied to fool the prediction of models. Adversarial examples are produced by character-level modifications [4], paraphrasing sentences [37], or substituting a word with a synonym [54, 65]. The main difference between these methods and our work is that while the previous works generate human imperceptible noises and texts that preserve their original meaning, we intentionally create images and texts that are perceptibly different and disrupt their original meaning. Our intuition is straightforward: a robust model should not be confused by such simple alterations, which are too obvious for humans.

Multimodal. As vision-language models have generated growing research interest, robustness work for cross-modal domain has been actively studied [8, 52, 59, 79]. Especially, in visual question answering (VQA), diverse robustness-evaluation benchmarks [26, 28, 53, 68, 69, 82] have been proposed. Recently proposed LANCE [62] shares similarity with ours in terms of manipulating words within text and introducing variations in images. LANCE generates counterfactual images with diverse transformations using a language guide by employing image-text models, such as CLIP [63] and BLIP [49]. Our purpose is different from LANCE as we aim to create a benchmark dataset for evaluating those models, CLIP and BLIP.

3 Robustness-Evaluation Benchmark

3.1 Observations motivating the proposed approach

Our goal is to quantitatively evaluate how well vision-language (VL) models understand both text and image. Specifically, we measure the robustness of a VL model through our proposed benchmark, which assesses how robustly the model retrieves the ground-truth image or caption instead of our visually altered image or semantically altered caption.

Based on the examples observed from the BLIP [49] model from Figure 1a, we generate altered captions and images that are capable of assessing the model’s vulnerability. Firstly, to assess the model’s ability for understanding the overall details between the image and text, we introduce semantically altered captions to make the image-to-text task challenging. Specifically, we create a semantically altered caption by replacing one word in the caption to change the meaning of the sentence. For example, replacing “umbrella” with “gun” as in Figure 1a (1).

Likewise, we generate a visually altered image with noticeable changes by mixing an unrelated image into an original image (Figure 1a (2)). Surprisingly, even though the visually altered image is easily discernible by humans, we observe that the VL model often favors a mixture of unintended images rather than the desired (ground-truth) ones. In the following sections, we describe the generation process in detail.

3.2 Semantically Altered Caption Generation

Source Word Selection via Embedding-Influence. To introduce discernable changes in the meaning of a caption, we focus on nouns for replacement. For effective confusion in a caption, we choose a word that has minimal impact on the embedding output. This idea is inspired by the common practice in which models measure similarity between the feature embedding of image and text encoders trained on image-text pairs [10, 49, 63]. We hypothesize that even with considerable changes in the semantic meaning, the model would be confused with the original (ground-truth) caption if the feature embedding changes little.

To estimate the influence of a word, we propose embedding-influence (EI) score. EI score measures the change in embedding when the word is removed from the caption. Given a text encoder f_T , and a caption $C = \{c_m \mid m = 1, \dots, M\}$, where M is the number of words in C , the embedding-influence (EI) score of a word, c_s , is defined by

$$EI(c_s) = 1 - \frac{\langle f_T(C), f_T(C \setminus c_s) \rangle}{\|f_T(C)\| \|f_T(C \setminus c_s)\|}, \quad (1)$$

where \langle, \rangle denotes the dot (inner) product operation. A low EI score means that the word has little influence on the embedding output of the caption. Given its limited influence compared to other words, substituting this word with a different word is expected to have low impact on the overall embeddings, which may lead to confusion between the modified caption and the original (ground-truth) caption. We demonstrate this choice in Section 4.3.

Table 1: Added Concept Groups. We include new concepts. Table shows added unique concepts and 3 random words from each group. Examples from Same-concept can test if the model can correctly recognize the semantic differences.

concept group	#concepts	concept lemmas (sampled)	examples from ‘Same-concept’ (Ground-truth / Semantically-altered)
material	32	metal, plastic, wooden	Five bagels are on a (metal / wood) rack.
color	28	black, white, brown	(Brown / Green) and black dog looking at a person holding a frisbee.
direction	50	front, top, bottom	A book sitting on (top / bottom) of a wooden desk.
vehicle_part	12	hood, wheel, tire	A cat by an overturned pot and a bicycle (wheel / timer).
shape	15	round, square, octagon	(Square/ Round) dishes hold main dishes while a banana is ...
event	11	Christmas, birthday, wedding	A table set for a traditional (Thanksgiving / Christmas) dinner.
number	14	one, five, hundreds	(Hundreds / Five) of people gathered in the park with ...

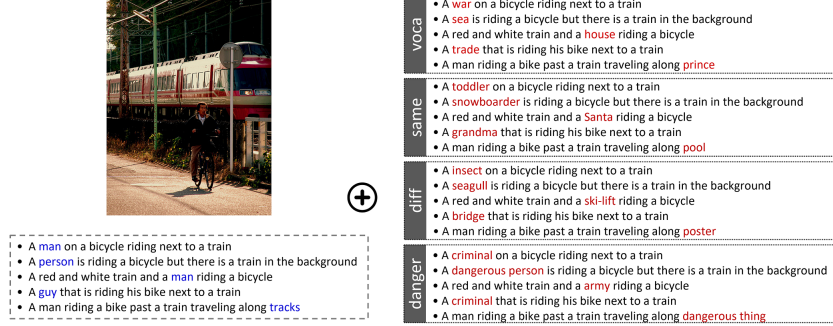
To ensure generalizability and not limit the choice of nouns to a specific model, we use four representative models in Table 2 (i.e., VSRN [51], CLIP [63], VSE ∞ [10], BLIP [49]), and measure the EI score of each word. We select a word with the least influence across the models. If the word is chosen by the majority of models, it is replaced by a target word (see Section 3.2). If there are multiple candidates of nouns, we randomly choose one. Interestingly, the words with the lowest embedding influence exhibit little variation across the models. We provide further details in Appendix Section A. In our evaluation, we not only include these four models but also add ten additional models not involved in noun selection process for a more comprehensive assessment.

Table 2: Models for EI scores.

Model	COCO-trained Text backbone	
VSRN	Finetuned	Bi-GRU [14]
CLIP	Zero-shot	Transformer [73]
VSE ∞	Finetuned	BERT [18]
BLIP	Finetuned	BERT [18]

Altered Word Selection for Diverse Benchmark Captions. To generate semantically altered captions covering various scenarios, we need to determine an “altered (replacement) word” replacing the source (original) word chosen in the previous section. To choose an altered word, we employ four different policies. First, we use concept groups from GRIT benchmark [31], which categorizes nouns from popular datasets including COCO into 24 concept groups such as food, people, and places. We add 7 concept groups (see Table 1) for words that are hard to be included in the given concept groups. Next, we categorize words according to the concept groups. Then, an original word is randomly replaced by any altered word inside **Same-concept** or **Diff-concept**. For example, **Same-concept** replaces “umbrella” with an altered word from the same concept (i.e., tools), which can be “rope” or “boxes”. **Diff-concept** replaces “umbrella” with an altered word selected randomly from different concepts, such as “pizza” from “food” concept, or “monkey” from “animal” concept.

Next, to stress-test with a wider range of words, we employ the BERT [18] vocabulary (**Rand-voca**). We randomly select an altered word consisting of only English letters, excluding those in other languages or special characters. Lastly, we create a special case (**Danger**) by using words related to public security. This allows us to evaluate the models’ ability to comprehend critical situations that could potentially pose a threat to human safety. For instance, we replace “umbrella” with “gun” or “weapon”. Examples of the semantically altered captions can be seen in Figure 3.



3.3 Visually Altered Image Generation

To generate visually altered images that preserve some visual context but have disrupted visual coherence, we employ two simple image mixing techniques. One is the Mixup-style approach [81], where two images are blended together in different proportion (Mix). The other method inserts a patch of an unrelated image onto the original image, as in CutMix [80] (Patch). For simplicity, the unrelated image is randomly selected from the COCO test set, but it can be selected from any other dataset. When inserting an unrelated image x^f into an original image x^o , we use two mixing ratios λ and \mathbf{M} for Mix and Patch, respectively, as follows:

$$\begin{aligned} \text{Mix} : \tilde{x} &= \lambda x^o + (1 - \lambda) x^f, \\ \text{Patch} : \tilde{x} &= \mathbf{M} \odot x^o + (\mathbf{1} - \mathbf{M}) \odot x^f, \end{aligned}$$

where $\mathbf{M} \in \{0, 1\}^{W \times H}$ denotes a binary mask at a randomly chosen location of the unrelated patch, where W is width and H is height of an image. The portion of 1 in \mathbf{M} is adjusted according to the mixing ratio of $\lambda = \frac{\sum_{i,j} \mathbf{M}_{i,j}}{W \times H}$. Figure 4 shows the examples of visually altered images. Creating these visually altered images and adding them to the gallery set provides an easy yet effective method to measure the robustness of the model.




Fig. 4: Example of visually altered images with different λ .

4 Experiments and Results

4.1 Experimental setting

In this section, we evaluate the existing image-text matching (ITM) models on our new dataset, RoCOCO. For Image-to-Text retrieval, we expand MS-COCO test data [40] by adding new 25,000 semantically altered captions to the existing 25,000 original captions, creating a gallery of 50,000 captions. We then retrieve text from this expanded gallery. Likewise, for Text-to-Image retrieval, we add new 5,000 visually altered images to the existing 5,000 original images, resulting in an image gallery of 10,000 images.

Evaluation Metrics. Recall@k, especially Recall@1 ($R@1$), is the most popular metric for evaluating the existing ITM methods. In this paper, we propose two metrics, *Drop Rate* and *Recall Score of Manipulated Samples* (RSMS) in addition to $R@1$. Drop rate measures the relative decrease in $R@1$ compared to the evaluation on the original COCO 5K test set. We calculate drop rate as $(R@1 - R_{New}@1)/R@1$. RSMS calculates the percentage of newly added semantically altered captions and visually altered images that are retrieved as top 1. This can quantitatively estimate the vulnerability of a model.

Models for Evaluation. We compare 14 state-of-the-art Image-Text Matching (ITM) models, whose trained weights are available to the public. They can be categorized into two groups; large-scale vision-language (VL) pre-training and visual semantic embedding groups. Large-scale VL pre-training group includes CLIP with ViT-B/32, ViT-B/16 and ViT-L/14 backbones [63], fine-tuned ALBEF [50], and zero-shot and fine-tuned BLIP with ViT-B and ViT-L backbones [49]. While ‘zero-shot’ and ‘fine-tuned’ models are both pre-trained on large-scale datasets, ‘zero-shot’ refers to not being fine-tuned with COCO train set. Visual semantic embedding group includes models using region features based on bottom-up attention [3] and SCAN [44]: VSRN [51], SAF, SGR [19], and VSE $_{\infty}$ [10].

4.2 Re-evaluation on RoCOCO

Image-to-Text Retrieval. Table 3 reports the image-to-text retrieval results on our new datasets. First, we observe the highest performance degradation on Rand-voca, with BLIP ViT-B showing a decrease of nearly 50%. This can be attributed to the fact that Rand-voca contains numerous unexpected words that do not commonly appear together in captions. This observation suggests that models are vulnerable to sentences comprising unfamiliar word combinations that rarely appear in the trained captions.

Furthermore, we can observe consistent degradation across all vision-language models, with drop rates ranging from 11% to 51% and RSMS from 13% to 55%, regardless of methods or the scale of pre-training datasets (e.g., 400M image pairs in CLIP [63], 129M in BLIP [49], 14M in ALBEF [50]). We assume that commonly used image-text matching loss might be vulnerable to a single-word change in the caption because the loss is used to minimize only the distance

Table 3: Image-to-Text retrieval results. Models are re-evaluated on our newly proposed RoCOCO datasets: Rand-voca, Same-concept, Diff-concept, and Danger. Recall@1 (R@1)(\uparrow), drop rate(\downarrow), Recall Score of Manipulated Samples (RSMS)(\downarrow) are shown. We can observe consistent degradation across all vision-language models.

	COCO 5K	Rand-voca		Same-concept		Diff-concept		Danger	
	R@1	R@1	drop rate	R@1	drop rate	R@1	drop rate	R@1	drop rate
Large-scale VL pre-training models									
CLIP ViT-B/32 (zero-shot) [63]	50.10	36.88	-26.39%	33.56	36.58	-26.99%	34.52	38.28	-23.59%
CLIP ViT-B/16 (zero-shot) [63]	52.44	40.16	-23.42%	31.26	39.84	-24.03%	31.68	40.56	-22.65%
CLIP ViT-L/14 (zero-shot) [63]	56.04	41.42	-26.09%	31.72	43.08	-23.13%	30.06	43.38	-22.59%
ALBEF [50]	77.58	59.38	-23.46%	26.76	60.92	-21.47%	24.94	60.44	-22.09%
BLIP ViT-B (zero-shot) [49]	70.54	34.84	-50.61%	55.62	48.52	-31.22%	37.26	43.30	-38.62%
BLIP ViT-B [49]	81.90	48.44	-40.85%	44.06	59.30	-27.59%	31.42	57.28	-30.06%
BLIP ViT-L (zero-shot) [49]	73.66	47.34	-35.73%	38.62	56.98	-22.64%	26.68	55.68	-24.41%
BLIP ViT-L [49]	82.36	64.56	-21.61%	24.00	70.46	-14.45%	16.88	70.32	-14.62%
Visual Semantic Embedding models									
VSRN [51]	52.66	44.98	-14.58%	16.28	47.38	-10.03%	12.32	48.36	-8.17%
SAF [19]	55.46	41.32	-25.50%	28.46	44.70	-19.40%	24.02	47.50	-14.35%
SGR [19]	57.22	43.66	-23.70%	26.86	45.68	-20.17%	23.78	48.00	-16.11%
VSE ∞ (BUTD region) [10]	58.02	32.72	-43.61%	46.76	41.78	-27.99%	31.60	37.90	-34.68%
VSE ∞ (BUTD grid) [10]	59.40	31.78	-46.50%	48.92	42.36	-28.69%	31.68	39.00	-34.34%
VSE ∞ (WSL grid) [10]	66.06	37.56	-43.14%	46.08	50.00	-24.31%	27.74	45.08	-31.76%



Fig. 5: Examples of incorrectly retrieved texts with BLIP from Same-concept (Image-to-Text). Captions with incorrect details are ranked at the top 1, while the correct captions are ranked lower.

between image-text pairs for learning multimodal representations. In addition, Figure 5 presents qualitative examples evaluated with BLIP (ViT-B) from Same-concept dataset. Our results highlight the importance of developing a robust training strategy for ITM models that can better capture word-level semantic meaning and align it with images.

Text-to-Image Retrieval. We evaluate VL methods on the new image set with $\lambda = 0.9, 0.8$ in Table 4. The images are generated using three random seeds, and the averaged results are reported. It can also be observed that all VL methods consistently exhibit performance degradation, with drop rates ranging from 9% to 39% and RSMS from 11% to 42%.

In addition, in Figure 6, we present examples of incorrect image retrievals using BLIP (ViT-B) when λ is set to 0.8. For instance, the face of the boy on the left is obscured, which is a crucial detail for relating to the query “Two black children.” While humans would generally prefer the ground-truth clean images over these mixed images, we observe that the models often overlook details and prefer the contaminated images. We argue that this evaluation is simple yet effective for assessing the robustness of the models. More results with different λ values can be found in Appendix (Section B).

Table 4: Text-to-Image retrieval. Models are evaluated with our new Mix and Patch with different λ . We report Recall@1 (R@1)(\uparrow), drop rate(\downarrow), Recall Score of Manipulated Samples (RSMS)(\downarrow), averaged over three random image generations with different seeds. We can see consistent degradation across all vision-language models.

	COCO 5K	Mix ($\lambda = 0.9$)			Mix ($\lambda = 0.8$)			Patch ($\lambda = 0.9$)			Patch ($\lambda = 0.8$)		
	R@1	R@1	drop rate	RSMS	R@1	drop rate	RSMS	R@1	drop rate	RSMS	R@1	drop rate	RSMS
Large-scale VL pre-training models													
CLIP ViT-B/32 (zero-shot) [63]	30.14	20.29	-32.68%	33.55	22.79	-24.39%	26.03	22.49	-25.38%	28.63	24.15	-19.87%	23.69
CLIP ViT-B/16 (zero-shot) [63]	33.03	20.05	-39.30%	39.00	23.57	-28.64%	29.88	22.58	-31.64%	35.18	24.70	-25.22%	29.41
CLIP ViT-L/14 (zero-shot) [63]	36.14	25.49	-29.47%	28.99	27.75	-23.22%	24.29	27.56	-23.74%	27.64	29.09	-19.51%	23.97
ALBEF [50]	60.67	44.13	-27.27%	26.60	48.02	-20.85%	21.11	48.86	-19.47%	19.58	51.80	-14.62%	15.30
BLIP ViT-B (zero-shot) [49]	56.36	39.03	-30.75%	31.54	43.94	-22.04%	22.28	41.96	-25.55%	27.56	45.05	-20.07%	22.79
BLIP ViT-B [49]	64.31	40.71	-36.70%	39.93	46.97	-26.96%	30.84	48.40	-24.74%	42.57	52.61	-18.19%	21.45
BLIP ViT-L (zero-shot) [49]	58.18	44.29	-23.87%	25.13	47.61	-18.17%	19.96	46.79	-19.58%	21.07	49.50	-14.93%	16.50
BLIP ViT-L [49]	65.06	41.87	-35.64%	42.45	48.92	-24.81%	33.91	48.55	-25.38%	29.17	49.50	-23.92%	22.10
Visual Semantic Embedding models													
VSRL [51]	40.34	27.04	-32.97%	39.05	31.36	-22.26%	28.87	30.08	-25.43%	31.11	32.50	-19.43%	24.80
SAF [19]	40.11	30.90	-22.96%	27.84	33.37	-16.80%	22.87	32.50	-18.97%	23.78	34.03	-15.16%	19.69
SGR [19]	40.45	30.71	-24.08%	28.08	33.41	-17.40%	22.57	32.40	-19.90%	23.95	34.08	-15.75%	19.90
VSE ∞ (BUTD region) [10]	42.46	31.57	-25.65%	30.74	35.61	-16.13%	20.45	34.17	-19.52%	23.51	36.48	-14.08%	17.28
VSE ∞ (BUTD grid) [10]	44.07	30.22	-31.43%	36.68	35.26	-19.99%	25.00	35.70	-18.99%	23.52	38.75	-12.07%	15.82
VSE ∞ (WSL grid) [10]	51.55	34.31	-33.44%	38.60	40.40	-21.63%	26.26	43.67	-15.29%	18.39	46.87	-9.08%	11.31

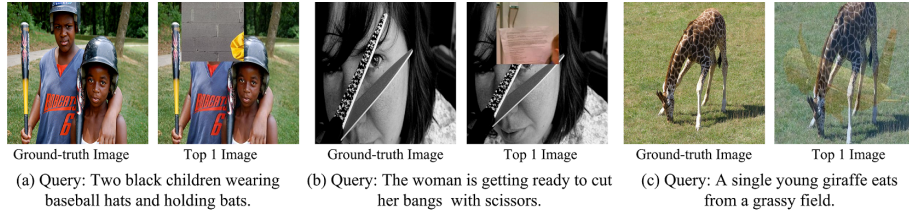


Fig. 6: Examples of images incorrectly retrieved by BLIP when $\lambda = 0.8$ (Text-to-Image). The first two examples are from the Patch, while the last one is from the Mix. In the Patch examples, some important areas are obscured. For instance, the face of the boy on the left is obscured, which is a crucial detail for relating to the query “Two black children.” In the Mix example, unrelated image of a “plane” is visible.

4.3 Analysis and Discussions

Each word within a caption has a different impact on the embedding. In Section 3.2, we introduce the Embedding-Influence (EI) score. Figure 7 illustrates the varying EI scores of words within each caption, with the red color indicating higher influence. The noun with the highest EI score is underlined in red, while the noun with the lowest score is underlined in gray. For instance, in the first caption, the noun “umbrella” has a relatively lower influence compared to the noun “tram,” even though it carries significant meaning within the caption. Therefore, replacing the word “umbrella” can lead to a notable change in semantic meaning without significantly affecting the original embedding.

A large woman holding a red umbrella standing next to a tram.
A man with a red helmet on a small moped on a dirt road.
A young girl inhales with the intent of blowing out a candle.
A man on a bicycle riding next to a train.
A kitchen is shown with a variety of items on the counters.
A bathroom that has a broken wall in the shower.

Fig. 7: Varying influence of a word in a caption. Red indicates higher influence.

Table 5: Effects of using EI scores. Deleting a source word with the lowest EI score results in the largest performance drop, causing the feature embedding to remain confusingly similar to the original.

	COCO	Random Deletion			High EI Deletion			Low EI Deletion		
	R@1(\uparrow)	R@1(\uparrow)	drop rate(\downarrow)	RSMS(\downarrow)	R@1(\uparrow)	drop rate(\downarrow)	RSMS(\downarrow)	R@1(\uparrow)	drop rate(\downarrow)	RSMS(\downarrow)
CLIP ViT-B/32 (zero-shot) [63]	50.10	38.58	-22.99%	29.66	42.76	-14.65%	21.84	36.04	-28.06%	32.30
CLIP ViT-L/14 (zero-shot) [63]	56.04	42.54	-24.09%	30.4	48.58	-13.31%	20.42	39.22	-30.01%	33.74
BLIP ViT-B (zero-shot) [49]	70.54	45.58	-35.38%	40.54	57.14	-19.00%	25.80	36.34	-48.48%	52.48
BLIP ViT-B [49]	81.90	65.54	-22.46%	19.98	72.74	-11.18%	14.06	59.28	-27.62%	30.10
VSRR [51]	52.66	44.7	-15.12%	18.02	43.46	-17.47%	22.56	38.56	-26.78%	29.36
VSE ∞ (BUTD region) [10]	58.02	34.2	-41.05%	45.58	40.58	-30.06%	38.06	30.02	-48.26%	50.72
VSE ∞ (BUTD grid) [10]	59.40	34.3	-42.26%	46.46	39.92	-32.79%	39.78	30.46	-48.72%	51.54
VSE ∞ (WSL grid) [10]	66.06	40.8	-38.24%	41.68	47.32	-28.37%	33.76	36.56	-44.66%	47.14




Fig. 8: Varying influence of spatial regions on the embedding. The similarity heatmap shows the cosine similarity between the embedding of the original image and the embedding when each part is masked. Dark colors indicate higher similarity (less influence), meaning the embedding changes little when that part is removed. White colors indicate important areas where significant changes occur. Even when specific parts are replaced with unrelated scenes (e.g., a kite), the model may overlook these changes because other, more influential parts (e.g., the boy’s face) remain.

Manipulating words with low EI scores proves to be an effective approach for semantic-shift confusion. To demonstrate this, we evaluate model performance by removing words in captions using different methods. “Random” method randomly removes a noun, while “Large EI” removes the noun with the highest EI score, and vice versa for “Low EI”. We create new captions by simply deleting the original word without replacement to mitigate the impact of the altered word. Table 5 shows that deleting words with low EI scores is the most effective approach for misleading the models, while deleting words with high EI scores results in minimal performance degradation. This finding supports our hypothesis that leveraging the influence of words on embedding features can effectively confuse the models. Thus, manipulating words with low EI scores can be an effective method for assessing the robustness of newly trained models.

The influence of each spatial part on the embedding varies within a single image. To examine why the model can be deceived by unrelated images, we analyze the impact of each spatial location on the feature embedding. We divide the image into 16 parts and mask each part to zeros, to observe the changes in the embedding. The heatmap in Figure 8 illustrates the cosine similarity between the embedding of the original image and the image embedding when each corresponding part is masked. In the cases where visually altered images are retrieved as top 1, we can observe that influential parts like “boy” or “toilet”

Table 6: Image-to-Text retrieval on dataset with multiple words substitutions. The results are averaged over generations with three different random seeds. Recall@1 (R@1)(\uparrow), drop rate(\downarrow), Recall Score of Manipulated Samples (RSMS)(\downarrow) are shown. Models can confuse sentences even when semantic meaning is largely damaged.

	COCO				2 words substitution				3 words substitution				4 words substitution				5 words substitution			
	R@1	R@1	drop rate	RSMS	R@1	R@1	drop rate	RSMS	R@1	R@1	drop rate	RSMS	R@1	R@1	drop rate	RSMS	R@1	R@1	drop rate	RSMS
Large-scale VL pre-training models																				
CLIP ViT-B/32 (zero-shot) [63]	50.10	42.89	-14.39%	19.71	46.07	-8.04%	12.67	47.45	-5.29%	8.15	48.37	-3.45%	5.46							
CLIP ViT-B/16 (zero-shot) [63]	52.44	45.35	-13.52%	19.07	48.43	-7.65%	11.89	49.97	-4.71%	8.01	50.61	-3.49%	5.95							
CLIP ViT-L/14 (zero-shot) [63]	56.04	47.35	-15.51%	22.18	50.22	-10.39%	15.78	51.99	-7.23%	11.56	53.07	-5.30%	8.27							
ALBEF [50]	77.58	72.43	-6.64%	2.40	73.03	-5.86%	0.88	73.23	-5.61%	0.43	73.26	-5.57%	0.32							
BLIP ViT-B (zero-shot) [49]	70.54	53.04	-24.81%	30.75	62.99	-10.70%	14.72	67.95	-3.67%	5.44	69.73	-1.15%	1.86							
BLIP ViT-B [49]	81.90	73.62	-10.11%	12.76	77.45	-5.43%	7.10	79.54	-2.88%	4.05	80.48	-1.73%	2.51							
BLIP ViT-L (zero-shot) [49]	73.66	60.35	-18.07%	21.66	67.99	-7.70%	10.16	71.63	-2.76%	3.93	72.87	-1.07%	1.61							
BLIP ViT-L [49]	82.36	73.93	-10.24%	12.65	77.93	-5.38%	7.45	79.81	-3.10%	4.23	80.98	-1.68%	2.54							
Visual Semantic Embedding models																				
VSRN [51]	52.66	45.07	-14.41%	17.79	47.89	-9.06%	11.33	49.89	-5.26%	7.08	50.99	-3.17%	4.29							
SAF [19]	55.46	44.06	-20.56%	20.29	47.22	-14.86%	26.71	50.02	-9.81%	15.12	51.71	-6.76%	10.85							
SCR [19]	57.22	43.57	-23.86%	28.53	46.98	-17.90%	22.79	49.81	-12.95%	17.49	51.91	-9.28%	13.09							
VSE ∞ (BUTD region) [10]	58.02	33.94	-41.50%	46.81	37.15	-35.98%	42.66	40.39	-30.39%	37.79	43.17	-25.60%	33.01							
VSE ∞ (BUTD grid) [10]	59.40	34.79	-41.44%	45.95	38.03	-35.98%	41.75	41.17	-30.68%	37.14	44.97	-24.30%	30.57							
VSE ∞ (WSL grid) [10]	66.06	39.95	-39.52%	43.79	44.04	-33.33%	38.44	48.29	-26.90%	32.85	51.73	-21.69%	27.51							

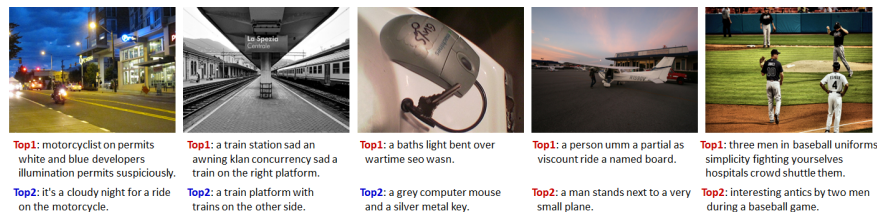


Fig. 9: Example of substituting four random words evaluated with BLIP (ViT-B). We find that the model can be confused by highly nonsensical and fragmented sentences. We hypothesize that the model focuses on specific words (e.g., “motorcycle” in the first image) rather than comprehending the language as a whole.

still remain despite obscuring some important parts like “kyte” or “a man’s face”. This finding indicates that certain parts of the image have a stronger impact on the retrieval outcomes than the other parts.

VL models can be misled by highly nonsensical sentences with multiple word replacements. To further investigate the vulnerability of VL models, we conduct experiments where two to five words in the captions are randomly replaced with words from the BERT vocabulary. Due to the brevity of some captions, we do not restrict the source words to nouns. Figure 9 presents examples of top-1 retrieval results from BLIP (ViT-B) for captions with four word replacements. Interestingly, these broken captions often include at least one correct keyword, such as “motorcyclist” in the first image. This suggests that the model may prioritize specific words over understanding the sentence as a whole. These findings are consistent with the observation by [79], that VL models behave like bags-of-words, ignoring word order in sentences. Our results with random word replacements further reveal that VL models may heavily rely on the presence of specific words, even when the sentences are nonsensical and fragmented. Additionally, the results in Table 6 quantitatively show meaningful

performance degradation across the model, when the original semantic meaning is significantly disrupted. Large-scale pretraining methods demonstrate better robustness than VSE models when multiple words are changed simultaneously.

5 Semantic and Visual Contrastive Loss for Robustness

Throughout our study, we have observed that VL models tend to overlook semantic and visual details. To address this issue, we propose the Semantic Contrastive (SC) loss and Visual Contrastive (VC) loss, which encourages the model to learn details in texts and images, respectively. Given a text encoder f_T , an image encoder f_I , an image x , a corrupted image $x_{corrupt}$, a caption c , and a corrupted caption $c_{corrupt}$, SC and VC losses are defined by:

$$L_{SC} = \frac{1}{2} \left(\underbrace{\frac{\langle f_T(c_{corrupt}), f_I(x) \rangle}{\|f_T(c_{corrupt})\| \|f_I(x)\|}}_{L(c_{corrupt}, x)} + \underbrace{\frac{\langle f_T(c_{corrupt}), f_T(c) \rangle}{\|f_T(c_{corrupt})\| \|f_T(c)\|}}_{L(c_{corrupt}, c)} \right), \quad (2)$$

$$L_{VC} = \frac{1}{2} \left(\underbrace{\frac{\langle f_I(x_{corrupt}), f_I(x) \rangle}{\|f_I(x_{corrupt})\| \|f_I(x)\|}}_{L(x_{corrupt}, x)} + \underbrace{\frac{\langle f_I(x_{corrupt}), f_T(c) \rangle}{\|f_I(x_{corrupt})\| \|f_T(c)\|}}_{L(x_{corrupt}, c)} \right). \quad (3)$$

Minimizing the losses will encourage the model to decrease the similarity (i.e., increase the difference) between the corrupted and original embeddings. Firstly, to improve robustness for semantic details, SC loss has two components. (1) $L(c_{corrupt}, x)$ reduces the similarity between the embeddings of the original images and their corresponding corrupted captions. This loss allows the model to capture shifts in meaning within similar sentences, with the aim of achieving a more solid alignment between the embeddings of the ground-truth images and text. (2) $L(c_{corrupt}, c)$ encourages the model to increase the difference between similar sentences, thereby capturing semantic details more effectively. Likewise, to learn visual details within images, VC loss has two components. (1) $L(x_{corrupt}, x)$ is designed to increase the distance between corrupted image and original image embeddings. (2) $L(x_{corrupt}, c)$ increases the distance between corrupted image and original caption embeddings to enhance the robust alignment between caption and visual details.

In each batch, we generate a corrupted caption $c_{corrupt}$ by randomly selecting words within the caption to be replaced with a probability of p (set to 0.3). These selected words are then substituted with random words from the BERT vocabulary with a probability of q (set to 0.5), or deleted with a probability of $1 - q$. Corrupted images are created with a 0.5 probability by randomly manipulating images within a batch using either the Mixup [81] or CutMix [80].

To demonstrate the effectiveness of the proposed Semantic and Visual Contrastive loss, we train BLIP ViT-B with the same training details as the original paper, except for the inclusion of SC and VC losses. As shown in Figure 10, employing SC and VC losses significantly improves Recall@1 and RSMS for both

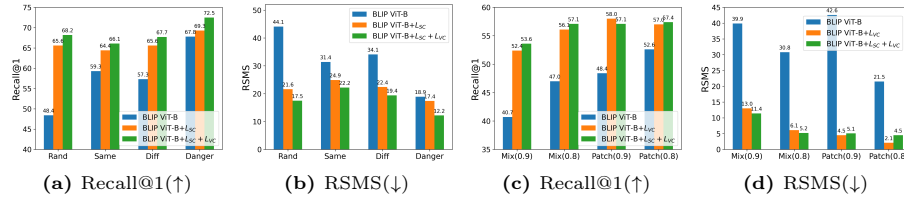


Fig. 10: Simple Semantic Contrastive Loss and Visual Contrastive Loss can significantly improve robustness of BLIP (ViT-B) model. (a) and (b): Image-to-Text Retrieval. (c) and (d): Text-to-Image Retrieval.

image-to-text (I2T) and text-to-image (T2I) retrieval tasks. Especially, a substantial improvement is observed in the Rand-voca setting (e.g., Recall@1: 48.4 \rightarrow 68.2 and RSMS: 44.1 \rightarrow 17.5). This improvement can be attributed to our strategy of constructing corrupted captions using random selections from the BERT vocabulary. In addition, interestingly, in image-to-text retrieval, using not only the SC loss, which is explicitly designed to learn semantic details, but also the VC loss, has led to further performance improvement. This observation suggests that enhancing the robustness of image embeddings can help training a robust alignment of embeddings across both modalities.

6 Conclusion and Discussion

The paper discovers that current vision-language models are vulnerable to capturing visual and semantic details, often focusing instead on specific words or spatial regions. To stress-test VL models, we propose a new methodology for creating a robustness benchmark and demonstrate consistent degradation across various state-of-the-art models using our RoCOCO dataset. To address this issue, we introduce semantic and visual contrastive losses. We hope that our findings provide insights for improving the robustness of VL models and for developing more diverse stress-test methods.

Limitations. One limitation is that the process of randomly replacing words can result in unnatural sentences, such as “A *war* on bicycle riding next to a train (man \rightarrow war).” Additionally, the altered images are not natural and can be considered out-of-distribution data. Therefore, future work should explore methods for generating more realistic and challenging texts and images, potentially using large language models like GPT-4 or generative models such as Stable Diffusion [66]. Another limitation is that the four models used in adversarial caption generation may be more susceptible to performance degradation. However, among the 14 models we evaluated, there was no significant difference between the 4 models used for dataset generation and the 10 models that were not involved in data generation. This could be attributed to the fact that many current models use common frozen language models like BERT. However, the trends in future models may differ, and further research should explore more generalizable methods for generating sentences.

Acknowledgments

I would like to express my appreciation to Professor Jin Young Choi for his invaluable guidance and support.

References

1. Alayrac, J.B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al.: Flamingo: a visual language model for few-shot learning. arXiv preprint arXiv:2204.14198 (2022) 4
2. Alzantot, M., Sharma, Y., Elgohary, A., Ho, B.J., Srivastava, M., Chang, K.W.: Generating natural language adversarial examples. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (2018) 4
3. Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L.: Bottom-up and top-down attention for image captioning and visual question answering. In: Proceedings of the IEEE conference on computer vision and pattern recognition (2018) 8
4. Belinkov, Y., Bisk, Y.: Synthetic and natural noise both break neural machine translation. In: International Conference on Learning Representations (2018) 4
5. Biggio, B., Nelson, B., Laskov, P.: Poisoning attacks against support vector machines. In: Proceedings of the 29th International Conference on International Conference on Machine Learning (2012) 4
6. Boucher, N., Shumailov, I., Anderson, R., Papernot, N.: Bad characters: Imperceptible nlp attacks. In: 2022 IEEE Symposium on Security and Privacy (SP) (2022) 4
7. Byun, J., Hwang, T., Fu, J., Moon, T.: Grit-vlp: Grouped mini-batch sampling for efficient vision and language pre-training. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel (2022) 4
8. Carlini, N., Terzis, A.: Poisoning and backdooring contrastive learning. In: International Conference on Learning Representations (2022) 4
9. Carlini, N., Wagner, D.: Adversarial examples are not easily detected: Bypassing ten detection methods. In: Proceedings of the 10th ACM workshop on artificial intelligence and security (2017) 4
10. Chen, J., Hu, H., Wu, H., Jiang, Y., Wang, C.: Learning the best pooling strategy for visual semantic embedding. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (2021) 2, 3, 4, 5, 6, 8, 9, 10, 11, 12, 22
11. Chen, X., Liu, C., Li, B., Lu, K., Song, D.: Targeted backdoor attacks on deep learning systems using data poisoning. arXiv preprint arXiv:1712.05526 (2017) 4
12. Chen, Y.C., Li, L., Yu, L., El Kholi, A., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Universal image-text representation learning. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX (2020) 4
13. Cheng, M., Sun, Y., Wang, L., Zhu, X., Yao, K., Chen, J., Song, G., Han, J., Liu, J., Ding, E., et al.: Vista: vision and scene text aggregation for cross-modal retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022) 4
14. Cho, K., van Merriënboer, B., Bahdanau, D., Bengio, Y.: On the properties of neural machine translation: Encoder–decoder approaches. In: 8th Workshop on Syntax, Semantics and Structure in Statistical Translation, SSST 2014 (2014) 6

15. Chun, S., Kim, W., Park, S., Chang, M., Oh, S.J.: Eccv caption: Correcting false negatives by collecting machine-and-human-verified image-caption associations for ms-coco. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VIII (2022) [4](#)
16. Chun, S., Oh, S.J., de Rezende, R.S., Kalantidis, Y., Larlus, D.: Probabilistic embeddings for cross-modal retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2021) [2](#), [4](#)
17. Croce, F., Hein, M.: Sparse and imperceivable adversarial attacks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2019) [4](#)
18. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers) (2019) [6](#)
19. Diao, H., Zhang, Y., Ma, L., Lu, H.: Similarity reasoning and filtration for image-text matching. In: Proceedings of the AAAI conference on artificial intelligence (2021) [4](#), [8](#), [9](#), [10](#), [12](#), [22](#)
20. Dong, X., Luu, A.T., Ji, R., Liu, H.: Towards robustness against natural language word substitutions. In: International Conference on Learning Representations (2021) [4](#)
21. Ebrahimi, J., Rao, A., Lowd, D., Dou, D.: Hotflip: White-box adversarial examples for text classification. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers) (2018) [4](#)
22. Frome, A., Corrado, G.S., Shlens, J., Bengio, S., Dean, J., Ranzato, M., Mikolov, T.: Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems* (2013) [4](#)
23. Gal, R., Patashnik, O., Maron, H., Bermano, A.H., Chechik, G., Cohen-Or, D.: Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)* (2022) [1](#), [2](#)
24. Galatolo, F.A., Cimino, M.G., Vaglini, G.: Generating images from caption and vice versa via clip-guided generative latent space search. *arXiv preprint arXiv:2102.01645* (2021) [1](#)
25. Garg, S., Ramakrishnan, G.: Bae: Bert-based adversarial examples for text classification. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP) (2020) [4](#)
26. Gokhale, T., Banerjee, P., Baral, C., Yang, Y.: Vqa-lol: Visual question answering under the lens of logic. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16 (2020) [4](#)
27. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014) [4](#)
28. Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., Parikh, D.: 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 (2017) [4](#)
29. Gu, T., Dolan-Gavitt, B., Garg, S.: Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733* (2017) [4](#)
30. Guo, C., Gardner, J., You, Y., Wilson, A.G., Weinberger, K.: Simple black-box adversarial attacks. In: Proceedings of the 36th International Conference on Machine Learning (2019) [4](#)
31. Gupta, T., Marten, R., Kembhavi, A., Hoiem, D.: Grit: general robust image task benchmark. *arXiv preprint arXiv:2204.13653* (2022) [6](#)

32. Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al.: The many faces of robustness: A critical analysis of out-of-distribution generalization. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021) 4
33. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. In: *International Conference on Learning Representations* (2019) 4
34. Hendrycks, D., Mazeika, M., Wilson, D., Gimpel, K.: Using trusted data to train deep networks on labels corrupted by severe noise. In: *Advances in Neural Information Processing Systems* (2018) 4
35. Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., Song, D.: Natural adversarial examples. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021) 4
36. Hu, P., Peng, X., Zhu, H., Zhen, L., Lin, J.: Learning cross-modal retrieval with noisy labels. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2021) 4
37. Iyyer, M., Wieting, J., Gimpel, K., Zettlemoyer, L.: Adversarial example generation with syntactically controlled paraphrase networks. In: *Proceedings of NAACL-HLT* (2018) 4
38. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: *International Conference on Machine Learning* (2021) 1, 4
39. Jia, R., Raghunathan, A., Göksel, K., Liang, P.: Certified robustness to adversarial word substitutions. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (2019) 4
40. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for generating image descriptions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition* (2015) 8
41. Kim, D., Kim, N., Kwak, S.: Improving cross-modal retrieval with set of diverse embeddings. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2023) 2
42. Kim, W., Son, B., Kim, I.: Vilt: Vision-and-language transformer without convolution or region supervision. In: *Proceedings of the 38th International Conference on Machine Learning*. PMLR (2021) 4
43. Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236* (2016) 4
44. Lee, K.H., Chen, X., Hua, G., Hu, H., He, X.: Stacked cross attention for image-text matching. In: *Proceedings of the European conference on computer vision (ECCV)* (2018) 4, 8
45. Li, D., Zhang, Y., Peng, H., Chen, L., Brockett, C., Sun, M.T., Dolan, W.B.: Contextualized perturbation for textual adversarial attack. In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (2021) 4
46. Li, D., Li, J., Hoi, S.: BLIP-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. In: *Thirty-seventh Conference on Neural Information Processing Systems* (2023) 1
47. Li, J., Ji, R., Liu, H., Hong, X., Gao, Y., Tian, Q.: Universal perturbation attack against image retrieval. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4899–4908 (2019) 4

48. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597* (2023) [1](#)
49. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: *International Conference on Machine Learning*. PMLR (2022) [2](#), [3](#), [4](#), [5](#), [6](#), [8](#), [9](#), [10](#), [11](#), [12](#), [21](#), [22](#)
50. Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems* (2021) [4](#), [8](#), [9](#), [10](#), [12](#), [22](#)
51. Li, K., Zhang, Y., Li, K., Li, Y., Fu, Y.: Visual semantic reasoning for image-text matching. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2019) [2](#), [6](#), [8](#), [9](#), [10](#), [11](#), [12](#), [22](#)
52. Li, L., Gan, Z., Liu, J.: A closer look at the robustness of vision-and-language pre-trained models. *arXiv preprint arXiv:2012.08673* (2020) [4](#)
53. Li, L., Lei, J., Gan, Z., Liu, J.: Adversarial vqa: A new benchmark for evaluating the robustness of vqa models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2021) [4](#)
54. Li, L., Ma, R., Guo, Q., Xue, X., Qiu, X.: Bert-attack: Adversarial attack against bert using bert. In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2020) [4](#)
55. Li, X., Yin, X., Li, C., Zhang, P., Hu, X., Zhang, L., Wang, L., Hu, H., Dong, L., Wei, F., et al.: Oscar: Object-semantics aligned pre-training for vision-language tasks. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX 16* (2020) [4](#)
56. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13* (2014) [2](#)
57. Lu, H., Fei, N., Huo, Y., Gao, Y., Lu, Z., Wen, J.R.: Cots: Collaborative two-stream vision-language pre-training model for cross-modal retrieval. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15692–15701 (2022) [4](#)
58. Parekh, Z., Baldrige, J., Cer, D., Waters, A., Yang, Y.: Crisscrossed captions: Extended intramodal and intermodal semantic similarity judgments for ms-coco. *arXiv preprint arXiv:2004.15020* (2020) [4](#)
59. Park, D.H., Darrell, T., Rohrbach, A.: Robust change captioning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019) [4](#)
60. Patashnik, O., Wu, Z., Shechtman, E., Cohen-Or, D., Lischinski, D.: Styleclip: Text-driven manipulation of stylegan imagery. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021) [1](#)
61. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In: *Proceedings of the IEEE international conference on computer vision* (2015) [2](#)
62. Prabhu, V., Yenamandra, S., Chattopadhyay, P., Hoffman, J.: Lance: Stress-testing visual models by generating language-guided counterfactual images. *arXiv preprint arXiv:2305.19164* (2023) [4](#)
63. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from

- natural language supervision. In: International conference on machine learning. PMLR (2021) [1](#), [4](#), [5](#), [6](#), [8](#), [9](#), [10](#), [11](#), [12](#), [22](#)
64. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 (2022) [1](#)
 65. Ren, S., Deng, Y., He, K., Che, W.: Generating natural language adversarial examples through probability weighted word saliency. In: Proceedings of the 57th annual meeting of the association for computational linguistics (2019) [4](#)
 66. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2022) [14](#)
 67. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M.: Photorealistic text-to-image diffusion models with deep language understanding. In: Advances in Neural Information Processing Systems. Curran Associates, Inc. (2022) [1](#), [2](#)
 68. Shah, M., Chen, X., Rohrbach, M., Parikh, D.: Cycle-consistency for robust visual question answering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2019) [4](#)
 69. Sheng, S., Singh, A., Goswami, V., Magana, J., Thrush, T., Galuba, W., Parikh, D., Kiela, D.: Human-adversarial visual question answering. Advances in Neural Information Processing Systems (2021) [4](#)
 70. Shridhar, M., Manuelli, L., Fox, D.: Cliport: What and where pathways for robotic manipulation. In: Conference on Robot Learning. PMLR (2022) [1](#), [2](#)
 71. Song, Y., Soleymani, M.: Polysemous visual-semantic embedding for cross-modal retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [4](#)
 72. Steinhardt, J., Koh, P.W.W., Liang, P.S.: Certified defenses for data poisoning attacks. Advances in neural information processing systems (2017) [4](#)
 73. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems (2017) [6](#)
 74. Wallace, E., Zhao, T., Feng, S., Singh, S.: Concealed data poisoning attacks on nlp models. In: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (2021) [4](#)
 75. Wang, H., He, D., Wu, W., Xia, B., Yang, M., Li, F., Yu, Y., Ji, Z., Ding, E., Wang, J.: Coder: Coupled diversity-sensitive momentum contrastive learning for image-text retrieval. In: European Conference on Computer Vision (2022) [2](#)
 76. Wang, H., Zhang, Y., Ji, Z., Pang, Y., Ma, L.: Consensus-aware visual-semantic embedding for image-text matching. In: Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV 16 (2020) [4](#)
 77. Wang, Y., Zhang, T., Zhang, X., Cui, Z., Huang, Y., Shen, P., Li, S., Yang, J.: Wasserstein coupled graph learning for cross-modal retrieval. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021) [4](#)
 78. Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., Wu, Y.: Coca: Contrastive captioners are image-text foundation models. arXiv preprint arXiv:2205.01917 (2022) [4](#)

- 79. Yuksekgonul, M., Bianchi, F., Kalluri, P., Jurafsky, D., Zou, J.: When and why vision-language models behave like bags-of-words, and what to do about it? In: The Eleventh International Conference on Learning Representations (2023) 4, 12
- 80. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: International Conference on Computer Vision (ICCV) (2019) 7, 13
- 81. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: International Conference on Learning Representations (2018) 7, 13
- 82. Zhang, P., Goyal, Y., Summers-Stay, D., Batra, D., Parikh, D.: Yin and yang: Balancing and answering binary visual questions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) 4
- 83. Zhang, P., Li, X., Hu, X., Yang, J., Zhang, L., Wang, L., Choi, Y., Gao, J.: Vinvl: Revisiting visual representations in vision-language models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 2, 4

Appendix

A Statistics of Source Words Selected by EI scores.

Figure 11 represents the statistics of source words selected based on the lowest EI scores across four models. Figure 11 (a) displays the level of agreement among models in selecting the word with the lowest EI scores. The x-axis represents the maximum number of agreements among the four models in selecting the word with the lowest EI score, while the y-axis represents the number of captions. Interestingly, in over 70% of cases, two or more models select the same word, despite being trained using different architectures and datasets (e.g., more pre-training data). Furthermore, Figure 11 (b) illustrates the distribution of the selected source word, which exhibits a long-tailed distribution. This highlights a common vulnerability in the current image-text matching approach, suggesting that attacks can have a universal impact.




Fig. 11: Statistics of source words selected by EI scores. (a) Consensus among models. x -axis indicates the number of consensus models that pick the same word, y -axis the frequency of consensus for each case in x -axis over all captions. (b) The source word distribution exhibits a long-tailed distribution.

B Text-to-Image Retrieval

We report the results on new image set with $\lambda = 0.7, 0.6$ in Table 7. We can still observe meaningful performance drop, when the added images are more significantly perturbed that seemed less confusing. In most cases, Incorrect Recall@1 exceeded 10%. In BLIP [49], performance degradation occurred more in fine-tuned models than in zero-shot models. We conjecture that this is because the models overfitted to COCO dataset during finetuning. We display the examples of retrieving incorrect images with BLIP ViT-B when $\lambda = 0.6, 0.7$ in Figure 12.

Table 7: Text-to-Image retrieval. Models are evaluated with our new benchmark: Mix and Patch with different λ . Recall@1 (R@1)(\uparrow), drop rate(\downarrow), Incorrect Recall@1 (IR@1)(\downarrow) are shown. The results are averaged over image generations with three different random seeds. We can see consistent degradation across all vision-language models.

	COCO 5K	Mix ($\lambda = 0.7$)			Mix ($\lambda = 0.6$)			Patch ($\lambda = 0.7$)			Patch ($\lambda = 0.6$)		
	R@1	R@1	drop rate	IR@1	R@1	drop rate	IR@1	R@1	drop rate	IR@1	R@1	drop rate	IR@1
Large-scale VL pre-training models													
CLIP ViT-B/32 (zero-shot) [63]	30.14	25.24	-16.25%	19.16	26.87	-10.84%	14.34	25.18	-16.45%	20.27	25.96	-13.86%	17.97
CLIP ViT-B/16 (zero-shot) [63]	33.03	26.60	-19.46%	22.48	28.67	-13.19%	16.90	26.14	-20.85%	25.64	27.12	-17.88%	22.76
CLIP ViT-L/14 (zero-shot) [63]	36.14	30.45	-15.75%	18.59	32.01	-11.43%	15.17	30.33	-16.08%	21.16	30.96	-14.34%	19.30
ALBEF [50]	60.67	52.53	-13.42%	14.44	55.91	-7.85%	9.19	53.71	-11.47%	12.54	54.75	-9.76%	10.92
BLIP ViT-B (zero-shot) [49]	56.36	48.12	-14.62%	16.52	50.81	-9.85%	11.70	46.97	-16.66%	18.74	48.38	-14.16%	16.33
BLIP ViT-B [49]	64.31	53.56	-16.72%	20.15	57.77	-10.18%	13.45	55.20	-14.17%	16.79	56.68	-11.87%	14.82
BLIP ViT-L (zero-shot) [49]	58.18	51.21	-11.98%	13.82	53.69	-7.71%	9.38	51.05	-12.25%	13.96	52.28	-10.14%	11.95
BLIP ViT-L [49]	65.06	52.06	-19.98%	24.43	57.08	-12.27%	16.19	55.58	-14.57%	18.05	57.19	-12.10%	15.41
Visual Semantic Embedding models													
VSRN [51]	40.34	34.80	-13.72%	19.24	37.04	-8.17%	12.69	34.01	-15.68%	21.31	34.99	-13.25%	18.59
SAF [19]	40.11	35.55	-11.37%	16.57	36.91	-7.98%	12.32	35.22	-12.20%	16.81	35.75	-10.87%	15.24
SGR [19]	40.45	35.59	-12.01%	16.54	37.23	-7.96%	11.96	35.23	-12.90%	17.12	35.85	-11.37%	15.53
VSE (BUTD region) [10]	42.46	38.74	-8.76%	11.99	40.38	-4.90%	7.20	38.18	-10.08%	13.57	38.99	-8.17%	11.42
VSE (BUTD grid) [10]	44.07	39.01	-11.47%	15.39	41.22	-6.46%	9.26	40.10	-9.00%	12.12	40.94	-7.09%	9.94
VSE (WSL grid) [10]	51.55	45.13	-12.46%	15.70	47.97	-6.95%	9.30	48.37	-6.17%	8.02	48.93	-5.08%	6.58

C Substituting more words

The results in Table 6 show meaningful performance degradation across the entire model, even when the original semantic meaning is significantly disrupted. When more words are substituted to create a stranger sentence (e.g., 5 words substitution), a large-scale vision-language pre-training models demonstrate relative robustness compared to the visual semantic embedding models.

Figure 13 shows examples of newly added captions that BLIP ViT-B model has retrieved as top 1. While the created captions are not natural, they include some keywords. Thus, we can conclude that the model is focusing on some nouns rather than the whole sentence.



Fig. 12: Text-to-Image retrieval examples.

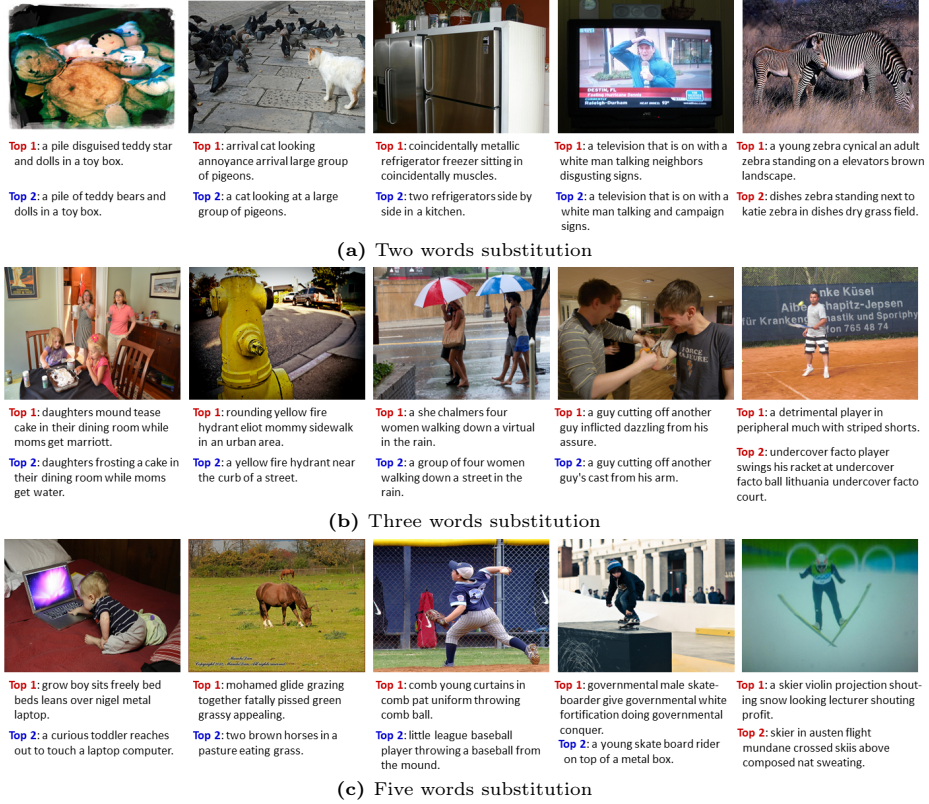


Fig. 13: Examples of substituting multiple random words with BLIP (ViT-B).