

Enhancing Adversarial Transferability in Visual-Language Pre-training Models via Local Shuffle and Sample-based Attack

Xin Liu*, Aoyang Zhou*, Kun He†

School of Computer Science and Technology

Huazhong University of Science and Technology, Wuhan, China

liuxin_jhl@hust.edu.cn, aoyangzhou@hust.edu.cn, brooklet60@hust.edu.cn

Abstract

Visual-Language Pre-training (VLP) models have achieved significant performance across various downstream tasks. However, they remain vulnerable to adversarial examples. While prior efforts focus on improving the adversarial transferability of multimodal adversarial examples through cross-modal interactions, these approaches suffer from overfitting issues, due to a lack of input diversity by relying excessively on information from adversarial examples in one modality when crafting attacks in another. To address this issue, we draw inspiration from strategies in some adversarial training methods and propose a novel attack called Local Shuffle and Sample-based Attack (LSSA). LSSA randomly shuffles one of the local image blocks, thus expanding the original image-text pairs, generating adversarial images, and sampling around them. Then, it utilizes both the original and sampled images to generate the adversarial texts. Extensive experiments on multiple models and datasets demonstrate that LSSA significantly enhances the transferability of multimodal adversarial examples across diverse VLP models and downstream tasks. Moreover, LSSA outperforms other advanced attacks on Large Vision-Language Models.

1 Introduction

Visual-Language Pre-training (VLP) models have achieved outstanding performance in various downstream visual-and-language tasks, including image-text retrieval (Chen et al., 2020a; Wang et al., 2019), visual grounding (Sadhu et al., 2019), and image captioning (Vinyals et al., 2015). Despite their success, recent works (Zhang et al., 2022; Lu et al., 2023) have revealed that VLP models remain vulnerable to multimodal adversarial examples, which add malicious perturbations to the original image-text pairs. Moreover, the multimodal adversarial

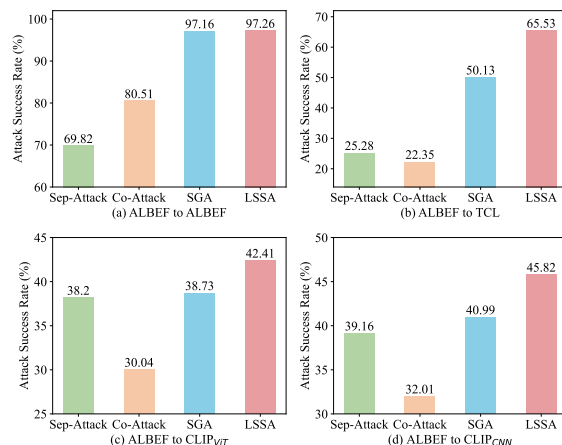


Figure 1: Comparison of attack success rate (%) using our LSSA method and existing advanced attacks in image-text retrieval tasks. The multimodal adversarial examples are crafted on the ALBEF model to attack ALBEF, TCL, CLIP_{VIT} and CLIP_{CNN}, respectively.

examples crafted for one visual-and-language task are often transferable to other downstream tasks. Therefore, it is crucial to explore multimodal adversarial attacks, as they provide valuable insights into the robustness of various VLP models (Li et al., 2021; Yang et al., 2022; Radford et al., 2021).

Recent research focuses on attacking VLP models in white-box setting (Zhang et al., 2022), where the attacker has access to the target model’s architecture and weights. However, it is more practical to explore multimodal adversarial attacks in black-box setting (Lu et al., 2023; Yu et al., 2023), where the inner information of target model is not accessible. As shown in Figure 1, while existing works have improved the white-box attack success rate on VLP models through modality interaction, their transferability remains unsatisfied, which is due to overfitting issue caused by limited input diversity, relying excessively on adversarial examples in one modality when crafting attacks in another.

Input transformations have been demonstrated to

*The first two authors contributed equally.

†Corresponding author.

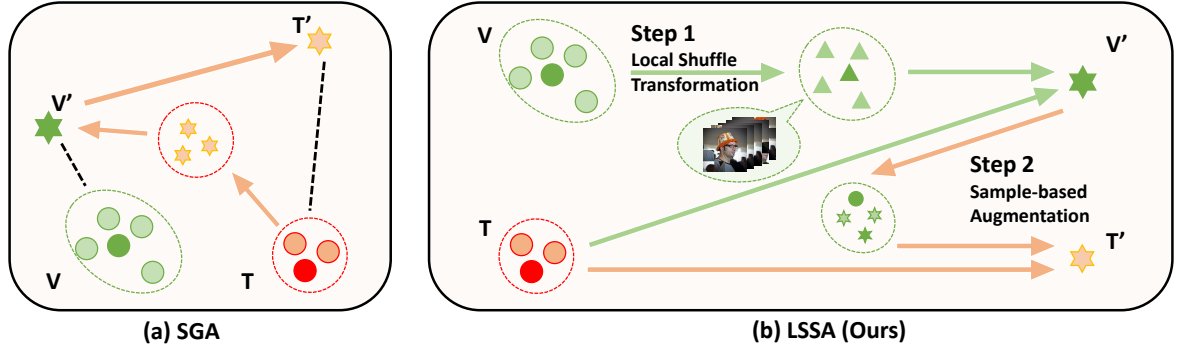


Figure 2: Comparison of SGA and LSSA. (a) and (b) show the multimodal adversarial examples generation process of SGA and LSSA, respectively. V' and T' represent the corresponding image and text adversarial examples. Dashed lines enclose the expanded set. Arrows indicate the process or guidance for generating adversarial examples.

boost adversarial examples’ transferability by enhancing input diversity in unimodal attacks. However, recent research (Lu et al., 2023) reveals that existing unimodal input transformation attacks not only fail to enhance adversarial transferability significantly but also considerably degrade white-box attack performance in multimodal case. It is because multimodal retrieval tasks rely more on the spatial information in images and feature alignment (Cao et al., 2022), while unimodal classification tasks depend more on global features. The data augmentation techniques used in adversarial training also contribute to enhancing input diversity. However, directly applying unimodal data augmentation to a multimodal attack harms the white-box attack performance due to the fundamental differences between multimodal and unimodal tasks. On the other hand, the generalization of models can be improved with diverse training data, and the transferability of adversarial examples can be enhanced by leveraging more data information. Nonetheless, existing attacks typically consider the adversarial image in generating adversarial text, ignoring the valuable information in the original image.

To address these issues, we draw inspiration from strategies in adversarial training methods (Song et al., 2020; Li et al., 2023b) and propose a novel attack called Local Shuffle and Sample-based Attack (LSSA), which improves input diversity while preserving spatial information. As in Figure 2, LSSA randomly shuffles one of the local image blocks and utilizes the shuffled images and original texts to generate adversarial images. Then, we sample the neighborhoods around the generated adversarial images and craft adversarial texts using sampled images, original images, and texts.

We conduct experiments on Flickr30K (Plum-

mer et al., 2017) and MSCOCO (Lin et al., 2014) to evaluate our LSSA across various downstream tasks. Experimental results demonstrate that LSSA boosts the attack performance in multimodal learning, outperforming the advanced multimodal attacks in black-box setting. Additionally, LSSA also surpasses advanced attacks in image captioning and visual grounding. Especially, to our knowledge, this is the first work to evaluate existing multimodal adversarial transferability performance on Large Vision-Language Models (LVLMs). Our contributions are summarized as follows:

- We observe that multimodal adversarial attacks rely on input diversity. For images, transformations should maintain spatial information to preserve white-box performance. For texts, using more data can disrupt feature alignment and enhance attack effectiveness.
- Based on the observations, we propose a novel attack called Local Shuffle and Sample-based Attack (LSSA), which randomly shuffles one of the local image blocks and samples around adversarial examples to generate adversarial images and texts.
- Extensive experiments demonstrate the effectiveness of LSSA, showing a significant improvement in the transferability of adversarial examples across various VLP models and downstream tasks. Moreover, LSSA outperforms advanced attacks on various LVLMs.

2 Related Work

2.1 Vision-Language Pre-training Models

Vision-Language Pre-training (VLP) models boost the performance of various Vision-and-Language

tasks (Li et al., 2022) through pre-training on extensive image-text pairs. Previous works primarily depend on the pre-trained object detectors to obtain the vision-language representations (Chen et al., 2020b; Li et al., 2020b; Wang et al., 2022; Zhang et al., 2021). Recently, with the great success of Vision Transformer (ViT) (Dosovitskiy et al., 2021; Touvron et al., 2021; Yuan et al., 2021) in various tasks, recent works have introduced ViT as an image encoder to convert the input into patches, substituting original computationally expensive object detectors (Dou et al., 2022; Li et al., 2022, 2021; Wang et al., 2023; Yang et al., 2022).

According to the architectures, VLP models can be categorized into fused and aligned models. Fused VLP models (*e.g.*, ALBEF (Li et al., 2021), TCL (Yang et al., 2022)) initially utilize separate unimodal encoders to obtain visual and text features. Then, it utilizes a multimodal encoder to output multimodal embeddings by further fusing the embeddings of images and texts. In contrast, the aligned VLP model (*e.g.*, CLIP (Radford et al., 2021)) extracts visual and text features by separate unimodal encoders and directly aligns their embeddings. Furthermore, the LVLMs (Li et al., 2023a; Zhu et al., 2023) also show remarkable performance in numerous multimodal tasks.

2.2 Downstream Vision-and-Language Tasks

Given input from one modality, Image-Text Retrieval (ITR) (Chen et al., 2020a; Wang et al., 2019) is a retrieval task where the goal is to retrieve the most relevant instances from a gallery database in the other modality.

For fused VLP models (Li et al., 2021; Yang et al., 2022; Li et al., 2022), the similarity scores are calculated for all image-text pairs to retrieve top-k candidates, which are then processed by a multimodal encoder to compute final image-text matching scores for ranking. For aligned VLP model (Radford et al., 2021), the final ranking can be directly determined by the image and text embedding similarity.

Image Captioning (IC) (Vinyals et al., 2015) generates suitable and logical textual descriptions for the visual images. The evaluation metrics for image captioning models, including BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE (Lin, 2004), CIDEr (Vedantam et al., 2015), and SPICE (Anderson et al., 2016), compare the generated text and the reference text in terms of quality and relevance. Visual Ground-

ing (VG) (Sadhu et al., 2019) is a localization task that identifies the corresponding visual region in images based on the textual descriptions.

2.3 Transferability of Adversarial Examples

Existing adversarial attacks can be categorized into white-box and black-box attacks. In white-box setting, the attacker has full access to the information of the target model, while such access is not available in black-box settings. In computer vision, numerous attacks have been proposed to craft adversarial examples by utilizing gradient information or input transformations, such as FGSM (Goodfellow et al., 2015), MIM (Dong et al., 2018), and TIM (Dong et al., 2019). In natural language processing, existing attacks, such as BERT-Attack (Li et al., 2020a) and FGPM (Wang et al., 2021), modify characters and words in the inputs.

In the multimodal vision-language field, Co-attack (Zhang et al., 2022) is the first and typical white-box attack for popular VLP models on downstream tasks. SGA (Lu et al., 2023) diversifies image-text pairs and extends the embedding distance of image-text pairs to improve the transferability of multimodal adversarial examples in black-box setting. However, SGA suffers from the overfitting issue due to a lack of input diversity by relying excessively on information from adversarial examples in one modality when crafting attacks in another. To further improve the adversarial transferability, we enrich the diversity of adversarial examples for image-text pairs and consider both original and sampled image-text pairs, enabling the transferability of adversarial examples.

3 Methodology

This section first introduces our motivations and observations. Then, we propose a novel multimodal adversarial attack called Local Shuffle and Sample-based Attack, providing a detailed description.

3.1 Motivations and Observations

Existing unimodal adversarial attacks, such as DIM (Xie et al., 2019), SIM (Lin et al., 2020), and multimodal attack SGA (Lu et al., 2023) use input transformation to effectively increase the diversity of inputs, thereby improving the performance of the attacks. However, SGA shows that existing unimodal adversarial attacks based on input transformation significantly harm the white-box attack performance. On the other hand, we observe that

Table 1: The attack success rate (%) of multimodal adversarial examples against different VLP models. The source column represents the VLP models used for crafting multimodal adversarial examples on the Flickr30K dataset. * indicates white-box attacks. The higher attack success rate indicates the better performance.

Source	Attack	ALBEF		TCL		CLIP _{ViT}		CLIP _{CNN}	
		TR R@1	IR R@1	TR R@1	IR R@1	TR R@1	IR R@1	TR R@1	IR R@1
ALBEF	SGA	97.24*	97.08*	44.68	55.57	33.25	44.20	35.38	46.59
	Global Shuffle	91.97*	91.93*	48.26	58.02	34.87	46.79	38.85	48.95
	Local Shuffle	96.45*	96.40*	56.16	64.38	35.95	48.16	40.38	50.23

data augmentation in adversarial training can also effectively increase input diversity. Therefore, inspired by the adversarial training method of RLFAT (Song et al., 2020), we conduct experiments to explore the impact of data augmentation on the transferability of adversarial examples.

Specifically, we combine the Global Shuffle of RLFAT with SGA. Following their settings, we also divide the original images into blocks and randomly shuffle them to generate adversarial examples using the gradients of shuffled images. Detailed information is described in Appendix A.1 and Appendix A.2. As shown in Table 1, Global Shuffle improves the adversarial transferability, yet the performance of white-box attacks significantly decays. It is because multimodal retrieval tasks rely more on the spatial information in images and feature alignment while Global Shuffle excessively disrupts them in the original images. Therefore, a suitable input transformation needs to increase diversity while preserving the original spatial information. To address this issue, we block the original image and randomly shuffle one of the local image blocks. As shown in Table 1, Local Shuffle transformation significantly enhances the black-box attack while maintaining white-box performance, boosting the diversity of original image-text pairs (v, t) .

In the text modality, previous attacks only consider either the original image or the adversarial image in generating adversarial text, without considering both information. Similar to improving the generalization of models with more training data, the transferability of adversarial examples can be enhanced by utilizing more data information (Lin et al., 2020). Therefore, it is natural to consider the information from multiple samples to improve the adversarial transferability. Inspired by another adversarial training method of STAT (Li et al., 2023b), which utilizes both the original and the neighbors of the adversarial examples to enhance the model robustness, we conduct experiments to investigate

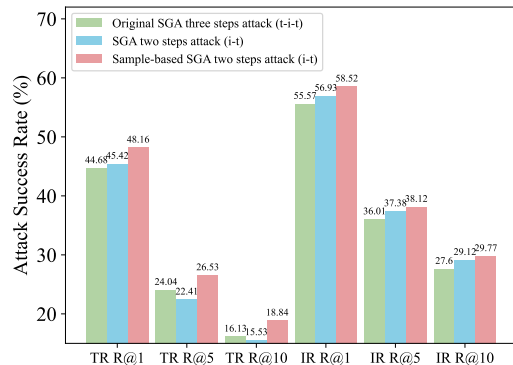


Figure 3: The attack success rate of multimodal adversarial examples against TCL model, which are crafted on ALBEF model. We explore three methods: **a) Original SGA three steps attack (t-i-t)**, **b) SGA two steps attack (i-t)**: generating adversarial images using original text and then generating adversarial text using the adversarial images, **c) Sample-based SGA two steps attack (i-t)**: generating adversarial images using original text and then using the original and neighbours of the adversarial image to generate the adversarial images.

the multimodal adversarial examples generated by different images and text information. As illustrated in Figure 3, the experimental results indicate that directly using the original text outperforms using adversarial text for generating transferable adversarial images. Additionally, combining the original image with the sampled points around the adversarial image shows better transferability compared to individually using the adversarial image.

3.2 Local Shuffle and Sample-based Attack

Local Shuffle Transformation. Based on the analysis in Section 3.1, in each iteration, we randomly shuffle one of the local image blocks of the current adversarial images (v_i^{adv}, t) , and repeat to get N samples. By utilizing the above input transformation method, we can obtain an expanded dataset $(V, T) = \{(v'_1, t), (v'_2, t), \dots, (v'_N, t)\}$. Then, the expanded dataset (V, T) is used to generate the ad-

versarial images (v_{i+1}^{adv}, t), which can be formulated as follows:

$$g_{i+1} = \mu g_i + \frac{1}{N} \sum_{j=1}^N \frac{\nabla_v J(f_I(v'_j), f_T(t))}{\|\nabla_v J(f_I(v'_j), f_T(t))\|}, \quad (1)$$

$$v_{i+1}^{adv} = \text{Clip}_{\epsilon_v}(v_i^{adv} + \alpha \cdot \text{sign}(g_{i+1})),$$

where $J(\cdot, \cdot)$ is the loss function, g_i is the accumulated gradient, ϵ_v is the perturbation boundary of the image, α is the step size and μ is the decay factor. $f_I(\cdot)$ and $f_T(\cdot)$ are visual and text encoder.

Sample-based Augmentation. Based on the analysis in Section 3.1, it is evident that the original and neighbors of adversarial images can effectively enhance the transferability of multimodal adversarial examples. It ensures the adversarial text has significant differences not only from the original image features but also from the adversarial image features. The adversarial text generation process can be formalized as follows:

$$t^{adv} = \arg \max_{t' \in B[t, \epsilon_t]} \lambda \cdot J(f_I(v), f_T(t')) + (1 - \lambda) \cdot \frac{1}{M} \sum_{i=1}^M J(f_I(v_i), f_T(t')), \quad (2)$$

where $J(\cdot, \cdot)$ is the loss function, M is the sampled number, v is the original image, and v_i is the neighbours of adversarial image. The overall algorithm of LSSA is summarized in Appendix A.3.

Difference of LSSA. Since our LSSA attack draws inspiration from the strategies of adversarial training methods RLFAT and STAT, we highlight the differences as follows:

a) The goal of adversarial training is to boost model generalization, while our LSSA aims to generate more transferable adversarial examples with input diversity. **b)** The local shuffle in our LSSA is designed for multimodal retrieval tasks, whereas the global shuffle in RLFAT is proposed for unimodal classification tasks. **c)** The sample-based attack enhances performance by randomly sampling around the adversarial images, whereas STAT simultaneously crafts different adversarial images to improve model robustness.

With different goals, tasks, and methods, our LSSA is a new and novel input transformation based multimodal adversarial attack to improve the transferability of attack performance.

4 Experiments

In this section, we conduct experiments on two benchmark datasets using popular VLP models and

LVLMS. First, we describe the experimental settings in Section 4.1. Next, we evaluate the performance of various advanced multimodal adversarial attacks in Section 4.2. We then analyze the cross-task transferability between different V+L tasks in Section 4.3 and present the results of further study in Section 4.4. Finally, we evaluate the LSSA method on various LVLMS to verify its generalization in Section 4.5.

4.1 Experimental Settings

Datasets and VLP Models. Following the setting on SGA (Lu et al., 2023), we conduct experiments on two benchmark datasets, namely Flickr30K (Plummer et al., 2017) and MSCOCO (Lin et al., 2014). Flickr30K and MSCOCO dataset contains 31,783 and 123,287 images, respectively, with five corresponding captions. We evaluate two typical architectures of VLP models: the fused VLP models and aligned models. For the fused VLP models, we choose ALBEF (Li et al., 2021) and TCL (Yang et al., 2022), which are pre-trained by different pre-trained tasks. For the aligned VLP models, we choose the CLIP model with different visual encoders. More details are provided in Appendix B.1.

Adversarial Attack Settings. We follow the multimodal adversarial attack setting of SGA. Moreover, we extend the dataset by random shuffling one of the local image blocks, where the size of the original images is $h \times w$, and the size of the local image block is $h/2 \times w/2$, which is blocked to four $h/4 \times w/4$ subblocks. We randomly shuffle $N = 20$ times for each image. For the sampled image number, we choose $M = 20$ to generate text adversarial examples. More details are provided in Appendix B.1.

4.2 Experiments Results

We conduct extensive experiments to evaluate our LSSA method on two widely adopted architectures of models: fused VLP and aligned VLP models.

As shown in Table 2, our proposed LSSA outperforms existing advanced multimodal attacks in white-box and black-box settings. When the type of source model is the same as the target model, LSSA has significantly improved adversarial transferability. Specifically, LSSA surpasses SGA by 14.75% and 10.05% in terms of attack success rate on TR and IR tasks, respectively, when transferring the multimodal adversarial examples from ALBEF to TCL. We observe the same phe-

Table 2: The attack success rate (%) of multimodal adversarial examples against different VLP models compared with state-of-the-art methods on image-text retrieval task. The source column represents the VLP models used for crafting multimodal adversarial examples on the Flickr30K dataset. * indicates white-box attacks.

Source	Attack	ALBEF		TCL		CLIP _{ViT}		CLIP _{CNN}	
		TR R@1	IR R@1	TR R@1	IR R@1	TR R@1	IR R@1	TR R@1	IR R@1
ALBEF	PGD	52.45*	58.65*	3.06	6.79	8.96	13.21	10.34	14.65
	BERT-Attack	11.57*	27.46*	12.64	28.07	29.33	43.17	32.69	46.11
	Sep-Attack	65.69*	73.95*	17.60	32.95	31.17	45.23	32.82	45.49
	Co-Attack	77.16*	83.86*	15.21	29.49	23.60	36.48	25.12	38.89
	SGA	97.24*	97.08*	44.68	55.57	33.25	44.20	35.38	46.59
	LSSA (Ours)	97.29*	97.22*	59.43	65.62	36.32	48.49	40.87	50.77
TCL	PGD	6.15	10.78	77.87*	79.48*	7.48	13.72	10.34	15.33
	BERT-Attack	11.89	26.82	14.54*	29.17*	29.69	44.49	33.46	46.07
	Sep-Attack	20.13	36.48	84.72*	86.07*	31.29	44.65	33.33	45.80
	Co-Attack	23.15	40.04	77.94*	85.59*	27.85	41.19	30.74	44.11
	SGA	49.32	59.92	98.52*	98.79*	33.87	45.07	38.44	47.92
	LSSA (Ours)	59.85	67.59	98.63*	98.80*	35.46	48.39	41.12	51.05
CLIP _{ViT}	PGD	2.50	4.93	4.85	8.17	70.92*	78.61*	5.36	8.44
	BERT-Attack	9.59	22.64	11.80	25.07	28.34*	39.08*	30.40	37.43
	Sep-Attack	9.59	53.25	11.38	25.60	79.75*	86.79*	30.78	39.76
	Co-Attack	10.57	24.33	11.94	26.69	93.25*	95.86*	32.52	41.82
	SGA	13.56	27.01	14.54	30.07	98.80*	98.94*	40.61	47.55
	LSSA (Ours)	16.89	32.11	17.18	33.07	98.85*	98.97*	48.28	55.44
CLIP _{CNN}	PGD	2.09	4.82	4.00	7.81	1.10	6.60	86.46*	92.25*
	BERT-Attack	8.86	23.27	12.33	25.48	27.12	37.44	30.40*	40.10*
	Sep-Attack	8.55	23.41	12.64	26.12	28.34	39.43	91.44*	95.44*
	Co-Attack	8.79	23.79	13.10	26.07	28.79	40.03	94.76*	96.89*
	SGA	10.74	25.02	14.54	27.26	30.80	41.82	99.23*	99.42*
	LSSA (Ours)	14.08	26.89	15.28	30.12	37.67	46.91	99.74*	99.97*

nomenon that the multimodal adversarial examples have better adversarial transferability from CLIP_{ViT} to CLIP_{CNN}, which outperforms SGA by 7.67% and 7.89% on TR and IR tasks, respectively. Moreover, we also evaluate the performance of LSSA on another benchmark dataset MSCOCO, which is more challenging. Our LSSA still exceeds the adversarial transferability of baseline methods with a clear margin. The details of experiments on the MSCOCO dataset are provided in Appendix B.3.

To explore the effectiveness of LSSA on the different architecture of VLP models, we craft multimodal adversarial examples from ALBEF (CLIP) to CLIP (ALBEF) model. As shown in Table 2, our proposed LSSA consistently improves the transferability of adversarial attacks on both ViT-based CLIP and ViT-based CLIP, generated on ALBEF model, surpassing the performance of other advanced attacks. Specifically, LSSA significantly enhances the adversarial transferability, attaining an improvement of 3.07%, 4.69% and 5.59%, 4.18% on CLIP_{ViT} and CLIP_{CNN}.

4.3 Cross-Task Transferability

We also conduct extensive experiments to evaluate the effectiveness of our LSSA attack on Image Captioning (IC) and Visual Grounding (VG).

In the IC task, we craft adversarial images in the ALBEF model and then attack the BLIP (Li et al., 2022) model on the MSCOCO dataset. As shown in Table 3, LSSA enhances the adversarial transferability compared with SGA, gaining improvements of 2.87%, 4.93%, 2.49%, 6.29% and 6.54%, respectively. In the VG task, we generate adversarial image-text pairs on the RefCOCO+ dataset. The source and target models are ALBEF models pre-trained on the ITR and VG tasks, respectively. As shown in Table 4, LSSA still performs better than existing attacks, significantly improving 8.58%, 12.42%, and 6.18%, respectively. It might be because the adversarial perturbations crafted by LSSA contain more spacial adversarial information compared with other modal interaction attacks. The results demonstrate the effectiveness of our LSSA on cross-task transferability.

Table 3: **Cross-Task Transerability: ITR \rightarrow IC.** The adversarial image-text pairs are generated from Image-Text Retrieval (ITR) task to attack Image Caption (IC) task on MSCOCO dataset. The source and target model are ALBEF and BLIP, respectively. The Baseline is the performance of IC on the original data, where the lower value indicates better cross-task transerability of adversarial examples.

Attack	B@4	METEOR	ROUGE_L	CIDEr	SPICE
Baseline	39.7	31.0	60.0	133.3	23.8
Co-Attack	37.4	29.8	58.4	125.5	22.8
SGA	34.8	28.4	56.3	116.0	21.4
LSSA (Ours)	33.8	27.0	54.9	108.7	20.0

Table 4: **Cross-Task Transerability: ITR \rightarrow VG.** The adversarial image-text pairs are generated from the Image-Text Retrieval (ITR) task to attack the Visual Grounding (VG) task in the ALBEF model on the ReFCOCO+ dataset. The Baseline is the performance of VG on the original data, where the lower value indicates better cross-task transferability of adversarial examples.

Attack	Val	TestA	TestB
Baseline	58.5	65.9	46.3
Co-Attack	54.3	61.8	43.8
SGA	53.6	61.2	43.7
LSSA (Ours)	49.0	53.6	41.0

4.4 Further Study

We conduct various experiments on image-text retrieval to explore the impact of hyper-parameters in the LSSA method, including the number of local shuffles, the positions of local shuffles, and the weight of the loss. Detailed information of the ablation study is provided in Appendix B.5. Specifically, we use ALBEF as the white-box model to craft multimodal adversarial examples, while other models are attacked as black-box models.

Number of local shuffles N . We conduct experiments to explore the effectiveness of local shuffle number N . Specifically, when the number of local shuffles $N = 0$, it degenerates into a sample-based attack. As shown in Figure 4, local shuffle can effectively enhance the transferability of multimodal adversarial examples, and the enhancement improves with the increase of the number of shuffled samples, reaching its maximum at approximately $N = 20$. Specifically, compared to no local shuffle, shuffling gains an improvement of 5.59%, 1.96%, 1.15% on IR task and 4.69%, 1.33%, 0.72% on TR task, crafted multimodal adversarial examples in TCL, CLIP_{VIT} and CLIP_{CNN}. To achieve optimal performance, we choose $N = 20$.

Table 5: The attack success rate (%) of LSSA for black-box attacks on LVLMs using the Flickr30K dataset, where the surrogate model is the ALBEF model.

Attack/Model	BLIP2	VisualGLM	MiniGPT4	PandaGPT
Baseline	0.84	0.68	0.92	0.86
Co-Attack	19.82	10.38	25.80	19.06
SGA	20.45	10.72	25.54	19.52
LSSA (Ours)	21.64	11.52	27.56	20.72

Positions of local shuffle. We conduct experiments to investigate the impact of the shuffle position on the diversity of image-text pairs and the transferability of multimodal adversarial examples. Specifically, we only shuffle the left top, right top, left bottom, and right bottom of the image. Moreover, we compare the performance of the above strategies with a random shuffle. As shown in Figure 5, shuffling different positions results in variations in the generated multimodal adversarial examples. Moreover, the impact of local shuffling differs across different VLP models, possibly due to variations in the model’s attention to different regions of the image-text pairs. To mitigate this influence and improve the attack performance, we choose the random shuffle strategy.

Weight of the loss λ . The weight of loss λ controls the trade-off between the loss of original image-pairs and sampled image-pairs. To investigate the impact of weight loss λ , we evaluate the effectiveness with different λ . Specifically, we let λ vary from 0 to 1. When $\lambda = 0$, the loss disregards the loss of the original image-text pairs, while when $\lambda = 1$, the loss disregards the information from the sampled image-text pairs. As shown in Figure 6, the performance of LSSA increases and reaches the peak when $\lambda = 0.5$, which indicates that excessive reliance on either the information from the original image-text pairs or the sampled image-text pairs can lead to a decrease in the attack success rate. Therefore, we select an intermediate value of $\lambda = 0.5$ to obtain the best performance.

4.5 Attack Performance on LVLMs

We first evaluate the attack performance of existing attacks and our LSSA on LVLMs, which have achieved remarkable success on various tasks. Specifically, we conduct attacks on BLIP-2 (Li et al., 2023a), VisualGLM (Du et al., 2022), MiniGPT4 (Zhu et al., 2023) and PandaGPT (Su et al., 2023) in black-box setting. Due to the difference between VLP and LVLm tasks, we utilize




Figure 4: The attack success rate (%) of LSSA with different shuffle number N on the Flickr30K dataset. The source model is ALBEF, and the target model is other VLP models.




Figure 5: The attack success rate (%) of LSSA with different shuffle position on the Flickr30K dataset. The source model is ALBEF, and the target model is other VLP models.

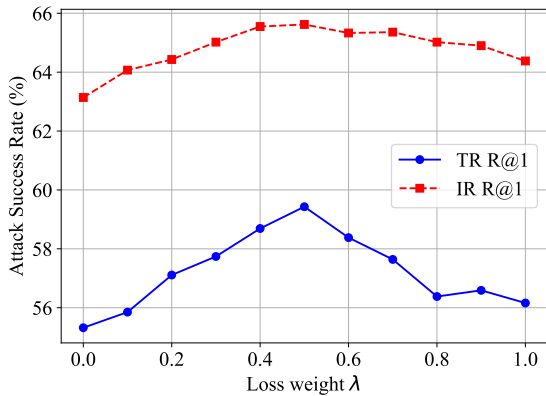


Figure 6: The attack success rate (%) of CSSA with different loss weights λ on the Flickr30K dataset. The source and target model are ALBEF and TCL model.

dialogues to compare the images and texts. For each image, we use a language template ("Question: Does the picture depict that *[text]*? Answer:") to ask LVLMs to obtain judgments. As shown in Table 5, LSSA still surpasses baselines, gaining an improvement of 1.19%, 0.80%, 2.02%, and 1.20%. It indicates the effectiveness and robustness

of LSSA compared with other advanced attacks.

5 Conclusion

In this work, we first investigate the adversarial transferability of multimodal attacks. We observe that existing attacks suffer from the overfitting issue due to a lack of input diversity by relying excessively on information from adversarial examples in one modality when crafting attacks in another. To address this issue, we propose a novel multimodal adversarial attack called Local Shuffle and Sample-based Attack (LSSA), which utilizes local shuffle to enrich image-text pairs and generate adversarial text by original and sampled images. Extensive experiments on popular benchmark datasets and VLP models demonstrate that LSSA significantly improves the transferability of multimodal adversarial examples across different V+L downstream tasks. Additionally, the multimodal adversarial examples crafted by LSSA are more threatening than other attacks on LVLMs. We hope this work could inspire more works that connects strategies for adversarial attack and defense to evaluate the adversarial robustness of VLP models.

Limitations

LSSA focuses on the two popular and typical vision-language pre-training model in the multi-modal task. However, it does not consider other types of modalities, such as video and signal. These data can also be crafted adversarial examples by applying our method. We will continue to explore the potential of LSSA in our further work.

Acknowledgments

This work is supported by National Natural Science Foundation (U22B2017) and International Cooperation Foundation of Hubei Province, China (2024EHA032).

References

- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. SPICE: semantic propositional image caption evaluation. In *ECCV*, pages 382–398.
- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: an automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization@ACL 2005, Ann Arbor, Michigan, USA, June 29, 2005*, pages 65–72.
- Min Cao, Shiping Li, Juntao Li, Liqiang Nie, and Min Zhang. 2022. Image-text retrieval: A survey on recent research and development. *arXiv preprint arXiv:2203.14713*.
- Hui Chen, Guiguang Ding, Xudong Liu, Zijia Lin, Ji Liu, and Jungong Han. 2020a. IMRAM: iterative matching with recurrent attention memory for cross-modal image-text retrieval. In *CVPR*, pages 12652–12660.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020b. UNITER: universal image-text representation learning. In *ECCV*, pages 104–120.
- Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. 2018. Boosting adversarial attacks with momentum. In *CVPR*, pages 9185–9193.
- Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. 2019. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *CVPR*, pages 4312–4321.
- 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. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*.
- Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Pengchuan Zhang, Lu Yuan, Nanyun Peng, Zicheng Liu, and Michael Zeng. 2022. An empirical study of training end-to-end vision-and-language transformers. In *CVPR*, pages 18145–18155.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. GLM: general language model pretraining with autoregressive blank infilling. In *ACL*, pages 320–335.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations, ICLR*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*, pages 770–778.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. 2023a. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, pages 19730–19742.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. 2022. BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, pages 12888–12900.
- Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. In *NeurIPS*, pages 9694–9705.
- Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020a. BERT-ATTACK: adversarial attack against BERT using BERT. In *EMNLP*, pages 6193–6202.
- Qizhang Li, Yiwen Guo, Wangmeng Zuo, and Hao Chen. 2023b. Squeeze training for adversarial robustness. In *ICLR*.
- Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. 2020b. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, pages 121–137.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Jiadong Lin, Chuanbiao Song, Kun He, Liwei Wang, and John E. Hopcroft. 2020. Nesterov accelerated gradient and scale invariance for adversarial attacks. In *ICLR*.

- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: common objects in context. In *ECCV*, pages 740–755.
- Dong Lu, Zhiqiang Wang, Teng Wang, Weili Guan, Hongchang Gao, and Feng Zheng. 2023. Set-level guidance attack: Boosting adversarial transferability of vision-language pre-training models. In *ICCV*, pages 102–111.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318.
- Bryan A. Plummer, Liwei Wang, Chris M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. 2017. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. *Int. J. Comput. Vis.*, 123(1):74–93.
- 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. 2021. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763.
- Arka Sadhu, Kan Chen, and Ram Nevatia. 2019. Zero-shot grounding of objects from natural language queries. In *ICCV*, pages 4693–4702.
- Chuanbiao Song, Kun He, Jiadong Lin, Liwei Wang, and John E. Hopcroft. 2020. Robust local features for improving the generalization of adversarial training. In *ICLR*.
- Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. 2023. Pandagpt: One model to instruction-follow them all. *CoRR*, abs/2305.16355.
- Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. 2021. Training data-efficient image transformers & distillation through attention. In *ICML*, pages 10347–10357.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *CVPR*, pages 4566–4575.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *CVPR*, pages 3156–3164.
- Teng Wang, Yixiao Ge, Feng Zheng, Ran Cheng, Ying Shan, Xiaohu Qie, and Ping Luo. 2023. Accelerating vision-language pretraining with free language modeling. In *CVPR*, pages 23161–23170.
- Teng Wang, Wenhao Jiang, Zhichao Lu, Feng Zheng, Ran Cheng, Chengguo Yin, and Ping Luo. 2022. Vlmixer: Unpaired vision-language pre-training via cross-modal cutmix. In *ICML*, pages 22680–22690.
- Xiaosen Wang, Yichen Yang, Yihe Deng, and Kun He. 2021. Adversarial training with fast gradient projection method against synonym substitution based text attacks. In *AAAI*, pages 13997–14005.
- Zihao Wang, Xihui Liu, Hongsheng Li, Lu Sheng, Junjie Yan, Xiaogang Wang, and Jing Shao. 2019. CAMP: cross-modal adaptive message passing for text-image retrieval. In *ICCV*, pages 5763–5772.
- Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L. Yuille. 2019. Improving transferability of adversarial examples with input diversity. In *CVPR*, pages 2730–2739.
- Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liqun Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. 2022. Vision-language pre-training with triple contrastive learning. In *CVPR*, pages 15650–15659.
- Zhen Yu, Zhou Qin, Zhenhua Chen, Meihui Lian, Haojun Fu, Weigao Wen, Hui Xue, and Kun He. 2023. Sparse black-box multimodal attack for vision-language adversary generation. In *EMNLP*, pages 5775–5784.
- Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zihang Jiang, Francis E. H. Tay, Jiashi Feng, and Shuicheng Yan. 2021. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *ICCV*, pages 538–547.
- Jiaming Zhang, Qi Yi, and Jitao Sang. 2022. Towards adversarial attack on vision-language pre-training models. In *ACM MM*, pages 5005–5013.
- Pengchuan Zhang, Xiujuan Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. Vinvl: Revisiting visual representations in vision-language models. In *CVPR*, pages 5579–5588.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigt-4: Enhancing vision-language understanding with advanced large language models. *CoRR*, abs/2304.10592.

Appendix

A More Details

A.1 Visualization of Shuffle Images

To better visualize the differences between images after a local shuffle and global shuffle, we present the corner shuffled images in LSSA and the global shuffled images in RLFAT to show the varying degrees of spatial information disruption. As shown in Figure 7, the global shuffled image has been totally disrupted while the corner shuffled image only loses part of the information.




Figure 7: Visualization of original images (left), corresponding global shuffle images (middle) and local shuffled images (right). It indicates that local shuffled images preserve more spacial information than global shuffle images while enhancing the diversity.

A.2 Visualization of Adversarial Examples

To better show the differences between the generated multimodal adversarial examples and the original samples, we present the image-text pairs and the corresponding adversarial examples generated by LSSA in Figure 8. Due to the small magnitude

of the adversarial perturbations, we have amplified the adversarial perturbations by 40 times.

A.3 Algorithm

The detailed attack process of our proposed LSSA is described in Algorithm 1.

Algorithm 1 The Local Shuffle and Sample-based Attack

Input: Image encoder f_I , Text encoder f_T , Dataset D , Image-caption pair (v, t) , iteration steps T , random shuffle number N , sampled number M , loss function J , decay factor λ , adversarial step α , image perturbation boundary ϵ_v , text perturbation boundary ϵ_t .

$v_1^{adv} \leftarrow v$

$g_1 \leftarrow 0$

// Generate adversarial image v^{adv}

for $i = 1$ to T do

$$g_{i+1} = \lambda \cdot g_i + \frac{1}{N} \sum_{j=1}^N \frac{\nabla_v J(f_I(v'_j), f_T(t))}{\|\nabla_v J(f_I(v'_j), f_T(t))\|}$$

$$v_{i+1}^{adv} \leftarrow \text{Clip}_{\epsilon_v}(v_i^{adv} + \alpha \cdot \text{sign}(g_{i+1}))$$

end for

$v^{adv} \leftarrow v_{T+1}^{adv}$

// Generate adversarial text t^{adv}

$$t^{adv} = \arg \max_{t' \in B[t, \epsilon_t]} \lambda \cdot J(f_I(v), f_T(t')) + (1 -$$

$$\lambda) \cdot \frac{1}{M} \sum_{i=1}^M J(f_I(v_i), f_T(t'))$$

Output: adversarial image v^{adv} , adversarial caption t^{adv}

B More Experiments Details

B.1 Experimental Setting

Datasets and VLP Models. We conduct experiments on two benchmark datasets, namely Flickr30K (Plummer et al., 2017) and MSCOCO (Lin et al., 2014). Flickr30K and MSCOCO dataset contains 31,783 and 123,287 images, respectively, with each image paired with five corresponding captions. We evaluate two typical architectures of VLP models: the fused VLP models and aligned models. For the fused VLP models, we choose ALBEF (Li et al., 2021) and TCL (Yang et al., 2022). The visual encoder of ALBEF and TCL is ViT-B/16 (Dosovitskiy et al., 2021), and the text encoder and multimodal encoder are two 6-layer transformers. However, they are pre-trained by different pre-trained tasks. For the aligned VLP models, we choose the CLIP model with different visual encoders. Specifically, we choose ViT-B/16 (Dosovitskiy et al., 2021)



Figure 8: Visualization of original and adversarial image-text pairs generated by LSSA. The rows are the original images, corresponding adversarial images and adversarial perturbations, original and corresponding adversarial texts. The perturbation magnitude ensures that the perturbations in the image are imperceptible, and only a few words in the text have changed.

and ResNet-101 (He et al., 2016) as the base visual encoder and BERT as the base text encoder, respectively.

Adversarial Attack Settings. We follow the multimodal adversarial attack setting of SGA (Lu et al., 2023). Specifically, we use MI-FGSM attack (Dong et al., 2018) with perturbation boundary $\epsilon_v = 2/255$, the step size $\alpha = 0.5/255$, and the iteration steps $T = 10$ to generate image adversarial examples. For text adversarial examples, we adopt BERT-Attack (Li et al., 2020a) with the perturbation bound $\epsilon_t = 1$ and the maximum candidate words $W = 10$. We enlarge the image dataset by resizing the original image as SGA.

B.2 Performance Metric

We utilize the Attack Success Rate (ASR) as the metric to evaluate the adversarial robustness and transferability of our models in both white-box and black-box settings. The ASR measures the percentage of attacks that successfully generate adversarial examples. A higher ASR indicates better transferability of adversarial attacks across differ-

ent models. By focusing on ASR, we can assess the effectiveness of our models against adversarial attacks and their ability to generalize beyond the training data, providing a comprehensive evaluation of their robustness and transferability.

B.3 More Experiments Results

We also conduct comprehensive evaluations of our LSSA and other baselines on the MSCOCO dataset. The experimental settings are consistent with those on the Flickr30K dataset. As shown in Table 6, our LSSA method performs well on two different architecture of models and shows significant improvement compared to the baselines. Specifically, LSSA further improves the attack success rate in the white-box setting. Moreover, in the black-box model setting, it boost the transferable attack performance of multimodal adversarial examples with a clear margin. Similar to previous observations, adversarial examples generated by models with similar architecture (such as ALBEF and TCL, CLIP_{VIT} and CLIP_{CNN}) perform better on other models.

Table 6: The attack success rate (%) of multimodal adversarial examples against different VLP models compared with state-of-the-art methods on image-text retrieval task. The source column represents the VLP models used for crafting multimodal adversarial examples on the MSCOCO dataset. * indicates white-box attacks. The higher attack success rate indicates the better performance.

Source	Attack	ALBEF		TCL		CLIP _{VIT}		CLIP _{CNN}	
		TR R@1	IR R@1	TR R@1	IR R@1	TR R@1	IR R@1	TR R@1	IR R@1
ALBEF	PGD	76.70*	86.30*	12.46	17.77	13.96	23.10	17.45	23.54
	BERT-Attack	24.39*	36.13*	24.34	33.39	44.94	52.28	47.73	54.75
	Sep-Attack	82.60*	89.88*	32.83	42.92	44.03	54.46	46.96	55.88
	Co-Attack	79.87*	87.83*	32.62	43.09	44.89	54.75	47.30	55.64
	SGA	96.75*	96.95*	58.56	65.38	57.06	65.25	58.95	66.52
	LSSA (Ours)	97.89*	97.58*	68.97	73.35	60.85	67.99	65.01	73.33
TCL	PGD	10.83	16.52	59.58*	69.53*	14.23	22.28	17.25	23.12
	BERT-Attack	35.32	45.92	38.54*	48.48*	51.09	58.80	52.23	61.26
	Sep-Attack	41.71	52.97	70.32*	78.97*	50.74	60.13	51.90	61.26
	Co-Attack	46.08	57.09	85.38*	91.39*	51.62	60.46	52.13	62.49
	SGA	65.93	73.30	98.97*	99.15*	56.34	63.99	59.44	65.70
	LSSA (Ours)	74.04	80.36	99.23*	99.33*	60.70	67.43	62.40	69.63
CLIP _{VIT}	PGD	7.24	10.75	10.19	13.74	54.79*	66.85*	7.32	11.34
	BERT-Attack	20.34	29.74	21.08	29.61	45.06*	51.68*	44.54	53.72
	Sep-Attack	23.41	34.61	25.77	36.84	68.52*	77.94*	43.11	49.76
	Co-Attack	30.28	42.67	32.84	44.69	97.98*	98.80*	55.08	62.51
	SGA	33.41	44.64	37.54	47.76	99.69*	99.69*	58.93	65.83
	LSSA (Ours)	38.88	49.66	40.63	52.95	99.72*	99.77*	69.02	73.17
CLIP _{CNN}	PGD	7.01	10.62	10.08	13.65	4.88	10.70	76.99*	84.20*
	BERT-Attack	23.38	34.64	24.58	29.61	51.28	57.49	54.43*	62.17*
	Sep-Attack	26.53	39.29	30.26	41.51	50.44	57.11	88.72*	92.49*
	Co-Attack	29.83	41.97	32.97	43.72	53.10	58.90	96.72*	98.56*
	SGA	31.61	43.00	34.81	45.95	56.62	60.77	99.61*	99.80*
	LSSA (Ours)	37.92	47.90	34.81	47.18	61.30	68.68	99.88*	99.99*

We extend the perturbation magnitude of adversarial images and conduct more experiments on the Flickr30K dataset. As shown in Table 7, our LSSA method shows significant enhancements compared to other advanced attacks, achieving high attack success rates even on models with different architectures. Specifically, in the white-box attack setting, the attack success rate of LSSA nearly approaches 100%. In the black-box attack setting, the transferable attack success rate for the ITR task increases by over 15% even when the model architecture differs. It indicates that even with larger perturbations, LSSA remains the most advanced transferable attack, demonstrating the robustness of our method.

B.4 Parameter Study

In this section, we present the impact of the sampled number of images in adversarial text generation on the transferable attack. The specific experimental results are shown in the following.

Sampled number in text generation. We con-

duct experiments to explore the performance of different sampled numbers around the adversarial images. As shown in Table 8, the increase of sampled number can effectively boost the white-box attack performance and slightly improve the black-box attack success rate, indicating that the performance of adversarial attacks can be enhanced by introducing additional sample information. To balance between the effectiveness and efficiency, we set the sampled number to 20.

Momentum decay μ . We conduct experiments to investigate the impact of momentum decay. As shown in Figure 9, the large momentum coefficients μ might lead to unstable convergence, while the small one might make the attacker hard to escape from the suboptimal area. Therefore, the attack performance of LSSA reaches peak when $\mu = 1.0$.

Sampling boundary ϵ_0 . We conduct detailed experiments to further study the impact of sampling boundary ϵ_0 . As shown in Figure 10, the larger sampling boundary results in a slight decrease in

Table 7: The attack success rate (%) of multimodal adversarial examples with perturbation boundary $\epsilon_v = 8/255$ against different VLP models compared with state-of-the-art methods on image-text retrieval task. The source column represents the VLP models used for crafting multimodal adversarial examples on Flickr30K dataset. * indicates white-box attacks. The higher attack success rate indicates the better performance.

Source	Attack	ALBEF		TCL		CLIP _{ViT}		CLIP _{CNN}	
		TR R@1	IR R@1	TR R@1	IR R@1	TR R@1	IR R@1	TR R@1	IR R@1
ALBEF	PGD	93.74*	94.43*	24.03	27.9	10.67	15.82	14.05	19.11
	BERT-Attack	11.57*	27.46*	12.64	28.07	29.33	43.17	32.69	46.11
	Sep-Attack	95.72*	96.14*	39.30	51.79	34.11	45.72	35.76	47.92
	Co-Attack	97.08*	98.36*	39.52	51.24	29.82	38.92	31.29	41.99
	SGA	99.79*	99.95*	88.30	87.81	37.30	46.17	39.85	50.67
	LSSA (Ours)	99.90*	99.98*	96.00	96.07	54.36	60.60	54.79	64.63
TCL	PGD	35.77	41.67	99.37*	99.33*	10.18	16.30	14.81	21.10
	BERT-Attack	11.89	26.82	14.54*	29.17	29.69	44.49	33.46	46.07
	Sep-Attack	52.45	61.44	99.58*	99.45*	37.06	45.81	37.42	49.91
	Co-Attack	49.84	60.36	91.68*	95.48*	32.64	42.69	32.06	47.82
	SGA	93.12	92.89	98.42*	98.76*	35.34	45.75	40.10	50.22
	LSSA (Ours)	98.23	98.50	100.0*	100.0*	53.50	60.60	56.45	66.52
CLIP _{ViT}	PGD	3.13	6.48	4.43	8.83	69.33*	84.79*	13.03	17.43
	BERT-Attack	9.59	22.64	11.80	25.07	28.34*	39.08*	30.40	37.43
	Sep-Attack	7.61	20.58	10.12	20.74	76.93*	87.44*	29.89	38.32
	Co-Attack	8.55	20.18	10.01	21.29	78.53*	87.50*	29.50	38.49
	SGA	22.63	35.15	26.55	37.26	99.26*	99.10*	54.92	61.41
	LSSA (Ours)	45.88	57.18	45.10	56.57	100.0*	100.0*	77.27	80.96
CLIP _{CNN}	PGD	2.29	6.15	4.53	8.88	5.40	12.08	89.78*	93.04*
	BERT-Attack	8.86	23.27	12.33	25.48	27.12	37.44	30.40*	40.10*
	Sep-Attack	9.38	22.99	11.28	25.45	26.13	39.24	93.61*	95.30*
	Co-Attack	10.53	23.62	12.54	26.05	27.24	40.62	95.91*	96.50*
	SGA	16.37	28.74	18.76	33.14	38.52	52.00	99.11*	99.49*
	LSSA (Ours)	31.49	43.92	35.72	47.19	64.17	69.10	100.0*	100.0*

attack success rate. It may be because the too-large sampling boundary causes the sampled images to be too far from the adversarial image. However, the attack success rate does not decrease significantly, possibly because the absolute magnitude of the sampling boundary is not large, keeping the sampling images and the original image within the same area. Consequently, the differences between the generated adversarial text and the original and sampled image regions remain substantial.

B.5 Ablation study

To verify the effectiveness of each module, we conduct a series of ablation studies. The following conclusions can be drawn as shown in table 9. First, SGA can achieve the desired effect by performing only the first two steps. Furthermore, combined with the sampling module, the attack performance of adversarial examples is boosted on the TCL model, improving 3.48% (TR R@1) and 2.93% (IR R@1). On the other hand, integrated with the local shuffle module can also enhance the attack

success rate, gaining an improvement of 8.11% and 5.5%. After combining the sampling and local shuffle modules, the success rate of white-box attacks also reaches the highest 99.06% and 98.55%, and the success rate of black-box attacks increases by 11.48% and 8.95% compared with SGA. Finally, momentum further improves the adversarial transferability, and LSSA obtains the best adversarial transferable performance while the while-box attack performance only slightly decays.

Table 8: The attack success rate (%) of adversarial text generation with different sampled numbers against different VLP models. The source column represents the VLP models used for crafting multimodal adversarial examples on the Flickr30K dataset. * indicates white-box attacks. The higher attack success rate indicates the better performance.

Source	Number	ALBEF		TCL		CLIP _{ViT}		CLIP _{CNN}	
		TR R@1	IR R@1	TR R@1	IR R@1	TR R@1	IR R@1	TR R@1	IR R@1
ALBEF	0	96.45*	96.40*	56.16	64.38	35.95	48.16	40.38	50.23
	1	97.50*	97.10*	59.33	65.88	35.83	48.13	40.45	50.33
	2	97.50*	97.10*	59.33	65.90	35.83	48.20	40.45	50.33
	5	97.50*	97.10*	59.33	65.88	35.98	48.22	40.47	50.52
	10	97.50*	97.10*	59.33	65.90	35.98	48.22	40.87	50.52
	20	97.50*	97.10*	59.33	65.88	36.32	48.49	40.87	50.77
	30	97.50*	97.10*	59.33	65.90	36.32	48.49	40.91	50.80

Table 9: Ablation study of the proposed LSSA. SGA₂ indicates a variant of SGA which performs only the first two steps. SA and LS denote the sampling and local shuffle modules, respectively. The difference between SGA₂+SA+LS and LSSA is the former does not contain momentum while the latter contains. ALBEF is used as the white-box model to craft multimodal adversarial examples on the Flickr30K dataset. * indicates white-box attacks. The higher attack success rate indicates the better performance.

Source	Attack	ALBEF						TCL					
		TR R@1	TR R@5	TR R@10	IR R@1	IR R@5	IR R@10	TR R@1	TR R@5	TR R@10	IR R@1	IR R@5	IR R@10
ALBEF	SGA	97.24*	94.59*	92.20*	97.08*	94.30*	92.38*	44.68	24.02	16.13	55.57	36.01	27.60
	SGA ₂	96.14*	92.69*	90.20*	96.66*	93.31*	91.41*	45.42	22.41	15.53	56.90	37.34	29.14
	SGA ₂ +SA	97.81*	95.89*	93.80*	97.26*	94.67*	92.98*	48.16	26.53	18.84	58.52	38.12	29.77
	SGA ₂ +LS	98.12*	95.89*	94.46*	98.22*	96.37*	95.13*	52.79	31.16	23.85	61.07	41.68	32.86
	SGA ₂ +SA+LS	99.06*	97.70*	96.30*	98.55*	96.86*	95.67*	56.16	35.18	26.25	64.52	45.34	36.21
	LSSA	97.29*	94.59*	92.50*	97.22*	94.09*	91.99*	59.43	36.58	26.45	65.62	47.14	38.40

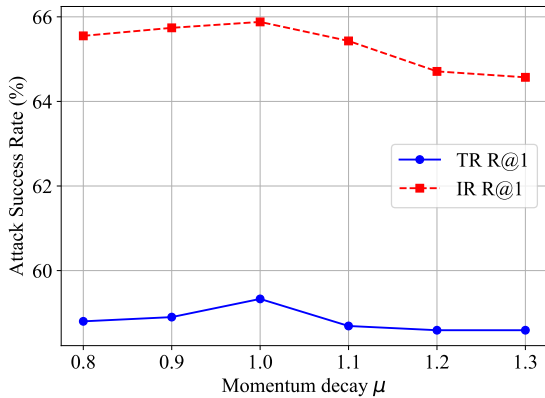


Figure 9: The attack success rate (%) of LSSA with different momentum decay μ on the Flickr30K dataset. The source and target model are ALBEF and TCL model, respectively.




Figure 10: The attack success rate (%) of LSSA with different sample magnitude ϵ on the Flickr30K dataset. The source and target model are ALBEF and TCL model, respectively.