Targeted Adversarial Self-Supervised Learning

Minseon Kim[1], **Hyeonjeong Ha**[1], **Sooel Son**[1], **Sung Ju Hwang**[1,2]
[1]Korea Advanced Institute of Science and Technology (KAIST), [2]AITRICS
{minseonkim, hyeonjeongha, sl.son, sjhwang82}@kaist.ac.kr

Abstract

Recently, unsupervised adversarial training (AT) has been extensively studied to attain robustness with the models trained upon unlabeled data. To this end, previous studies have applied existing supervised adversarial training techniques to self-supervised learning (SSL) frameworks. However, all have resorted to untargeted adversarial learning as obtaining targeted adversarial examples is unclear in the SSL setting lacking of label information. In this paper, we propose a novel targeted adversarial training method for the SSL frameworks. Specifically, we propose a target selection algorithm for the adversarial SSL frameworks; it is designed to select the most confusing sample for each given instance based on similarity and entropy, and perturb the given instance toward the selected target sample. Our method is readily applicable to general SSL frameworks that only uses positive pairs. We validate our method on benchmark datasets, on which it obtains superior robust accuracies, outperforming existing unsupervised adversarial training methods.

1 Introduction

Recent self-supervised learning (SSL) frameworks [17, 8, 34] mostly resort to maximizing the consistency across two differently augmented samples of the same instance, using an additional momentum encoder [17], without any negative pairs or additional networks [8, 34]. Such non-contrastive SSL frameworks using only positive pairs are shown to obtain representations with superior generalization performance compared to contrastive counterparts in a more efficient manner. However, leveraging untargeted adversarial attacks in these SSL frameworks results in a suboptimal performance. BYORL [15], an adversarial SSL framework using only positive pairs, obtains much lower robust accuracies than those of adversarial contrastive-learning SSL methods on the benchmark datasets (Table 3) compare to RoCL [21] and ACL [20] that have proposed adversarial SSL methods based on contrastive learning frameworks. Then, what is the cause of such suboptimal robustness in a non-contrastive adversarial SSL framework?

We observe that this limited robustness mainly comes from the suboptimality of untargeted attacks; adversarial examples generated by the deployed untargeted attacks are ineffective in improving robustness in non-contrastive adversarial SSL frameworks. As shown in Figure 1, the attack in the inner loop of the adversarial training loss, designed to minimize the distance between two differently augmented samples, perturbs a given example to a random position in the latent space. Thus, the ineffectiveness of the untargeted attacks in non-contrastive SSL frameworks mostly comes from their inconsideration of other instances.

To tackle this issue, we propose **T**argeted **A**ttack for **RO**bust **S**elf-**S**upervised learning (TAROSS). TAROSS is designed to enhance robustness of a non-contrastive SSL framework with only positive pairs, such as BYOL [17] and SimSiam [8], by conducting **targeted attacks**, which perturbs the given instance toward a target. However, this leads to the question of which direction we want to perform the targeted attack to, that is unclear in unsupervised adversarial learning without class labels. To address this point, we propose a simple yet effective target selection algorithm based on the similarity and entropy between instances.

ML Safety Workshop, 36th Conference on Neural Information Processing Systems (NeurIPS 2022).

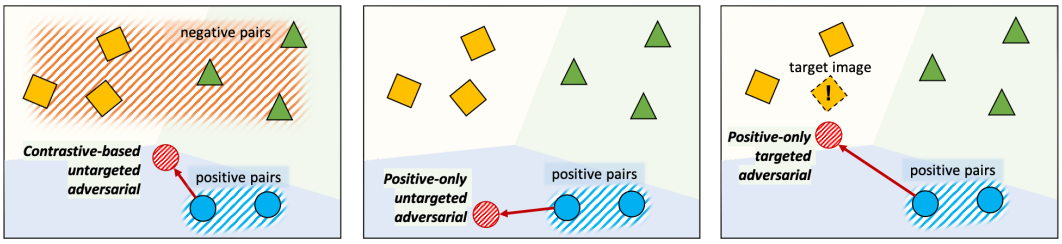

(a) contrastive adversarial attack (b) positive-only similarity untargeted attack (c) positive-only similarity targeted attack

Figure 1: **Motivation.** In supervised AT, perturbation is generated to maximize the cross-entropy loss. In adversarial contrastive SSL (a), perturbation is generated to minimize the similarity between positive pairs while maximizing between negative pairs. Then, the adversarial examples may be pushed to the space of other classes as negative samples may mostly contain samples from other classes. However, in positive-only adversarial SSL (b), minimizing the similarity between positive pairs have weaker constraints in generating effective adversarial examples than (a). To overcome such a limitation, we suggest a selective targeted attack that maximizes the similarity to the most confusing target instance (c).

Our main contributions can be summarized as follows:

- We demonstrate that achieving comparable robustness in the positive-only self-supervised learning (SSL) with contrastive-based SSL is difficult due to ineffective adversarial inputs generated by untargeted attacks.
- We propose a novel targeted adversarial attack scheme which perturbs the target sample in the direction to the most confusing instance to it, based on similarity and entropy.
- We experimentally confirm that the proposed targeted adversarial SSL framework is able to obtain significantly high robustness, outperforming previous adversarial SSL methods.

2 Related Work

Adversarial Self-Supervised Learning The first adversarial SSL method employed contrastive learning to achieve a high level of robustness [21, 20] without any class labels. Adversarial self-supervised contrastive learning [21, 20] generated an instance-wise adversarial example that maximizes the contrastive loss against their positive and negative samples by conducting untargeted attacks. Both methods achieved robustness with the cost of requiring large computation power due to a large batch size for contrastive learning. On the other hand, Gowal et al. [15] utilized only positive samples to obtain adversarial examples by maximizing the similarity loss between the latent vectors from the online and target networks, enabling this method free for batch size. However, it exhibited relatively poorer robustness than that of self-supervised contrastive learning, even with an advanced SSL framework. Despite this advanced framework, i.e., non-contrastive SSL, that employs only positive pairs, robustness is not guaranteed with a simple combination of untargeted adversarial learning and advanced SSL. To overcome such vulnerability in non-contrastive SSL, we propose a targeted attack leveraging a novel score function designed to improve robustness.

3 Targeted Adversarial Self-supervised Learning

Targeted adversarial self-supervised learning We argue that leveraging untargeted adversarial attacks in an SSL framework using positive-only pairs still leaves large room for better robustness. We can simply improve robustness in non-contrastive SSL by changing the inner loop to a randomly targeted attack as follows:

$$\delta^{t+1} = \Pi_{B(0,\epsilon)}\Big(\delta^t + \alpha \texttt{sign}\Big(\nabla_{\delta^t}\mathcal{L}_{\texttt{target}}\big(f(\theta, x+\delta^t), f(\theta, x')\big)\Big)\Big), \tag{1}$$

where $\mathcal{L}_{\texttt{target}} = -\mathcal{L}_{\texttt{loss}}$, x' is a randomly selected target within the batch. To precisely elaborate our observation in the latter section, we introduce a score function ($\mathcal{S}(x, \texttt{batch})$) to select the target ($x'$) within the batch. The score function ($\mathcal{S}$) can be a random-sampling function or a designed function as Eq. 3. In other words, the score function outputs the target (x') corresponding to the base image (x), then the targeted attack generates perturbation that maximizes the similarity to target x'.

In Table 1, an untargeted attack makes the model more vulnerable when the non-contrastive SSL frameworks utilize only positive pairs, such as in BYOL [17], and SimSiam [8]. However, when we modify the inner loop of the untargeted attack to the targeted attack with the simple random sampling score function as Eq. 1., non-contrastive SSL achieves better robustness.

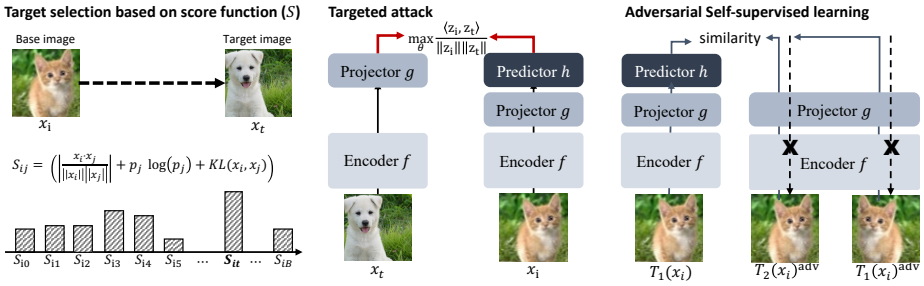

Figure 2: **Overview of TAROSS.** Our approach consists of three parts: target selection, targeted attack, and adversarial self-supervised learning (SSL). We propose a simple similarity- and entropy-based target selection algorithm that selects the maximum score target based on score function ($\mathcal{S}$) (left). When we select the target image (x_t) for each instances, we conduct a targeted attack (middle). Then, we train the model with targeted adversarial examples and clean images as SSL (right).

Positive-paired self-supervised learning. We first describe an SSL method, SimSiam [8], that learns visual representations with only positive pairs using a stop-gradient. Let us denote the dataset $\mathcal{D} = \{X\}$, and transformation set $\mathbf{T}$ that augments the images $x \in X$. SimSiam consists of the encoder f, followed by the projector g, and the predictor h; each of g and h is a multi-layer perceptron (MLPs). To learn visual representations, SimSiam maximizes the cosine similarity between the positive pairs as follows:

Table 1: Experimental results against PGD ℓ_∞ attacks on ResNet18 trained on the CIFAR-5. The CIFAR-5 consists of airplane, automobile, bird, cat, and deer from CIFAR-10.

Method	Attack	Clean	PGD ℓ_∞
BYOL	Untargeted	75.40	21.00
BYOL	Randomly targeted	83.50	**36.50**
SimSiam	Untargeted	66.36	36.53
SimSiam	Randomly targeted	77.08	**47.58**

$$\mathcal{L}_{\texttt{SimSiam}}(\mathbf{T}_1(x), \mathbf{T}_2(x)) = -\frac{1}{2}\frac{p_1}{||p_1||_2} \cdot \frac{z_2}{||z_2||_2} - \frac{1}{2}\frac{p_2}{||p_2||_2} \cdot \frac{z_1}{||z_1||_2}, \tag{2}$$

where $p_1 = h(g(f(\mathbf{T}_1(x))))$ and $z_2 = g(f(\mathbf{T}_2(x)))$ are output vectors of the predictor h and the projector g, respectively. Before calculating the loss, SimSiam detaches the gradient on the z which called stop-gradient. Stop-gradient helps the model to form proper visual representations without any momentum network but makes the single model acts like a momentum network.

Entropy and similarity based target selection. From the observation in the previous paragraph, we argue that selecting target instances that are most confused and are adequately close to the base instance is effective for targeted attacks. For this, we design the score function based on the similarity and entropy values; it does not need any class information, as follows:

$$\mathcal{S}(\mathbf{T}_1(x)_i, \mathbf{T}_2(x)_j) = \frac{e_i}{||e_i||_2} \cdot \frac{e_j}{||e_j||_2} + (p_j/\tau)\log(p_j/\tau), \tag{3}$$

where $p_i = h(g(f(\mathbf{T}_1(x)_i)))$ and $e_i = f(\mathbf{T}_1(x)_i)$ are output vectors of predictor h and encoder f, respectively. Overall, the score function $\mathcal{S}$ consists of a cosine similarity term, and an entropy term. The cosine similarity is calculated between features of base images and candidate images in the differently augmented batch ($\mathbf{T}_2$). The entropy is calculated with the assumption that the vector p is an instance's logit as Caron et al. [6]. We design the score function to naturally selects the target that is mostly confused with base images without any class labels (Figure 1). We also verify our score function selects targets as our intention in the experiment section (Figure 3b).

Robust self-supervised learning with target attack (TAROSS). In non-contrastive SSL, we make a positive pair, i.e., $\mathbf{T}_1(x), \mathbf{T}_2(x)$, with differently transformed augmentation. As shown in Figure 2, to generate adversarial examples, we first select the target images $\mathbf{T}_1(x)', \mathbf{T}_2(x)'$ for each base image $\mathbf{T}_1(x), \mathbf{T}_2(x)$, respectively, which have the maximum score from score function ($\mathcal{S}$). Here, we employ score function ($\mathcal{S}$) in Eq. 3. Then, we generate adversarial examples, i.e., $\mathbf{T}_1(x)^{adv}, \mathbf{T}_2(x)^{adv}$, for each transformed input with our suggested targeted attack (Eq. 1) where $\mathcal{L}_{\texttt{target}} = -\mathcal{L}_{\texttt{SimSiam}}$. Finally, we maximize the agreement between adversarial images $\mathbf{T}_1(x)^{adv}$,

and $\mathbf{T}_2(x)^{adv}$ with clean image $\mathbf{T}_1(x)$ as follows:

$$\mathcal{L}_{\texttt{TAROSS}} = \mathcal{L}(\mathbf{T}_1(x), \mathbf{T}_1(x)^{adv}) + w \cdot \mathcal{L}(\mathbf{T}_1(x)^{adv}, \mathbf{T}_2(x)^{adv}) \tag{4}$$

where $\mathcal{L}$ is Eq. 2. Since all three instances have the same identity, we maximize the similarity between the clean image and adversarial examples.

4 Experiment

4.1 Evaluating Robustness

Robustness with targeted/untargeted attack. In Table 3, as we observed in the previous section, when replacing untargeted attacks with targeted attacks in positive-only SSL (BYOL, SimSiam), it obtains robustness comparable to contrastive adversarial SSL (RoCL, ACL), which verifies our motivation. Therefore, it is feasible to improve robustness even in positive-only SSL while attaining the pros of positive-only SSL frameworks. Interestingly, the targeted attacks helped improve the clean accuracy in SSL when compared with those using untargeted attacks in positive-only SSL. We conjecture that untargeted attacks are not only inadequate to learn robust features but also hinder to learn good visual representation of natural images.

Table 2: Experimental results against white box attacks on ResNet18 trained on the CIFAR10. We evaluate adversarial SSL with linear evaluation and robust linear evaluation.

SSL framework	Method	Clean	PGD ℓ_∞	AutoAttack
SimCLR	RoCL [21]	78.14	42.89	27.19
SimCLR	ACL [20]	79.96	39.37	35.97
BYOL	Untargeted attack	72.65	16.20	0.01
BYOL	**TAROSS**	**84.52**	**31.19**	**21.01**
SimSiam	Untargeted attack	71.78	37.28	32.41
SimSiam	**TAROSS**	**74.06**	**44.71**	**36.39**

Robustness compared to contrastive-based approaches. We compare our TAROSS to contrastive adversarial SSL methods to show that our approach could make positive-only SSL to achieve robustness comparable to previous works. Notably, as shown in Table 3, the targeted attacks allow the positive-only SSL model to have better robustness than that of contrastive adversarial SSL (RoCL, ACL) in linear evaluation. Specifically, when we train the fc layer, our model achieves 36.39% of robustness against strong AutoAttack which shows the effectiveness of our proposed method.

4.2 Analysis of Targeted Attacks

Analyze the distribution of the target class. We examine which class is selected as a target by the score function ($\mathcal{S}$). We test the selected target of the single class (airplane) to the adversarial supervised trained models [25] to obtain the probability and prediction for analyzing the class distribution. As shown in Figure 3b, half of the target images are the same class as the base image. And the next most frequently selected class was the class with the second highest probability that the images of the corresponding class were most confused (Figure 3a). The score function we proposed selects images of the same class when the image in the center of the cluster, and images near the decision boundary that have a high probability to predict to a wrong class select images of a different class.

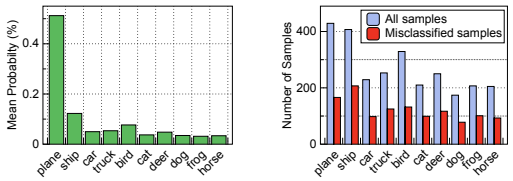

(a) Mean predict probability of base images (b) Distribution of class of targeted images

Figure 3: Analysis of target from score function ($\mathcal{S}$)

5 Conclusion

In this paper, we showed that a simple combination of supervised adversarial training with SSL is highly suboptimal due to the ineffectiveness of adversarial examples generated by the untargeted adversarial attacks perturbed to random places without the consideration of decision boundaries. Based on this observation, we propose an instance-wise targeted attack scheme for an adversarial SSL framework. This scheme selects the target instance based on the similarity and entropy, such that the given instance is perturbed to be confused with the selected target. Our targeted adversarial SSL framework obtains representation that achieves better robustness than the state-of-the-art adversarial SSL frameworks, including contrastive ones, without using any negative pairs or additional networks. We believe that our work opens a door to future research on the search for more effective targeted attacks, for adversarial SSL.

Acknowledgement

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2020-0-00153) and by Institute of Information communications Technology Planning Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2019-0-00075, Artificial Intelligence Graduate School Program(KAIST)). We thank Jihoon Tack, and Hayeon Lee for providing helpful feedbacks and suggestions in preparing an earlier version of the manuscript.

References

[1] M. Andriushchenko, F. Croce, N. Flammarion, and M. Hein. Square attack: a query-efficient black-box adversarial attack via random search. In *European Conference on Computer Vision*, pages 484–501. Springer, 2020.

[2] A. Athalye, N. Carlini, and D. Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International Conference on Machine Learning*, 2018.

[3] J. Buckman, A. Roy, C. Raffel, and I. Goodfellow. Thermometer encoding: One hot way to resist adversarial examples. 2018.

[4] N. Carlini and D. Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE symposium on security and privacy (sp)*, pages 39–57. IEEE, 2017.

[5] Y. Carmon, A. Raghunathan, L. Schmidt, P. Liang, and J. C. Duchi. Unlabeled data improves adversarial robustness. *Advances in Neural Information Processing Systems*, 2019.

[6] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin. Emerging properties in self-supervised vision transformers. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 9650–9660, 2021.

[7] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, 2020.

[8] X. Chen and K. He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021.

[9] F. Croce and M. Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International Conference on Machine Learning*, pages 2206–2216. PMLR, 2020.

[10] F. Croce and M. Hein. Minimally distorted adversarial examples with a fast adaptive boundary attack. In *International Conference on Machine Learning*, pages 2196–2205. PMLR, 2020.

[11] G. S. Dhillon, K. Azizzadenesheli, Z. C. Lipton, J. Bernstein, J. Kossaifi, A. Khanna, and A. Anandkumar. Stochastic activation pruning for robust adversarial defense. *International Conference on Learning Representations*, 2018.

[12] A. Dosovitskiy, P. Fischer, J. T. Springenberg, M. Riedmiller, and T. Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(9):1734–1747, 2015.

[13] L. Fan, S. Liu, P.-Y. Chen, G. Zhang, and C. Gan. When does contrastive learning preserve adversarial robustness from pretraining to finetuning? *Advances in Neural Information Processing Systems*, 34, 2021.

[14] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.

[15] S. Gowal, P.-S. Huang, A. van den Oord, T. Mann, and P. Kohli. Self-supervised adversarial robustness for the low-label, high-data regime. In *International Conference on Learning Representations*, 2021.

[16] S. Gowal, S.-A. Rebuffi, O. Wiles, F. Stimberg, D. A. Calian, and T. A. Mann. Improving robustness using generated data. *Advances in Neural Information Processing Systems*, 34: 4218–4233, 2021.

[17] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doersch, B. A. Pires, Z. D. Guo, M. G. Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 2020.

[18] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.

[19] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick. Momentum contrast for unsupervised visual representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.

[20] Z. Jiang, T. Chen, T. Chen, and Z. Wang. Robust pre-training by adversarial contrastive learning. In *Advances in Neural Information Processing Systems*, 2020.

[21] M. Kim, J. Tack, and S. J. Hwang. Adversarial self-supervised contrastive learning. *Advances in Neural Information Processing Systems*, 2020.

[22] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.

[23] A. Kurakin, I. Goodfellow, and S. Bengio. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*, 2016.

[24] Y. Le and X. Yang. Tiny imagenet visual recognition challenge. 2015.

[25] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.

[26] M. Noroozi and P. Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pages 69–84. Springer, 2016.

[27] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2536–2544, 2016.

[28] A. Shafahi, P. Saadatpanah, C. Zhu, A. Ghiasi, C. Studer, D. Jacobs, and T. Goldstein. Adversarially robust transfer learning. *International Conference on Learning Representations*, 2020.

[29] Y. Song, T. Kim, S. Nowozin, S. Ermon, and N. Kushman. Pixeldefend: Leveraging generative models to understand and defend against adversarial examples. *International Conference on Learning Representations*, 2017.

[30] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

[31] Y. Tian, D. Krishnan, and P. Isola. Contrastive multiview coding. In *European Conference on Computer Vision*, 2020.

[32] Y. Tian, C. Sun, B. Poole, D. Krishnan, C. Schmid, and P. Isola. What makes for good views for contrastive learning. In *Advances in Neural Information Processing Systems*, 2020.

[33] D. Wu, S.-T. Xia, and Y. Wang. Adversarial weight perturbation helps robust generalization. *Advances in Neural Information Processing Systems*, 33, 2020.

[34] J. Zbontar, L. Jing, I. Misra, Y. LeCun, and S. Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021.

[35] H. Zhang, Y. Yu, J. Jiao, E. P. Xing, L. E. Ghaoui, and M. I. Jordan. Theoretically principled trade-off between robustness and accuracy. In *International Conference on Machine Learning*, 2019.

[36] R. Zhang, P. Isola, and A. A. Efros. Colorful image colorization. In *European Conference on Computer Vision*, pages 649–666. Springer, 2016.

Supplementary Material

Targeted Adversarial Self-Supervised Learning

A Related works

Adversarial Training Szegedy et al. [30] showed that imperceptible perturbation to an input image may lead a DNN model to misclassify a given input into a false label, demonstrating the vulnerabilities of DNN models against adversarial attacks. Goodfellow et al. [14] proposed the Fast Gradient Sign Method (FGSM), which perturbs a given input to add imperceptible noise in the gradient direction of decreasing the loss of a target model. They also demonstrated that training a DNN model over perturbed as well as clean samples improves the robustness of the model against FGSM attacks. Follow-up works [23, 4] proposed diverse gradient-based strong attacks, and Madry et al. [25] proposed a projected gradient descent (PGD) attack and a robust training algorithm leveraging a minimax formulation; it finds an adversarial example that achieves a high loss while minimizing the adversarial loss over given data points. Due to a surge of interest in achieving robustness, various defense mechanisms [29, 3, 11] have been proposed.

However, Athalye et al. [2] showed that many of the previous studies depend on gradient masking, which results in obfuscated gradients in the representation space. At the same time, this renders a target model highly vulnerable to stronger attacks that circumvent obfuscated gradients. TRADES [35] proposed to minimize the Kullback-Leibler divergence (KLD) over clean examples and their adversarial counterparts, thus enforcing consistency between their predictions. They further showed that there is a theoretical trade-off relationship in achieving both clean accuracy and robustness. Recently, leveraging additional unlabeled data [5] and conducting additional attacks [33] have been proposed. Carmon et al. [5] proposed using Tiny ImageNet [24] as pseudo labels, and Gowal et al. [16] proposed using generated images from generative models to learn richer representations with additional data.

Self-Supervised Learning Due to the high annotation cost of labeling data, SSL has recently gained a large attention [12, 36, 31, 32, 34]. Previously, SSL focused on solving a pre-task problem of collaterally obtaining visual representation, such as solving the jigsaw puzzle [26], predicting the relative position of two regions [12], or impainting the masked area [27]. Contrastive SSL coined with data augmentation [7, 19] has achieved impressive performance in SSL. On the other hand, other previous studies employed a momentum network to learn visual differences between two augmented images [17] or to mimic the momentum network with stop-gradient [8].

B Baselines.

- **RoCL [21].** RoCL is SimCLR [7] based adversarial self-supervised learning methods. We experiment with the official code[1]. To make a fair comparison, we set the attack step to 10 as other baselines. We train the model with 1,000 epochs under the LARS optimizer with weight decay $2e-6$ and momentum with 0.9. For the learning rate schedule, we also followed linear warmup with cosine decay scheduling. We set a batch size of 512 for all datasets (CIFAR10, CIFAR100, STL10). For data augmentation, we use a random crop with 0.08 to 1.0 size, horizontal flip with a probability of 0.5, color jitter with a probability of 0.8, and grayscale with a probability of 0.2 for RoCL training.
- **ACL [20]** ACL is SimCLR [7] based adversarial self-supervised learning methods. We conduct the experiment with the official code[2]. To make a fair comparison, we set the attack step to 10 as other baselines. We train the model with 1,000 epochs. We set a batch size of 512 for STL10 dataset. For CIFAR10, and CIFAR100, we use the official pretrained checkpoints. For data augmentation, we use a random crop with 0.08 to 1.0 size, horizontal flip with a probability of 0.5, color jitter with a probability of 0.8, and grayscale with a probability of 0.2 for ACL training. We set PGD dual mode which calculates both clean and adversarial during the training.

[1] https://github.com/Kim-Minseon/RoCL

[2] https://github.com/VITA-Group/Adversarial-Contrastive-Learning

- **BYORL [15]** BYORL is BYOL [17] based adversarial self-supervised learning methods for low label regime. Since there is no official code for BYORL we implement the BYORL by ourselves. Though we mentioned as BYOL with untargeted attack in the Table 1, 3, 4 and 5, please note that BYOL with untargeted approach stands for Gowal et al. [15]. We implement based on BYOL from a self-supervised learning library[3]. We use the same CIFAR-10 setting in the library except for normalization. We exclude normalization in the data augmentation. To make a fair comparison, we implement on the ResNet18 with attack step 10 of PGD. As shown in supplementary materials in Gowal et al. [15], when the model is trained with 10 steps in ResNet34 it shows 37.88% of robustness. We conjecture that we have a different performance from the original paper because the original paper employs 40 steps of PGD in WideResNet34 to obtain the reported robustness which requires extraordinary computation power.

- **AdvCL [13]** AdvCL is SimCLR [7] based adversarial self-supervised learning which employ pseudo labels from the model that is pretrained on ImageNet [22] data. Even though the outstanding performance of AdvCL, we exclude this model as our baseline because the proposed methods require the model that is trained with the labels of ImageNet which we assume to have no label information for pretraining the model.

C Detailed description of experimental setups

C.1 Resource description.

All experiments are conducted with a two NVIDIA RTX 2080 Ti, except for the experiments with CIFAR100 experiments. For CIFAR100 experiments, two NVIDIA RTX 3080 are used. All experiments are processed in Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz.

C.2 Dataset description.

For experiments, we use CIFAR 10, CIFAR 100, CIFAR 5, and STL10. CIFAR 10 and CIFAR 100[4] consist of 50,000 training images and 10,000 test images with 10 and 100 classes, respectively. All CIFAR images are $32\times32\times3$ resolution (width, height, and channel). CIFAR 5 is a subset of CIFAR 10 which contains all images from 5 classes: airplane, automobile, bird, cat, and deer. We also test with determined classes. The STL10[5] consists of 5,000 training images and 8,00 test images with 10 classes. All STL10 images are $96\times96\times3$ resolution (width, height, and channel.) However, for our experiment in Table 4, 5, we resize the images into $32\times32\times3$ resolution (width, height, and channel).

C.3 Training detail.

For all methods, we train on the ResNet18 [18] with ℓ_∞ attacks with the attack strength of $\epsilon = 8/255$ and the step size of $\alpha = 2/255$, with the number of inner maximization iteration set to $K = 10$. For the optimization, we train every model for 800 epochs using the SGD optimizer with the learning rate of 0.05, weight decay of $5\text{e}{-4}$, and the momentum of 0.9. For data augmentation, we use a random crop with 0.08 to 1.0 size, horizontal flip with a probability of 0.5, color jitter with a probability of 0.8, and grayscale with a probability of 0.2. We exclude normalization for adversarial training. We set the weight of adversarial similarity loss w as 2.0. We use batch size 512 with two GPUs.

In the score function, we calculate the similarity score term and entropy term as shown in Equation 3. First, to exclude the positive pairs' similarity score we set the similarity score between positive pairs to -1. Then, to calculate the overall score, after obtaining the similarity score and entropy of each sample, we normalize each component with Euclidean normalization to balance each component to score function.

[3] https://github.com/vturrisi/solo-learn

[4] http://www.cs.toronto.edu/~kriz/cifar.html.

[5] https://ai.stanford.edu/~acoates/stl10/

Table 3: Experimental results against white box attacks on ResNet18 trained on the CIFAR10. We evaluate adversarial SSL with linear evaluation and robust linear evaluation. Clean is the accuracy of clean images. All models are evaluated with PGD ℓ_∞ with 20 steps of $\epsilon = 0.0314$ and AutoAttack [9]. Untargeted attack maximizes the training loss between the differently augmented but the same instances to generate adversaries. To see the effectiveness of our methods, we test TAROSS on positive-only SSL, i.e., BYOL, SimSiam (grey box).

Train type	Self-supervised framework	Method	Clean	PGD ℓ_∞	AutoAttack
Supervised	-	$\mathcal{L}_{\text{CE}}$	92.82	0.00	0.00
	-	AT [25]	83.07	53.78	49.23
Self-supervised linear evaluation	SimCLR	RoCL [21]	78.14	42.89	27.19
	SimCLR	ACL [20]	79.96	39.37	35.97
	BYOL	Untargeted attack	72.65	16.20	0.01
	BYOL	**TAROSS**	**84.52**	**31.19**	**21.01**
	SimSiam	Untargeted attack	71.78	37.28	32.41
	SimSiam	**TAROSS**	**74.06**	**44.71**	**36.39**
Self-supervised robust linear evaluation	SimCLR	RoCL [21]	76.53	47.51	30.22
	SimCLR	ACL [20]	77.17	40.67	39.13
	BYOL	Untargeted attack	54.01	27.24	4.49
	BYOL	**TAROSS**	**74.33**	**40.84**	**29.91**
	SimSiam	Untargeted attack	68.88	37.84	31.44
	SimSiam	**TAROSS**	**76.19**	**45.57**	**37.25**

C.4 Evaluation details.

PGD ℓ_∞ attack. For all PGD ℓ_∞ attacks used in the test time, we use the projected gradient descent (PGD) attack with the strength of $\epsilon = 8/255$, with the step size of $\alpha = 8/2550$, and with the number of inner maximization iteration set to $K = 20$ with the random start.

AutoAttack. We further test against strong gradient based attack, i.e., AutoAttack (AA) [9]. AutoAttack is an ensemble attack of four different attacks (APGD-CE, APGD-T, FAB-T [10], and Square [1]). AGPD-CE is an untargeted attack, APGD-T and FAB-T are targeted attacks. The Square is a black box attack. We use official code to test models[6].

Self-supervised learning. For self-supervised learning, we denote linear evaluation when we use only clean images to train the fully connected (fc) layer after the pretraining phase. When we denote robust linear evaluation, we train the fc layer with adversarial examples. While ACL uses partial fine-tuning to obtain their reported accuracy and robustness, to make a fair comparison we freeze the encoder and train only the fc layer. Robust fine-tuning is training all parameters including parameters of the encoder with adversarial examples. For linear evaluation, we followed baselines hyperparameters for each model. We train the baseline models with 150 epochs, 25 epochs, and 50 epochs for RoCL, ACL, and BYORL, respectively. We also followed their learning rate 0.1, 0.1, and $2 \times 10^{-}3$ for RoCL, ACL, and BYORL, respectively. On the other hand, we train our model with 100 epochs with a learning rate of 0.5 for linear evaluation. We use AT loss for robust linear evaluation except for ACL. For ACL, we use TRADES loss as the official code.

D Additional Experiment of Our Targeted Attack

Robustness compared to contrastive-based approaches. We compare our TAROSS to contrastive adversarial SSL methods to show that our approach could make positive-only SSL to achieve robustness comparable to previous works. To this end, we evaluate adversarial SSL with linear evaluation and robust linear evaluation against AutoAttack [9] and PGD [25] ℓ_∞ attacks. Notably, as shown in Table 3, the targeted attacks allow the positive-only SSL model to have better robustness than that of contrastive adversarial SSL (RoCL, ACL) in linear evaluation. Specifically, when we train the fc layer, our model achieves 36.39% of robustness against strong AutoAttack which shows the effectiveness of our proposed method. Interestingly, the SimSiam-based model has already obtained good robustness in the linear evaluation compared to the robust linear evaluation which takes $\times 4$ times to train. We conjecture that SimSiam-based TAROSS already has a robust well-cluster after the robust pretraining which does not need to find robust decision boundaries with the adversarial examples.

[6] https://github.com/fra31/auto-attack

Table 4: Results of self-supervised adversarial learning evaluated with linear evaluation in STL10 and CIFAR10. Rob. stands for robust accuracy against PGD ℓ_∞ attack with 20 steps and $\epsilon = 0.0314$.

		STL10		CIFAR100	
Method	Attack	Clean	Rob.	Clean	Rob.
RoCL	Untarget	52.63	19.19	45.99	17.17
ACL	Untarget	54.14	19.69	41.09	15.31
BYOL	Untarget	61.02	12.14	53.43	14.81
SimSiam	Untarget	43.88	12.70	27.53	14.20
SimSiam	**TAROSS**	**46.38**	**21.32**	**36.02**	**22.18**

Table 5: Results of adversarial transfer learning to CIFAR10 and STL10 from CIFAR100 pretrained models followed the experimental setting of Shafahi et al. [28]. Rob. stands for robust accuracy against PGD ℓ_∞ 20 attack with ϵ =0.0314.

CIFAR100 →		CIFAR10		STL10	
Method	Attack	Clean	Rob.	Clean	Rob.
RoCL	Untarget	73.93	18.62	74.06	19.06
ACL	Untarget	61.68	15.66	59.56	13.53
BYOL	Untarget	71.06	16.57	63.53	12.93
SimSiam	Untarget	50.05	25.43	34.54	**24.82**
SimSiam	**TAROSS**	**50.50**	**25.44**	**43.13**	22.46

Since the SimSiam directly maximizes the similarity between adversary and clean images only between the positive pairs with the single models, the gap between the representation of adversaries and clean examples may reduce.

E Evaluating the Quality of Robust Representation

Robustness on multiple benchmarks datasets. We validate our method on multiple benchmark datasets, such as CIFAR100 and STL10. In Table 4, our TAROSS consistently shows good robustness across different benchmark datasets when comparing with adversarial SSL frameworks using untargeted attacks. Especially, when the dataset becomes larger, such as CIFAR100, our method achieves even better robustness than contrastive learning-based approaches, e.g., RoCL, and ACL, where these approaches are sensitive to the number of same class samples in the single batch size. From this result, we argue that our targeted adversarial SSL contributes to learning better robustness in the larger dataset than contrastive-based SSL.

Transfer to different data. SSL is effective to utilize in several downstream tasks with a pretrained encoder. Therefore, we also evaluate how our robust pretrained features help in the different datasets in transfer tasks, which demonstrates the quality of our robust pretrained features in a different dataset. We followed the experimental setting as supervised adversarial transfer learning tasks [28] which freeze the encoder and train only the fc layer. We pretrained the model on CIFAR100 and evaluate the robust transferability to STL10, and CIFAR10. Once our model is trained, we transfer the robust features to different kinds of datasets with fewer epochs and parameters, which is efficient in that we do not need to conduct adversarial training from scratch. In Table 5, our model also shows impressive transferability both in CIFAR10 and STL10 compare to untargeted adversarial SSL. Moreover, while our model shows relatively low clean accuracy, TAROSS could obtain about 10% gain in robustness compared to contrastive-based adversarial SSL. From these results, we confirm that our methods generate robust representations that can be transferred to several downstream tasks, which previous adversarial SSL studies addressed.

F Experimental Details of Analysis

Analysis the distribution of target class To analyze the target from the score function ($\mathcal{S}$), we employ an adversarially supervised trained model. We calculate the score function that is trained with our TAROSS on SimSiam. We use a train set. For each class, we calculate the mean predict probability which is the average of all softmax outputs of target images from the supervised trained model. Further, we also count the number of samples that are predicted for each class. In Figure 3, the results are target images of the airplane as a base image. There is a similar tendency even though we change the base class to other classes as shown in the following Figure 4.

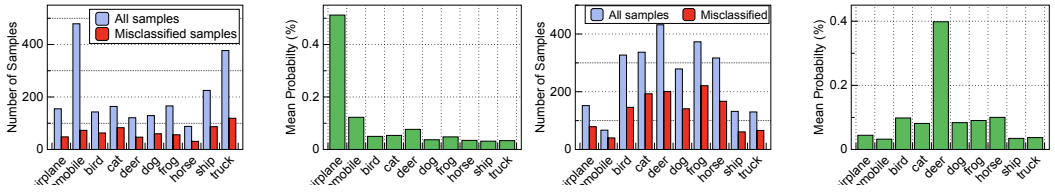

(a) Distribution of class of target of *automobile*. (b) Mean predict probability of *automobile* (c) Distribution of class of target of *deer*. (d) Mean predict probability of *deer*

Figure 4: Analysis of target distribution in different class

Visualization of embedding space To visualize the embedding of our targeted attack and untargeted attack, we use our TAROSS model that is trained on CIFAR10 as a feature extractor. We sample a few examples and conduct two types of attacks, untargeted attack, and targeted attack. To visualize more effectively we ignore the other seven classes in CIFAR10. We visualize clean examples from three classes and then visualize adversaries that are generated with our targeted attack and untargeted attack, respectively with dark blue.

G Additional Analysis of Targeted Attacks

Visualization of embedding space. To observe the differences between images that are generated with targeted versus untargeted attacks, we visualize the embedding space of targeted attack examples and untargeted attack examples. In Figure 5, adversarial examples are denoted with dark blue, and clean examples are denoted with light blue; both are instances of the same class. As shown in Figure 5a, untargeted adversarial examples are located near clean examples. On the other hand, targeted adversarial examples are located near the boundary of the cluster (Figure 5b). This visualization shows that our targeted attack generates relatively more effective adversarial examples than untargeted attacks, which is likely to push the decision boundary to have a better robust representation.

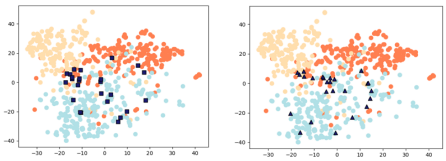

(a) Untargeted attack (b) Targeted attack

Figure 5: Visualize embedding