

Exploiting Supervised Poison Vulnerability to Strengthen Self-Supervised Defense

Jeremy Styborski^{1*}, Mingzhi Lyu^{2*}, Yi Huang^{1*}, and Adams Kong¹

¹ College of Computing and Data Science, Nanyang Technological University, Singapore

² Rapid-Rich Object Search (ROSE) Lab, Interdisciplinary Graduate Programme, Nanyang Technological University, Singapore
{styb0001,lyum0002,AdamsKong}@ntu.edu.sg, shellbyhuang@gmail.com

Abstract. Availability poisons exploit supervised learning (SL) algorithms by introducing class-related shortcut features in images such that models trained on poisoned data are useless for real-world datasets. Self-supervised learning (SSL), which utilizes augmentations to learn instance discrimination, is regarded as a strong defense against poisoned data. However, by extending the study of SSL across multiple poisons on the CIFAR-10 and ImageNet-100 datasets, we demonstrate that it often performs poorly, far below that of training on clean data. Leveraging the vulnerability of SL to poison attacks, we introduce adversarial training (AT) on SL to obfuscate poison features and guide robust feature learning for SSL. Our proposed defense, designated VESPR (Vulnerability Exploitation of Supervised Poisoning for Robust SSL), surpasses the performance of six previous defenses across seven popular availability poisons. VESPR displays superior performance over all previous defenses, boosting the minimum and average ImageNet-100 test accuracies of poisoned models by 16% and 9%, respectively. Through analysis and ablation studies, we elucidate the mechanisms by which VESPR learns robust class features. **Code:** <https://github.com/JStyborski/VESPR>

Keywords: Adversarial Training · Availability Poisoning · Poison Defense · Self-Supervised Learning

1 Introduction

The proliferation of open-source labeled data on the Internet has fueled the rapid development and application of deep neural networks across various domains. However, reliance on supervised learning in tasks like image classification exposes neural networks to shortcut learning, whereby models learn ‘easy’ features that are spuriously correlated with task labels in place of useful image features for accurate representation. Shortcut learning is an inherent behavior for neural networks [29, 56, 69, 72]. Availability poisons exploit this vulnerability by injecting malicious data into real-world data sources, thereby introducing shortcuts and

*Equal contribution

reducing the effectiveness of trained models. Crucially, availability poisons are imperceptible; they only perturb images slightly and do not modify image labels, such that poisoned images appear clean to human observers.

Models trained on poisoned data exhibit poor generalization on real-world clean data, posing significant security risks in critical applications like self-driving vehicles [61], anomaly detection [66] and medical AI [2, 3]. For instance, a self-driving system trained with poisoned images of stop signs may fail to recognize real images of stop signs in practical scenarios. A model trained on poisoned data will classify images based on the poisoned patterns rather than genuine features. Consequently, the model becomes incapable of correctly classifying benign images, jeopardizing its reliability in real-world applications.

Currently, there are various proposed defense methods against availability poisons [18, 21, 49, 53, 57, 63, 85, 87]. However, these methods are limited in their ability to defend against many popular availability poisons, leaving significant vulnerabilities to certain poisons while also suffering from reduced average performance. According to the surveys of poison defenses in Tables 7 and 8 (see Supplemental Material Sec. A), not all of these defense methods have been fully assessed on state-of-the-art poisons. Though most of these baseline defenses achieve robust performance on small-image datasets with few classes (e.g., CIFAR-10 [48] and SVHN [59]), our experiments reveal that they often fail to achieve comparable performance on high-resolution datasets with many classes (e.g., ImageNet-100 [20]), highlighting a critical research gap.

Poisoned perturbations establish shortcuts between labels and poisonous patterns in the images, prompting us to explore alternatives to traditional SL. SSL methods, such as contrastive learning [12, 34, 60], focus on learning image representations that are invariant to image transformations [58], rather than explicitly modeling the correlation between features and labels. In SSL, the combined effect of augmentations and the instance-discrimination objective filter out nuisance features introduced by poisons. Consequently, encoders trained with SSL tend to extract more semantic information from images and avoid class collapse [80]. In extending the study of SSL-based defense to multiple ImageNet-100 poisons, our experiments reveal that SSL unfortunately learns poison information alongside genuine features, challenging observations made largely on CIFAR-10 [32].

Motivated by the discrepancy in poison robustness between SL and SSL, we propose leveraging the poison weakness of SL to generate adversarial noise in order to improve augmentations for SSL, thereby preferentially capturing genuine image features and dismissing poisoned features. Thus, we introduce Vulnerability Exploitation of Supervised Poisoning for Robust SSL (VESPR). As illustrated in Fig. 2, VESPR trains an encoder using a multi-task loss function, while adversarial inputs are generated using gradients computed with the supervised cross-entropy loss. We evaluate the performance of VESPR against seven state-of-the-art poisoning methods [26, 27, 32, 40, 67, 79, 83] and compare it with six baseline defense methods [21, 32, 49, 53, 57, 85, 87] on the ImageNet-100 [20] and CIFAR-10 [48] datasets. Additionally, we conduct comprehensive ablation studies to verify and understand the effectiveness of VESPR.

Our work advances the state-of-the-art for poison defense methods through the following contributions:

- We extend the study of multiple defense methods, including SSL, to seven poisons on the CIFAR-10 and ImageNet-100 datasets. For all defenses, we find that the discrepancy between clean and poison training is exacerbated for large-image datasets.
- We exploit the vulnerability of SL to availability poisons in order to enhance the robustness of SSL training through our proposed method, VESPR. Integrating AT for SL within the SSL framework, VESPR generates optimized adversarial augmentations that preferentially emphasize robust image features and avoid poison shortcuts.
- Through extensive experiments, we demonstrate that VESPR achieves state-of-the-art performance against multiple prevalent poisons, improving both minimum and average poison defense beyond all other baselines.

The remainder of this paper is organized as follows: Sec. 2 reviews other works relevant to our research, Sec. 3 analyzes trends in shortcut learning and introduces the VESPR framework, and Sec. 4 covers experimental results, subsequent ablations, and corresponding analyses.

2 Related Works

2.1 Data Poisoning

The concept of data poisoning for computer vision was borne out of adversarial examples [30, 49, 57, 68], wherein images could be imperceptibly altered yet decimate supervised classification performance. Importantly, adversarial examples can transfer across models [19], such that generating poisons against unknown training methods becomes feasible. Adversarial Poisons (AP) is an early poisoning method that simply uses adversarial examples [26].

Synthetic poisons [67, 79, 83] are designed specifically for poisoning supervised classification methods. Synthetic poisons are crafted heuristically and commonly applied class-wise. They aim to introduce consistent class-wise shortcut features. The simplest synthetic poisoning method is the One-Pixel Shortcut [79], wherein a single pixel is altered to a consistent color for all images of a class. Min-max poisons [8, 25, 84] directly apply the data poisoning objective: find perturbations that minimize training loss on poison data while maximizing test loss on clean data. Min-min poisons [27, 32, 40] create ‘unlearnable examples’, such that a network trained on the poisoned samples will quickly learn subtle shortcuts without learning useful image features.

Unlike availability poisons, dirty-label attacks [6, 52] do not typically modify images but instead inject mislabelled images into datasets. Dirty-label are effective but easily detectable; they can only poison datasets with little or no human supervision. We do not consider dirty-label attacks in this paper and instead focus on poisons that perturb image features.

2.2 Poison Defenses

Poison defenses for SL are heretofore largely based on augmentations and adversarial training. Most poisons demonstrate reduced effectiveness against models trained with data augmentations like Cutout [21], Mixup [87], CutMix [85], and RandAugment [18]. Grayscale and JPEG-compression augmentations may be applied to counteract low-frequency and high-frequency poisons, respectively [53]. UEraser samples multiple augmentations and selects the one that maximizes classification loss [63]. Augmentation-based defenses filter out some non-robust poisonous patterns, but falter against poisons such as AP and CUDA [26, 67]. In fact, multiple poisoning papers [26, 32, 79] demonstrate their triumph over augmentation-based methods as a proof of their effectiveness. Gaussian noise augmentation applied to training data reduces the effectiveness of adversarial examples [16]. Unlike random noise, AT [57, 74, 82] generates optimized adversarial noise during training to improve model robustness. AT is largely considered a state-of-the-art poison defense, improving performance against multiple poisons. However, multiple poison methods still enforce a significant accuracy drop for models defended by AT [32, 67, 79], and some poisons even anticipate the use of AT during training [27]. Indeed, we demonstrate in Sec. 4.3 that these SL-based defenses underperform on multiple poisons.

SSL has also been uncovered as a method for learning robust image features from poisoned data sets [32]. Unfortunately, prior studies are limited in the extent to which they test various poisons, exploring SSL as poison defense against only one ImageNet-100 poison and three CIFAR-10 poisons (see Tables 7 and 8 in Supplemental Material Sec. A). We extend the study of SSL across seven poisons for CIFAR-10 and ImageNet-100. Furthermore, we are the first to investigate adversarial SSL methods [1, 11, 24, 42, 45, 46, 55] applied to poison defense.

2.3 Self-Supervised Learning

Autoencoders [23, 33, 36, 37, 47, 54, 75, 81] are considered the one of the first SSL methods, utilizing reconstruction-based loss with an encoder-decoder architecture to learn compressed image embeddings. Recently, invariance-based SSL methods have gained popularity. Contrastive methods [12, 15, 34, 58, 64, 70, 73, 77] typically rely on the InfoNCE loss function [60] to encourage all projections of views from the same source image to be similar, yet different from projections of views from other images. Non-contrastive methods [14, 31, 71, 78, 88–90] eliminate the requirement for negative samples, focusing their loss functions entirely on reducing the distance between same-source projections.

Like SL, SSL is vulnerable to feature suppression [43, 50, 65, 69, 78, 86, 90] and shortcut learning [13, 28]. Multiple works have proposed modifications to architecture or training objectives in order to mitigate feature suppression in SSL [39, 44, 51, 65, 76]. To ensure that SSL does not favor poison shortcuts, we employ AT on SL in order to generate adversarially-perturbed augmentations that shift focus towards robust image features.

2.4 Multi-Task Learning

More recently, multiple groups have proposed combining architectures and loss objectives in order to improve learned representation quality and downstream performance. Multiple works investigate the combination of instance discrimination and reconstruction objectives [4, 9, 22]. Islam et al. [41] empirically investigate the performance of supervised losses combined with contrastive SSL. Chen et al. [10] and Xue et al. [80] both demonstrate that combining supervised and self-supervised contrastive learning objectives avoids class-collapse and mitigates feature suppression. To our knowledge, no one has yet studied multi-task architectures or combined objectives for data poisoning.

3 Method

3.1 Formulation of Availability Poison

We describe the availability poison objective in Eq. 1. We assume a clean dataset $D_c = \{(x_{c,1}, y_1), (x_{c,2}, y_2), \dots, (x_{c,N}, y_N)\}$ with distribution \mathcal{P}_c , where $x_{c,i}$ is a benign image and y_i is the corresponding class label. The corresponding test dataset is $D'_t = \{(x'_{c,1}, y'_1), (x'_{c,2}, y'_2), \dots, (x'_{c,M}, y'_M)\}$, where each pair $(x'_{c,i}, y'_i)$ is drawn from distribution \mathcal{P}_c but not included in D_c . The goal of data poisoning is to generate poison dataset $D_p = \{(x_{p,1}, y_1), (x_{p,2}, y_2), \dots, (x_{p,N}, y_N)\}$ such that a classification model $CLS_{\theta_1^*}(\cdot)$ trained upon poison data D_p will perform poorly on clean test data D'_t . Poison sample $x_{p,i}$ is generated from corresponding clean sample $x_{c,i}$ via poisoning process $PSN_{\theta_2^*}(x_{c,i})$.

$$\begin{aligned} \theta_2^* &= \arg \max_{\theta_2} \sum_{i=0}^M L(CLS_{\theta_1^*}(x'_{c,i}), y'_i) \\ \text{s.t.} \quad \theta_1^* &= \arg \min_{\theta_1} \sum_{i=0}^N L(CLS_{\theta_1}(x_{p,i}), y_i), \\ x_{p,i} &= PSN_{\theta_2}(x_{c,i}), \quad d(x_{p,i}, x_{c,i}) \leq \epsilon \end{aligned} \tag{1}$$

L is the supervised learning loss function. In the context of classification, L is often set as the cross entropy (CE) loss, $L_{CE}(CLS_{\theta_1}(x), y) = y \log CLS_{\theta_1}(x)$, where $CLS_{\theta_1}(x)$ is the logit of input x with label y . Distance metric $d(\cdot)$ between $x_{p,i}$ and $x_{c,i}$ is bounded by a (small) threshold ϵ such that poisoned images retain high image quality and appear similar to their corresponding clean images. Eq. 1 says that poisoned images contain imperceptible perturbations yet maximally mislead poison-trained models to perform poorly on clean data.

3.2 Analysis of SL and SSL against Availability Poisons

Intuitively, if there exist features with high task correlation and low variance across training instances, deep learning models will grant greater weight on these




Fig. 1: (a) Mean of cosine similarity between representations of clean images and poison images (Psn-Cln Sim, $\text{cossim}[f(x_{p,i}), f(x_{c,i})]$). (b) Poison (Psn) and clean (Cln) classification accuracies tested with $x_{p,i}$ and $x_{c,i}$, respectively. We display curves for SL and SSL models, both trained on adversarial poison [26] data.

features. ‘Easy’ features that are spuriously correlated with the training task become shortcuts, which are quickly learned by SL models. Simplicity bias due to stochastic gradient descent may strengthen the preference for learning shortcuts [56, 80]. Availability poisons exploit the shortcut-learning weakness in order to mislead model training.

In order to demonstrate the effects of shortcut learning, we train a ResNet18 [35] encoder with an adversarial poison [26] via SL and compare representations of clean and poison images. As shown in Fig. 1a, representation similarity between a benign image $x_{c,i}$ and its poisoned version $x_{p,i}$ reduces rapidly (solid blue line in Fig. 1a), suggesting that the encoder quickly learns distinct representations based on the shortcut signals in the poisoned images. Though poison accuracy (solid blue line in Fig. 1b) quickly maximizes near 100%, clean accuracy never improves (dashed blue line in Fig. 1b).

Since availability poisons are often generated to confuse image classifiers, we explore SSL methods, which bypass the requirement for class labels. Invariance-based SSL techniques are designed to learn representations of images that are robust to input transformations (e.g., cropping, color shifts), yet distinct from representations of other images.

Contrastive learning is a classic invariance-based SSL framework, backed by extensive research [12, 15, 34, 58, 70] and proven improvements against some availability poisons [32]. Thus, we primarily utilize contrastive learning for SSL. As noted by Wang and Isola [77], the contrastive learning objective is optimized by two measures: representation alignment between positive samples and uniformity across all representations. Because SSL is **not** motivated by class labels, poison features shared by same-class images are less relevant. Additionally, the alignment objective of contrastive learning encourages augmentation invariance; poison features are obfuscated and representation learning favors clean features.

In Table 1, we explore feature suppression in SL and SSL by analyzing representation characteristics for their encoders across seven different poisons. We

Table 1: Models are trained on poisoned datasets with SL, SL with SSL-style augmentations, and SSL. Subsequent columns display average cosine similarity between representations of poison images of the same class (Psn In-Cls Sim, $\text{cossim}[f(x_{p,i}), f(x_{p,j})]$, where $y_i = y_j$), average cosine similarity between of poison image representations and their corresponding clean image representations (Psn-Cln Sim, $\text{cossim}[f(x_{p,i}), f(x_{c,i})]$), and poison representation feature diversity (Psn E-Rank, see [88, 90]). We average results across seven different poisons (see Supplemental Material Sec. B.1).

Method	Psn In-Cls Sim	Psn-Cln Sim	Psn E-Rank
SL	0.78	0.70	16.67
SL (SSL Aug)	0.79	0.79	16.43
SSL	0.47	0.86	59.61

instantiate the SSL objective L_{CL} with the commonly-used InfoNCE loss [60]. We begin by ablating the effect of augmentation: compared to SL, using SSL-style augmentations during SL barely affects Psn In-Cls Sim and Psn E-Rank, though there is a significant improvement in Psn-Cln Sim. SSL delivers significantly lower Psn In-Cls Sim than SL trained with or without SSL-style augmentation, demonstrating that the SSL training objective (and not SSL-style augmentation) is crucial for encoding instance-specific information rather than learning class-specific shortcuts. Likewise, higher Psn E-Rank and Psn-Cln Sim shows that SSL representations contain more features and are less influenced by poison signals, likely resulting in better generalization on clean test data.

Nevertheless, SSL models cannot inherently differentiate semantic information and poison information, and therefore cannot avoid poison features entirely. This is verified by Fig. 1b, where SSL still permits a significant gap between clean and poison accuracies. Therefore, the SSL model has learned some poisoned features along with true image features, leading the downstream classifiers to misclassify data lacking poisonous patterns.

3.3 VESPR

Motivated by these observations, we seek methods to improve augmentations for SSL that can better obscure poison features, thereby allowing the learning algorithm to learn true features. Noting the discrepancies between SL and SSL in representation similarities (Table 1) and performance (Fig. 1b), we conclude that the classification network largely learns poison features, distilling them from other image features. This betrays a weakness of availability poisons: once processed by a poisoned encoder, they are no longer subtle nor imperceptible! Drawing from theory for adversarial examples [26, 30], we hypothesize that poisoned classifiers can be exploited to generate adversarial examples that maximally muddle poison signals. We further hypothesize that incorporating these adversarial examples as augmentations will improve representations learned by SSL, thereby improving poison defense. Our final architecture is a combined SSL+SL

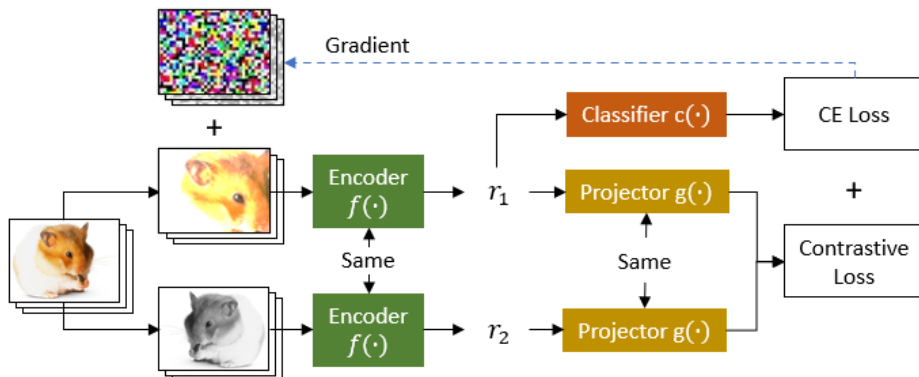


Fig. 2: Overview of VESPR architecture. For each input image, VESPR generates two views through augmentation and encodes them with encoder $f(\cdot)$ as representations r_1 and r_2 , respectively. Cross-entropy (CE) loss is calculated from the output of the classifier head, $c(r_1)$. The CE loss gradient is used to generate adversarial perturbations to the first view using projected gradient descent [49, 57]. The adversarial images are encoded to replace the original representations. Contrastive loss is calculated from projected representations, $g(r_1)$ and $g(r_2)$, using projections of other images as negative samples. The VESPR network is trained from the combined CE and contrastive losses.

framework, utilizing AT on the supervised loss. We designate our method as VESPR, Vulnerability Exploitation of Supervised Poisons for Robust SSL.

As shown in Fig. 2, VESPR contains feature encoder $f(\cdot)$, representation projector $g(\cdot)$, and classification head $c(\cdot)$. During training, each image of a batch is augmented into two different views. One view is taken as the source image to generate adversarial examples using the following objective function:

$$\begin{aligned} \delta^* &= \arg \max_{\delta} L_{CE}(c(f(x + \delta)), y) \\ \text{s.t. } & l_p(\delta) \leq \epsilon_{AT} \end{aligned} \quad (2)$$

where δ is an adversarial perturbation, x is an input view, $l_p(\cdot)$ denotes the p -norm, and ϵ_{AT} is the AT threshold, with $p = \infty$ in our case.

The encoder generates representations of the input views, $f(x + \delta^*)$, $f(x^+)$, and $f(x^-)$. The combined loss function L_{VESPR} is calculated from the outputs of the projector $g(\cdot)$ and classifier $c(\cdot)$ networks:

$$\begin{aligned} L_{VESPR}(x + \delta^*, x^+, x^-, y) &= \alpha L_{CL}[g(f(x + \delta^*)), g(f(x^+)), g(f(x^-))] \\ &+ \beta L_{CE}[c(f(x + \delta^*)), y] \end{aligned} \quad (3)$$

with inputs as view x with label y , corresponding positive view x^+ (same source image), and negative views x^- (different source image). α and β weight the balance between L_{CL} and L_{CE} .

As CE loss is integrated into L_{VESPR} , VESPR leverages class-specific features for training. Simultaneously, the contrastive loss component of L_{VESPR}

not only focuses on features that distinguish different images, but also mitigates class collapse and poison shortcut learning induced by CE loss [10, 80]. Finally, adversarial samples generated by the classification loss create perturbed augmentations that maximally obscure poison features.

4 Experiments

4.1 Experiment Setup

We evaluate the performance of VESPR and other baseline defenses on seven popular poisoning methods: Adversarial Poisons (AP8) [26], Contrastive Poisoning (CP8) [32], Convolution-based Unlearnable Datasets (CUDA) [67], Linearly Separable Poisons (LSP16) [83], One-Pixel Shortcut (OPS) [79], Robust Unlearnable Examples (RUE8) [27], and Unlearnable Examples (UE8) [40]. Further poison descriptions can be found in Supplemental Material Sec. B.1.

We primarily train models on poisoned versions of a 100-class subset of the ImageNet dataset [20], each with a total of 130,000 training images and 5,000 test images. For benchmarking purposes, we also train models on poisoned CIFAR-10 datasets [48], each with a total of 50,000 training images and 10,000 test images. Due to the lengthy optimization process required by some poisons on large images, AP8, CP8, and UE8 poisons for ImageNet-100 are generated with 20% of the dataset, containing 100 classes with a total of 26,000 training images.

To benchmark the effectiveness of VESPR, we compare against six state-of-the-art poison defense methods: adversarial training (SL+AT), Cutout [21], Mixup [87], CutMix [85], ISS [53], and SSL (SimCLR) [12].

For all models, we employ ResNet18 encoders [35] with an encoding dimension of 512. When running experiments on CIFAR-10, we follow the architecture of CIFAR ResNet models, as in [17]. For all datasets, all SSL methods utilize a 3-layer MLP projection network following the MoCo projector structure [34] with hidden and output dimensions of 2048. For classification, we attach a single linear layer to the encoder output. Adversarial examples are generated using projected gradient descent [49, 57] with 10 steps of size $0.6/255$, where perturbations are L_∞ -bounded at $\epsilon_{AT} = 4/255$. Additional training details are deferred to Supplemental Material Sec. B.

We evaluate performance of trained models by measuring classification accuracy on clean test data. For each defense method, we measure the minimum (Psn Min) and average (Psn Avg) accuracies across all poisons. An adversary will choose the most effective poison given state-of-the-art defenses, so increasing Psn Min is crucial to improving model robustness in real-world scenarios.

4.2 VESPR Performance

Table 2 presents clean test accuracy on ImageNet-100 for VESPR as well as multiple baseline defenses. VESPR debuts as a leading defense method, achieving the highest Psn Min and Psn Avg over all previous defenses by 16% and 9%,

Table 2: ImageNet-100 clean test set accuracies for VESPR and six defense benchmarks. Bold numbers indicate the best defense method for each poison, while underlined numbers indicate second-best methods.

Poison	SL	SL+AT	Cutout [21]	Mixup [87]	CutMix [85]	ISS [53]	SSL [12]	VESPR
Clean	77.66	69.76	78.04	<u>80.38</u>	81.08	71.58	70.96	74.84
AP8 [26]	8.18	<u>42.88</u>	7.68	9.68	8.38	24.22	41.80	48.68
CP8 [32]	57.46	48.76	58.20	<u>61.76</u>	62.28	50.18	14.80	54.40
CUDA [67]	6.10	<u>36.34</u>	8.66	5.16	5.70	3.58	26.12	45.20
LSP16 [83]	6.62	28.92	5.06	5.50	7.84	33.70	62.06	<u>56.34</u>
OPS [79]	51.00	52.56	53.86	48.82	65.12	57.36	62.22	<u>63.26</u>
RUE8 [27]	14.18	57.34	15.60	33.08	16.10	<u>67.52</u>	65.30	70.36
UE8 [40]	6.32	43.26	6.42	12.38	5.80	41.18	50.02	<u>49.34</u>
Psn Min	6.10	<u>28.92</u>	5.06	5.16	5.70	3.58	14.80	45.20
Psn Avg	21.41	44.29	22.21	25.20	24.46	39.68	<u>46.05</u>	55.37

Table 3: CIFAR-10 clean test set accuracies for VESPR and six defense benchmarks.

Poison	SL	SL+AT	Cutout [21]	Mixup [87]	CutMix [85]	ISS [53]	SSL [12]	VESPR
Clean	93.86	89.57	95.62	95.46	95.78	82.40	90.55	89.35
AP8 [26]	12.45	85.67	9.70	33.27	8.38	81.10	75.77	86.78
CP8 [32]	93.59	88.43	94.51	93.36	93.77	82.21	73.08	89.20
CUDA [67]	21.89	48.58	23.46	21.75	24.04	21.94	66.58	80.31
LSP16 [83]	21.47	85.67	18.47	22.76	21.68	79.33	88.87	88.96
OPS [79]	27.24	20.10	65.12	37.56	85.27	72.14	87.79	84.74
RUE8 [27]	18.91	34.45	19.67	23.27	27.46	80.04	86.70	87.82
UE8 [40]	32.53	89.31	36.18	54.19	38.40	82.12	88.45	89.39
Psn Min	18.91	20.10	9.70	21.75	8.38	21.94	66.58	80.31
Psn Avg	32.58	64.60	38.16	40.88	42.71	71.27	81.03	86.74

respectively, as well as maintaining strong performance when trained on clean data. VESPR frequently attains the highest clean test accuracy (AP8, CUDA, and RUE) or the second-highest clean test accuracy (LSP16, OPS, and UE8).

Though augmentation-based defenses occasionally achieve top performance (e.g., CutMix on CP8 and OPS), they are often overshadowed by other methods. Indeed the Psn Min values for augmentation-based methods are all less than 10%. SL+AT and SSL are consistently high performers, achieving average test accuracies of 44.29% and 46.05% across all poisons. However, their performance when trained on clean data alone is underwhelming, compared to other methods. This is unfortunate, as it suggests a smaller upper bound for clean test accuracy after poison training; intuitively, a model trained on poison data will not out-

perform the same model trained on clean data. As expected, SSL is particularly weak to CP8, a poison designed specifically to counter SSL.

Table 3 shows that the trends observed on ImageNet-100 performance are reflected in CIFAR-10. On CIFAR-10, VESPR outperforms other defense methods on multiple poisons (AP8, CUDA, LSP16, RUE8, and UE8). VESPR maintains the highest minimum and average performance across poisons by 13% and 5% respectively. For all defenses, the performance gap between clean and poison training is reduced relative to ImageNet-100, likely due to fewer classes and the ease of classifying smaller images.

4.3 Ablation Studies

Architecture Ablation In Table 4, we present the poison defense performance of all combinations for SL, SSL, and AT on ImageNet-100. Prior to this work, SSL+AT and SSL+SL methods have not been evaluated for poison defense. As measured by Psn Min and Psn Avg performance, VESPR outperforms all ablations by 13% and 4%, respectively.

Table 4: ImageNet-100 clean test set accuracies for VESPR and multiple ablations to architecture and training procedure.

Poison	SL	SL (SSL Aug)	SL+AT	SSL [12]	SSL+AT [42]	SSL+SL	VESPR
Clean	77.66	75.48	69.76	70.96	60.86	79.66	74.84
AP8 [26]	8.18	16.54	42.88	41.80	40.46	28.82	48.68
CP8 [32]	57.46	60.40	48.76	14.80	44.00	58.38	54.40
CUDA [67]	6.10	17.30	36.34	26.12	31.78	21.82	45.20
LSP16 [83]	6.62	49.30	28.92	62.06	49.00	52.36	56.34
OPS [79]	51.00	63.68	52.56	62.22	53.16	70.14	63.26
RUE8 [27]	14.18	49.08	57.34	65.30	54.30	70.02	70.36
UE8 [40]	6.32	34.46	43.26	50.02	40.94	56.68	49.34
Psn Min	6.10	16.54	28.92	14.80	31.78	21.82	45.20
Psn Avg	21.41	41.54	44.29	46.05	44.81	51.17	55.37

In order to assess the impact of augmentations, we additionally evaluate SL with SSL-style augmentations. For ease of comparison, we compile all augmentation-based defense method results for ImageNet-100 in Table 10 of the Supplemental Material Sec. C.1. As shown in Table 10, SL with SSL-style augmentations achieves the highest Psn Min and Psn Avg over the other augmentation-based methods by 10% and 2%, respectively. However, when compared to VESPR, SSL+SL, SSL, and SL+AT in Table 4, SL with SSL-style augmentations consistently underperforms. We hypothesize that existing augmentation-based defenses are weak to data poisoning due to their reliance on image cropping; most poisons

are spread throughout images such that all crops are poisoned. This hypothesis is further supported by the fact that all augmentation-based defenses are particularly resistant to OPS, the only poison which does not spread throughout images. This ablation underscores that augmentations cannot fully explain the advances observed with VESPR.

As illustrated in Table 4, integrating AT with SL enhances performance over SL alone for all poisons except CP8. Though SL+AT, which trains using bounded additive noise, is particularly effective against bounded additive poisons (e.g., AP8, UE8), improvement is relatively limited on non-additive poisons (e.g. OPS) or poisons with large perturbations (e.g., LSP16). Additionally, applying AT to SSL degrades performance relative to SSL for most poisons. This highlights the nuanced nature of AT application: as poisons aim to undermine supervised classification accuracy, the adversarial signals from SSL alone have limited capability to ignore poisonous patterns.

Aligning with expectations from prior literature, [10, 80], SSL+SL achieves top performance when trained on clean data, as the combined objectives of SSL and SL prevent the failure modes of either individual objective. Nevertheless, the intuition from prior literature does not necessarily apply for poisoned training, where SSL+SL often performs worse than other defenses, including SSL (e.g., AP8, CUDA, and LSP16). This signifies that the SL component of SSL+SL is still weak to poison patterns, dragging down the performance of SSL+SL. Furthermore, it emphasizes that the additional AT component of VESPR is crucial. For **all** poisons where SSL+SL underperforms SSL, VESPR utilizes SL-guided adversarial samples to outperform SSL+SL. This validates the claim that poisoned-SL can provide useful guidance for adversarial example generation by obfuscating poison features.

We defer the corresponding architecture ablation for CIFAR-10 to Table 11 of the Supplemental Material Sec. C.2. The trends seen for ImageNet-100 in Table 4 are largely reflected in CIFAR-10, and VESPR attains top performance for Psn Min and Psn Avg across all architecture ablations.

Architecture Ablation Analysis We expand our previous analysis on learned representations (Table 1) to all architecture ablations in Table 5. We additionally evaluate representation roughness (i.e., variability) using a local Lipschitz metric (Psn Local Lip) [82], where lower values correspond to smoother, more robust representation landscapes.

As illustrated in Table 5, VESPR maintains Psn In-Cls Sim values that are similar to those of SSL+SL and substantially higher than SSL or SSL+AT. Furthermore, utilizing AT consistently increases Psn-Cln Sim and reduces Psn Local Lip across all methods, including VESPR, compared to those without AT. In particular, the Psn Local Lip of VESPR decreases by an order of magnitude, from 8546 to 712, over SSL+SL. This improvement in smoothness is greater than that of any other architecture that incorporates AT, underscoring the superior effect of AT on VESPR. Taken together, the trends in Table 5 suggest that SSL+SL aids VESPR in establishing distinct class-wise clusters of representa-

Table 5: Columns display Psn In-Cls Sim, Psn-Cln Sim, and Psn E-Rank, as in Table 1. We also present poison image representation roughness (Psn Local Lip, see [82]). We average results across seven different poisons (see Supplemental Material Sec. B.1)

Method	Psn In-Cls Sim	Psn-Cln Sim	Psn E-Rank	Psn Local Lip
SL	0.78	0.70	16.67	6180
SL (SSL Aug)	0.79	0.79	16.43	5868
SL+AT	0.76	0.85	31.64	3312
SSL	0.47	0.86	59.61	317
SSL+AT	0.60	0.97	23.47	40
SSL+SL	0.77	0.83	34.07	8546
VESPR	0.72	0.91	30.29	712

tions while VESPR’s SL-guided AT ensures that these representations are reliant on robust, clean features. These findings justify the superior clean-classification performance of VESPR across multiple poisons.

Trends in representation similarity from Table 5 can also be identified visually. Fig 3 displays 2D t-SNE plots [38] of the 512D representations for three classes of poison and clean images and their CUDA-poisoned counterparts for various training methods. Slight mixing between encodings is expected as all three classes are types of fish, yet successful classification requires that clusters are largely distinct. Both SL and SSL tend to tightly cluster poison representations, yet corresponding clean images are encoded to distant representations. Furthermore, all clean image representations from SL and SSL are clustered together, regardless of class, suggesting that clean images are treated as an out-of-distribution class. In contrast, VESPR achieves superior poison-clean alignment, with representations of poison and clean samples overlapping for most samples. This verifies the high Psn-Cln Sim value for VESPR found in Table 5.




Fig. 3: 2-dimensional t-SNE plots [38] of clean and CUDA-poisoned image representations for SL, SSL, and VESPR models trained on CUDA data.

VESPR Adversarial Ablation In Table 6, we ablate the method of adversarial sample generation. Specifically, we compare the effect of random Gaussian noise (SSL+SL+GN) to methods of adversarial perturbation (VESPR and variants). VESPR remains the top-performing implementation.

For SSL+SL+GN, Gaussian noise with 4/255 standard deviation is applied to images with 50% probability. As the permissible Gaussian perturbation is higher than the AT perturbation bound, the domain of possible samples generated by Gaussian noise subsumes the domain of samples generated by AT. SSL+SL+GN slightly outperforms SSL+SL on most poisons, but overall performs similarly. Notably, the performance of SSL+SL+GN is conspicuously distinct from that of VESPR, suggesting that gradient-based perturbations from AT are crucial for learning robust features. We note that VESPR consistently outperforms VESPR-SSL and VESPR-Both, with notable improvements on CUDA and RUE8 in particular. VESPR also attains the highest minimum and average clean test accuracies. This suggests that VESPR better exploits the adversarial signal from SL, allowing the encoder to learn robust features with SSL.

Table 6: ImageNet-100 clean test set accuracy for different methods of sample perturbation. SSL+SL+GN is similar to SSL+SL, but with additional Gaussian noise augmentation. VESPR-SSL uses AT with SSL loss, VESPR-Both uses AT with the combined VESPR loss in Eq. 3, and VESPR uses AT with SL loss.

Poison	SSL+SL	SSL+SL+GN	VESPR-SSL	VESPR-Both	VESPR
Clean	79.66	78.86	73.74	73.96	74.84
AP8 [26]	28.82	29.92	47.56	48.02	48.68
CP8 [32]	58.38	59.14	54.20	53.64	54.40
CUDA [67]	21.82	24.02	36.64	41.26	45.20
LSP16 [83]	52.36	53.00	54.36	55.88	56.34
OPS [79]	70.14	71.12	64.30	63.14	63.26
RUE8 [27]	70.02	69.90	66.38	68.58	70.36
UE8 [40]	56.68	58.80	49.28	48.20	49.34
Psn Min	21.82	24.02	36.64	41.26	45.20
Psn Avg	51.17	52.27	53.25	54.10	55.37

5 Conclusion

We extend the investigation of poison defense to seven poisons on CIFAR-10 and ImageNet-100 datasets, revealing limited performance of current defense methods, including SSL. Drawing from insights on SL and SSL poison defense, we introduce VESPR, which leverages supervised poison vulnerability to enhance self-supervised defense. Through extensive experiments, we demonstrate that VESPR achieves state-of-the-art defense against multiple prevalent poisons, surpassing all other baselines in both minimum and average poison performance.

Acknowledgements

This research is supported by the National Research Foundation, Singapore and Infocomm Media Development Authority under its Trust Tech Funding Initiative and Strategic Capability Research Centres Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore and Infocomm Media Development Authority.

References

1. Addepalli, S., Jain, S., Babu, R.V.: Efficient and effective augmentation strategy for adversarial training (2022)
2. Adnan, M., Ioannou, Y., Tsai, C.Y., Galloway, A., Tizhoosh, H.R., Taylor, G.W.: Monitoring shortcut learning using mutual information (2022)
3. Alkhunaizi, N., Kamzolov, D., Takáč, M., Nandakumar, K.: Suppressing poisoning attacks on federated learning for medical imaging. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 673–683. Springer (2022)
4. Baier, F., Mair, S., Fadel, S.G.: Self-supervised siamese autoencoders (2023)
5. Bardes, A., Ponce, J., LeCun, Y.: Vicreg: Variance-invariance-covariance regularization for self-supervised learning (2022)
6. Carlini, N., Terzis, A.: Poisoning and backdooring contrastive learning. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=iC4UHbQ01Mp>
7. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers (2021)
8. Chan-Hon-Tong, A.: An algorithm for generating invisible data poisoning using adversarial noise that breaks image classification deep learning. *Machine Learning and Knowledge Extraction* **1**(1), 192–204 (2019), <https://www.mdpi.com/2504-4990/1/1/11>
9. Chen, K., Liu, Z., Hong, L., Xu, H., Li, Z., Yeung, D.Y.: Mixed autoencoder for self-supervised visual representation learning (2024)
10. Chen, M.F., Fu, D.Y., Narayan, A., Zhang, M., Song, Z., Fatahalian, K., Ré, C.: Perfectly balanced: Improving transfer and robustness of supervised contrastive learning (2022)
11. Chen, T., Liu, S., Chang, S., Cheng, Y., Amini, L., Wang, Z.: Adversarial robustness: From self-supervised pre-training to fine-tuning (2020)
12. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations (2020)
13. Chen, T., Luo, C., Li, L.: Intriguing properties of contrastive losses (2021)
14. Chen, X., He, K.: Exploring simple siamese representation learning (2020)
15. Chen, X., Xie, S., He, K.: An empirical study of training self-supervised vision transformers (2021)
16. Cohen, J.M., Rosenfeld, E., Kolter, J.Z.: Certified adversarial robustness via randomized smoothing (2019)
17. da Costa, V.G.T., Fini, E., Nabi, M., Sebe, N., Ricci, E.: Solo-learn: A library of self-supervised methods for visual representation learning (2022)

18. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space (2019)
19. Demontis, A., Melis, M., Pintor, M., Jagielski, M., Biggio, B., Oprea, A., Nita-Rotaru, C., Roli, F.: Why do adversarial attacks transfer? explaining transferability of evasion and poisoning attacks (2019)
20. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A Large-Scale Hierarchical Image Database. In: CVPR09 (2009)
21. DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout (2017)
22. Dippel, J., Vogler, S., Höhne, J.: Towards fine-grained visual representations by combining contrastive learning with image reconstruction and attention-weighted pooling (2022)
23. Eastwood, C., Williams, C.K.I.: A framework for the quantitative evaluation of disentangled representations. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=By-7dz-AZ>
24. Fan, L., Liu, S., Chen, P.Y., Zhang, G., Gan, C.: When does contrastive learning preserve adversarial robustness from pretraining to finetuning? (2021)
25. Feng, J., Cai, Q.Z., Zhou, Z.H.: Learning to confuse: Generating training time adversarial data with auto-encoder (2019)
26. Fowl, L., Goldblum, M., yeh Chiang, P., Geiping, J., Czaja, W., Goldstein, T.: Adversarial examples make strong poisons (2021)
27. Fu, S., He, F., Liu, Y., Shen, L., Tao, D.: Robust unlearnable examples: Protecting data privacy against adversarial learning. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=baUQQPwQiAg>
28. Geirhos, R., Narayanappa, K., Mitzkus, B., Bethge, M., Wichmann, F.A., Brendel, W.: On the surprising similarities between supervised and self-supervised models (2020)
29. Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In: International Conference on Learning Representations (2019), <https://openreview.net/forum?id=Bygh9j09KX>
30. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples (2015)
31. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B.A., Guo, Z.D., Azar, M.G., Piot, B., Kavukcuoglu, K., Munos, R., Valko, M.: Bootstrap your own latent: A new approach to self-supervised learning (2020)
32. He, H., Zha, K., Katabi, D.: Indiscriminate poisoning attacks on unsupervised contrastive learning. In: The Eleventh International Conference on Learning Representations (2023), https://openreview.net/forum?id=f0a_dWEYg-Td
33. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners (2021)
34. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning (2020)
35. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition (2015)
36. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A.: beta-VAE: Learning basic visual concepts with a constrained variational framework. In: International Conference on Learning Representations (2017), <https://openreview.net/forum?id=Sy2fzU9gl>

37. Hinton, G.E., Salakhutdinov, R.R.: Reducing the dimensionality of data with neural networks. *Science* **313**(5786), 504–507 (2006). <https://doi.org/10.1126/science.1127647>, <https://www.science.org/doi/abs/10.1126/science.1127647>
38. Hinton, G.E., Roweis, S.: Stochastic neighbor embedding. In: Becker, S., Thrun, S., Obermayer, K. (eds.) *Advances in Neural Information Processing Systems*. vol. 15. MIT Press (2002), https://proceedings.neurips.cc/paper_files/paper/2002/file/6150ccc6069bea6b5716254057a194ef-Paper.pdf
39. Hua, T., Wang, W., Xue, Z., Ren, S., Wang, Y., Zhao, H.: On feature decorrelation in self-supervised learning (2021)
40. Huang, H., Ma, X., Erfani, S.M., Bailey, J., Wang, Y.: Unlearnable examples: Making personal data unexploitable (2021)
41. Islam, A., Chen, C.F., Panda, R., Karlinsky, L., Radke, R., Feris, R.: A broad study on the transferability of visual representations with contrastive learning (2021)
42. Jiang, Z., Chen, T., Chen, T., Wang, Z.: Robust pre-training by adversarial contrastive learning (2020)
43. Jing, L., Vincent, P., LeCun, Y., Tian, Y.: Understanding dimensional collapse in contrastive self-supervised learning (2022)
44. Kahana, J., Hoshen, Y.: A contrastive objective for learning disentangled representations (2022)
45. Kim, M., Ha, H., Son, S., Hwang, S.J.: Effective targeted attacks for adversarial self-supervised learning (2023)
46. Kim, M., Tack, J., Hwang, S.J.: Adversarial self-supervised contrastive learning (2020)
47. Kingma, D.P., Welling, M.: Auto-encoding variational bayes (2022)
48. Krizhevsky, A.: Learning multiple layers of features from tiny images (2009), <https://api.semanticscholar.org/CorpusID:18268744>
49. Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial examples in the physical world (2017)
50. Li, A.C., Efros, A.A., Pathak, D.: Understanding collapse in non-contrastive siamese representation learning (2022)
51. Li, T., Fan, L., Yuan, Y., He, H., Tian, Y., Feris, R., Indyk, P., Katabi, D.: Addressing feature suppression in unsupervised visual representations (2021)
52. Liu, H., Jia, J., Gong, N.Z.: Poisonedencoder: Poisoning the unlabeled pre-training data in contrastive learning (2023), <https://arxiv.org/abs/2205.06401>
53. Liu, Z., Zhao, Z., Larson, M.: Image shortcut squeezing: Countering perturbative availability poisons with compression (2023)
54. Locatello, F., Bauer, S., Lucic, M., Rätsch, G., Gelly, S., Schölkopf, B., Bachem, O.: Challenging common assumptions in the unsupervised learning of disentangled representations (2019)
55. Luo, R., Wang, Y., Wang, Y.: Rethinking the effect of data augmentation in adversarial contrastive learning (2023)
56. Lyu, K., Li, Z., Wang, R., Arora, S.: Gradient descent on two-layer nets: Margin maximization and simplicity bias (2021)
57. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks (2019)
58. Misra, I., van der Maaten, L.: Self-supervised learning of pretext-invariant representations (2019)
59. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning. In: *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011* (2011), http://ufldl.stanford.edu/housenumbers/nips2011_housenumbers.pdf

60. van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding (2019)
61. Patel, N., Krishnamurthy, P., Garg, S., Khorrami, F.: Bait and switch: On-line training data poisoning of autonomous driving systems. arXiv preprint arXiv:2011.04065 (2020)
62. Qin, T., Gao, X., Zhao, J., Ye, K., Xu, C.Z.: Apbench: A unified benchmark for availability poisoning attacks and defenses (2023)
63. Qin, T., Gao, X., Zhao, J., Ye, K., Xu, C.Z.: Learning the unlearnable: Adversarial augmentations suppress unlearnable example attacks (2023)
64. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision (2021)
65. Robinson, J., Sun, L., Yu, K., Batmanghelich, K., Jegelka, S., Sra, S.: Can contrastive learning avoid shortcut solutions? (2021)
66. Rubinstein, B.I., Nelson, B., Huang, L., Joseph, A.D., Lau, S.h., Rao, S., Taft, N., Tygar, J.D.: Antidote: understanding and defending against poisoning of anomaly detectors. In: Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement. pp. 1–14 (2009)
67. Sadasivan, V.S., Soltanolkotabi, M., Feizi, S.: Cuda: Convolution-based unlearnable datasets (2023)
68. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks (2014)
69. Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., Isola, P.: What makes for good views for contrastive learning? (2020)
70. Tian, Y.: Understanding deep contrastive learning via coordinate-wise optimization (2022)
71. Tian, Y., Chen, X., Ganguli, S.: Understanding self-supervised learning dynamics without contrastive pairs (2021)
72. Tishby, N., Pereira, F.C., Bialek, W.: The information bottleneck method (2000)
73. Tsai, Y.H.H., Ma, M.Q., Yang, M., Zhao, H., Morency, L.P., Salakhutdinov, R.: Self-supervised representation learning with relative predictive coding (2021)
74. Uesato, J., Alayrac, J.B., Huang, P.S., Stanforth, R., Fawzi, A., Kohli, P.: Are labels required for improving adversarial robustness? (2019)
75. Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: Proceedings of the 25th International Conference on Machine Learning. pp. 1096–1103 (01 2008). <https://doi.org/10.1145/1390156.1390294>
76. Wang, T., Yue, Z., Huang, J., Sun, Q., Zhang, H.: Self-supervised learning disentangled group representation as feature (2021)
77. Wang, T., Isola, P.: Understanding contrastive representation learning through alignment and uniformity on the hypersphere (2022)
78. Wang, X., Chen, X., Du, S.S., Tian, Y.: Towards demystifying representation learning with non-contrastive self-supervision (2022)
79. Wu, S., Chen, S., Xie, C., Huang, X.: One-pixel shortcut: on the learning preference of deep neural networks (2023)
80. Xue, Y., Joshi, S., Gan, E., Chen, P.Y., Mirzasoleiman, B.: Which features are learnt by contrastive learning? on the role of simplicity bias in class collapse and feature suppression (2023)
81. Yang, W., Kirichenko, P., Goldblum, M., Wilson, A.G.: Chroma-vae: Mitigating shortcut learning with generative classifiers (2022)

82. Yang, Y.Y., Rashtchian, C., Zhang, H., Salakhutdinov, R., Chaudhuri, K.: A closer look at accuracy vs. robustness (2020)
83. Yu, D., Zhang, H., Chen, W., Yin, J., Liu, T.Y.: Availability attacks create shortcuts. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. ACM (aug 2022). <https://doi.org/10.1145/3534678.3539241>
84. Yuan, C.H., Wu, S.H.: Neural tangent generalization attacks. In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 12230–12240. PMLR (18–24 Jul 2021), <https://proceedings.mlr.press/v139/yuan21b.html>
85. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features (2019)
86. Zhang, C., Zhang, K., Zhang, C., Pham, T.X., Yoo, C.D., Kweon, I.S.: How does simsiam avoid collapse without negative samples? a unified understanding with self-supervised contrastive learning (2022)
87. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization (2018)
88. Zhang, Y., Tan, Z., Yang, J., Huang, W., Yuan, Y.: Matrix information theory for self-supervised learning (2023)
89. Zhou, J., Dong, L., Gan, Z., Wang, L., Wei, F.: Non-contrastive learning meets language-image pre-training (2022)
90. Zhuo, Z., Wang, Y., Ma, J., Wang, Y.: Towards a unified theoretical understanding of non-contrastive learning via rank differential mechanism (2023)