

Sparse Visual Counterfactual Explanations in Image Space

Valentyn Boreiko, Maximilian Augustin, Francesco Croce, Philipp Berens, and
Matthias Hein

University of Tübingen

Abstract. Visual counterfactual explanations (VCEs) in image space are an important tool to understand decisions of image classifiers as they show under which changes of the image the decision of the classifier would change. Their generation in image space is challenging and requires robust models due to the problem of adversarial examples. Existing techniques to generate VCEs in image space suffer from spurious changes in the background. Our novel perturbation model for VCEs together with its efficient optimization via our novel Auto-Frank-Wolfe scheme yields sparse VCEs which lead to subtle changes specific for the target class. Moreover, we show that VCEs can be used to detect undesired behavior of ImageNet classifiers due to spurious features in the ImageNet dataset. Code is available under https://github.com/valentyn1boreiko/SVCEs_code.

Keywords: Interpretability · Adversarial robustness · Trustworthy AI

1 Introduction

The black-box nature of decisions made by neural networks is one of the main obstacles for the widespread use of machine learning in industry and science. It is likely that future regulatory steps will strengthen the “right for an explanation”, which is currently already implemented in a weak form in the GDPR [63] and is included as “transparency of an AI system” in a draft for regulating AI of the European Union, at least concerning the use of AI in safety critical systems [12]. Apart from directly interpretable classifiers like linear models or decision trees, a variety of model-agnostic explanation techniques has been proposed: sensitivity based explanations [4], explanations based on feature attributions [3], saliency maps [55,54,21,64,57], Shapley additive explanations [38], and local fits of interpretable models [48], see [40] for a recent overview.

Another candidate are counterfactual explanations (CEs) introduced in [63] as a form of instance-specific explanations close to human reasoning [41]. Humans often justify decisions by counterfactual reasoning: “I would have decided for X , if Y had been true”. One of their biggest advantages in contrast to feature attribution and other mentioned above methods is that CEs [63,18,43,5,45,62,52] are actionable explanations [63] and thus are close to what the GDPR requires. Current approaches to generate CEs for classifier decisions can be summarized




Fig. 1: **VCEs for two different threat models.** VCEs together with difference maps for the change “cougar \rightarrow cheetah” for an adversarially robust ImageNet model [20,16]. Our novel $l_{1.5}$ -VCEs yield more sparse changes which are mainly focused on the object compared to the previously considered l_2 -VCEs [51,1].

by answering the question: “What is the minimal change δ of the input x , so that the perturbed input $x + \delta$ is classified as the desired target class with sufficiently high confidence and is realistic?”. From the developer’s perspective, counterfactuals are interesting for debugging as they allow to detect spurious features which the classifier has picked up. We refer to [62] for a recent extensive overview on the literature of counterfactual explanations who note five criteria for CEs: i) **validity**: the changed input $x + \delta$ should have the desired target class, ii) **actionability**: the change δ should be possible to be realized by the human, iii) **sparsity**: the change δ should be sparse so that the change is interpretable for humans, iv) **realism**: the changed input $x + \delta$ should lie close to the data manifold, v) **causality**: CEs should maintain causal relations between features. Interestingly, [62] noted that most papers they reviewed just evaluate on tabular datasets or MNIST. The reason for this is that the process of generating CEs for high-dimensional image datasets (which we will refer to as visual counterfactual explanations, or VCEs for short) is very similar to that of generating adversarial examples [58] which just exploit non-robust features of the classifier and thus show no class-specific changes required for VCEs. Thus, realistic VCEs require either (adversarially) robust models as in [51,1] or that the images are implicitly restricted via the usage of a generative model [28,29,49,10,25,53]. Very recently visual counterfactuals based on generative models have been proposed [37,25,50] but no code has been released so far or it is restricted to MNIST. These methods require to specify the amount of “classifier guidance” which might be difficult to be choose as we discuss in App. F. For this reason, in this work we investigate the generation of VCEs directly in image space, instead of working in the latent space, and purely based on the classifier, thus showing its behavior without the influence of an auxiliary model. We make the following contributions: i) we show that the l_2 -metric used for the generation of VCEs in [51,1] leads to changes all over the image (see Fig. 1) which are unrelated to the object. This is in particular true for ImageNet models; ii) we propose a new model for sparse VCEs based on the l_p -metric for $p = 1.5$. Since an efficient projection onto $l_{1.5}$ -balls is not available, we develop a novel Auto-Frank-Wolfe (AFW) optimization scheme

with an adaptive step-size for the generation of $l_{1.5}$ -VCEs. The resulting VCEs are more sparse and “subtle” as confirmed by a user study; iii) we illustrate that VCEs are useful to detect spurious features in ImageNet classifiers, e.g., we detect the spurious feature “watermark” in the class granny smith due to a bias in the training set and show that our findings transfer to other ImageNet classifiers. This shows the utility of VCEs as a “debugging tool” for ML classifiers.

2 Visual Counterfactual Explanations (VCEs)

In this section, we first discuss the previously considered formulation of Visual Counterfactual Explanations (VCEs) of [1] in the image space and the required kind of (adversarial) robustness of the classifier. Then we discuss a novel perturbation model which overcomes the partially non-object-related changes of the VCEs proposed in [1]. For the optimization over this perturbation model we provide in Sec. 3 a novel adaptive Frank-Wolfe scheme.

We assume in the paper that the classifier, $f : \mathbb{R}^d \rightarrow \Delta_K$, where $\Delta_K := \{w \in \mathbb{R}_{\geq 0}^K \mid \sum_{i=1}^K w_i = 1\}$ is the probability simplex, outputs for every input x a probability distribution $\hat{p}_f(y|x)$ ($y \in \{1, \dots, K\}$) over the classes. The l_p -distance on \mathbb{R}^d is defined as: $\|x - y\|_p = (\sum_{i=1}^d |x_i - y_i|^p)^{\frac{1}{p}}$.

2.1 Formulation and properties of VCEs

Counterfactual explanations for a given classifier are instance-wise explanations. Informally speaking, a visual counterfactual explanation for an input image x_0 is a new image \hat{x} which is visually similar and as close as possible to a real image, but class-specific features have been changed such that the classifier now assigns to \hat{x} a desired target class different from than one assigned to x_0 (counterfactual). In addition, it is often interesting which features appear if one aims to make the classifier maximally confident in its decision (same as for x_0).

VCEs via constrained optimization: In [63] (see also [43,62]) they suggest to determine counterfactuals by the following optimization problem:

$$\hat{x} = \arg \min_{x \in \mathbb{R}^d} L(k, f(x)) + \lambda d(x_0, x), \quad (1)$$

where L is a loss function, e.g. cross-entropy loss, $L(k, f(x)) = -\log \hat{p}_f(k|x)$, k is the desired target class and $d : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ a distance, measuring similarity of x_0 and \hat{x} . If the decision of the classifier for \hat{x} changes to the target class k , then the counterfactual is “valid”. The advantage of valid counterfactuals, compared to feature attribution methods or other instance-wise explanation techniques, is that the change $\hat{x} - x_0$ is actionable, in the sense that the user understands, how to influence and change the decision of the classifier. As λ has no direct interpretation, we employ the related and more interpretable objective of [1]

$$\arg \max_{x \in [0,1]^d \cap B(x_0, \epsilon)} \log \hat{p}_f(k|x), \quad (2)$$

where $B(x_0, \epsilon) = \{x \in \mathbb{R}^d \mid d(x, x_0) \leq \epsilon\}$. The constraint, $x \in [0, 1]^d$, is necessary as we want to generate valid images. The choice of the distance metric is crucial for the quality of the VCEs (see Sec. 2.3). The new free parameter ϵ can be interpreted as “perturbation budget” with respect to the chosen metric.



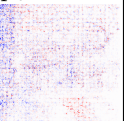


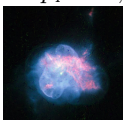
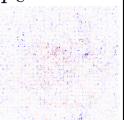

	Original	Non-robust	Madry[19]	Madry[19] + FT
ImageNet	→megalith	$p_i:0.00, p_e:1.00$	$p_i:0.00, p_e:1.00$	$p_i:0.00, p_e:1.00$
				
ImageNet-O	→jellyfish	$p_i:0.47, p_e:1.00$	$p_i:0.75, p_e:1.00$	$p_i:0.66, p_e:1.00$
				

Fig. 2: **Dependence of $l_{1.5}$ -VCEs on robustness of the model (for in- and out-of-distribution).** VCEs for a non-robust ResNet50, a robust ResNet trained with l_2 -Adversarial Training [19] and the same model finetuned for multiple norm robustness (Madry[19]+FT[16]). For the non-robust model, VCEs achieve high confidence without introducing any class-specific features or meaningful changes. For the robust models, VCEs only achieve high confidence by adding meaningful changes for the given target class (indicated by → in the original image). The multiple-norm robust model yields higher quality images than the original l_2 -robust model.

VCEs and Robustness: It has been noted in [63] that counterfactuals generated via Eq. (1) are equivalent to targeted adversarial examples. In [63] this did not cause problems as they only handled very low-dimensional problems. In fact, adversarial attacks often maximize a surrogate loss, in this case the log-probability, to induce misclassification into the target class. However, adversarial attacks [58,39] on non-robust image classifiers typically show no class-specific changes, see Fig. 2. The standard method to increase robustness to adversarial attacks is adversarial training [39] based on projected gradient descent (PGD). Notably, [61,51] have observed that adversarially robust models have strong generative properties, which is closely related to the explainability of the classifier decisions. In Fig. 2 we show the VCEs for a robust ResNet50[19] trained with adversarial training (Madry [19]) and the same model enhanced with multiple-norm finetuning [16] (Madry[19]+FT) and a non-robust model. The examples confirm that for meaningful VCE generation, a robust model is needed. Throughout the rest of the paper, we show VCEs for the robust Madry[19]+FT model. In the Appendix, we furthermore explore what kind of robustness is required for the VCE generation and how VCEs differ between robust models on both ImageNet and CIFAR10.

Table 1: **ImageNet:** Accuracy and $l_{1.5}$ -, l_2 -robust accuracy (RA) at $\epsilon_{1.5} = 12.5, \epsilon_2 = 2$ for the l_2 -adv. robust model of Madry[19] and the fine-tuned Madry[19]+FT for multiple-norm robustness [16], and FID scores for $l_1, l_{1.5}$ - and l_2 -VCEs generated on in(ID)- and out-distribution(OD) images and their average. The Madry[19]+FT model achieves the best FID score for $l_{1.5}$ -VCEs.

	Accuracies			FID scores (ID/OD/AVG)		
	Acc.	l_2 -RA	$l_{1.5}$ -RA	l_1 -VCE, $\epsilon = 400$	$l_{1.5}$ -VCE, $\epsilon = 50$	l_2 -VCE, $\epsilon = 12$
Madry[19]	57.9	45.7	37.4	13.6/41.6/27.6	8.4/24.3/16.4	8.4/22.8/15.6
[19] +FT	57.5	44.6	40.1	9.6/35.7/22.6	6.9/22.6/14.8	7.9/23.1/15.5

Properties of VCEs: Following [43,62], we aim to achieve the following main properties for our VCEs: i) **validity:** from Eq. (2), one sees that, for a given perturbation budget, we find the VCE with maximum probability $\hat{p}_f(k|\hat{x})$ in the target class k for f ; ii) **sparsity:** \hat{x} should be visually similar to x_0 and only contain sparse changes which is exactly the reason for our considered distance metric, see Sec. 2.3; iii) **realism:** \hat{x} should lie on the data manifold and look like a real image. For qualitative results we show examples of VCEs and for quantitative analysis we use the Frechet Inception Distance (FID) [32] both on VCEs generated from in-distribution test set, and from out-distribution samples, see Sec. 2.2, and a user-study in Sec. 2.3.

We stress that our primary goal is to explain the inner workings of a given classifier and not necessarily to generate the best looking images. We demonstrate in Sec. 4 that our VCEs can be successfully used to reveal undesired behavior of ImageNet classifiers due to biases in the ImageNet dataset.

2.2 Generation and evaluation of VCEs

We generate VCEs by approximately solving the non-convex problem Eq. (2) with a small computational budget. We thus use the efficient APGD [14] (only available for l_1, l_2 and l_∞) or our adaptive Frank-Wolfe scheme AFW (see Section 3). For both we use a budget of 5 random restarts each with 75 iterations. Typical deep learning classifiers are not calibrated, that is their decisions are either over- or underconfident [26]. We calibrate them using temperature rescaling by minimizing the expected calibration error (ECE) on a holdout validation set, so that confidence values are comparable, see App. D.1.

For the quantitative evaluation of the image quality of VCEs produced by different methods and classifiers, we use several metrics. First, FID scores [32] by generating 10.000 VCEs from the test set for the in-distribution (ID) evaluation where the target class is the second most likely class computed by using an ensemble of all classifiers used in the comparison, see Fig. 15 (top) in App. B. An evaluation using FID scores on the in-distribution test (FID ID) set only is in our setting problematic, as methods with no (or minimal) change would get the best FID-score. Thus, we also use VCEs generated from out-of-distribution images (ImageNet-A and ImageNet-O [31]) where the target label corresponds







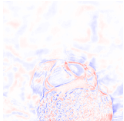
	Original	$l_1, \epsilon = 400$	Diff. map	$l_{1.5}, \epsilon = 50$	Diff. map	$l_2, \epsilon = 12$	Diff. map
ILSVRC2012	tiger beetle:0.85	dung beetle:0.88		dung beetle:0.99		dung beetle:1.00	
							

Fig. 3: **ImageNet: l_p -VCEs for $p \in \{1, 1.5, 2\}$.** l_p -VCEs into correct class for a misclassified Image of class “dung beetle”. for the multiple-norm adversarially robust model Madry[19]+FT. l_1 -VCEs are too sparse and introduce artefacts and l_2 -VCEs change the background. Our $l_{1.5}$ -VCEs are sparse and object-related (see difference maps right from the respective l_p -VCEs).

to the decision of an ensemble with all classifiers used in the comparison, see Fig. 15 (bottom) in App. B. As out-of-distribution images are not part of the data distribution, non-trivial changes are required to turn them into images of the in-distribution. Thus methods with almost no change will suffer here from large FID scores as the images are far from the in-distribution. In our experience from the qualitative inspection of the images, the average (AVG) of FID-scores on in-distribution (FID ID) and out-of-distribution images (FID OD) reflects best the realism and quality of the VCEs. Note that our FID scores cannot be directly compared to the ones of generative models as VCEs are not based on sampling. We just use the FID scores as a quantitative way to compare the different classifiers and perturbation models for VCEs. Moreover, we evaluate the utility of l_p -VCEs in a user study in Sec. 2.3.

2.3 Sparse VCEs via the $l_{1.5}$ -metric

The perturbation budget of VCEs in Eq. (2), in particular the chosen distance metric, is crucial for the generation of realistic VCEs. It might seem natural to use for l_2 -adversarially robust models also the l_2 -metric for the perturbation budget of the VCEs. However, as we show in Fig. 1 and Fig. 3, the problem of the l_2 -budget is that one typically gets non-sparse changes over the full image which are not centered on the object. Aiming at sparse VCEs it seems like the l_1 -metric might be well-suited as it is known to lead to sparse changes. However, as one can see in Fig. 3, the changes are in fact extremely sparse and often show color artefacts: *e.g.* for the dung beetle, single pixels are changed to non-natural extreme colors. As a compromise between l_1 (too sparse) and l_2 (non-sparse), we propose to use the $l_{1.5}$ -metric for the perturbation model in Eq. (2). In Fig. 1 and Fig. 3, for ImageNet the changes of $l_{1.5}$ -VCEs are sparse and localized on the object. For the generation of the FID scores for $l_{1.5}$ -VCEs, we used $\epsilon = 50$ for ImageNet. Apart from the better FID-scores of $l_{1.5}$ -VCEs we quantify in Tab. 16 in App. B for ImageNet the concentration of the changes on the actual object using the pixel-wise segmentations of ImageNet test images in [22].


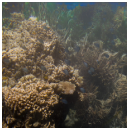
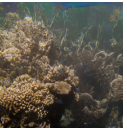



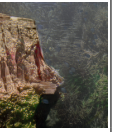


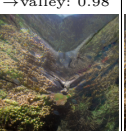






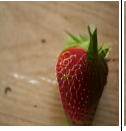









Original	$l_{1.5}$ -VCE			$l_{1.5}$ -VCE		
	$\epsilon = 50$	$\epsilon = 75$	$\epsilon = 100$	$\epsilon = 50$	$\epsilon = 75$	$\epsilon = 100$
Gila monster: 0.12 	→reef: 0.92	→reef: 0.98	→reef: 1.00	→cliff: 0.92	→cliff: 0.96	→cliff: 0.99
						
	→valley: 0.79	→valley: 0.94	→valley: 0.98	→volcano: 0.91	→volcano: 1.00	→volcano: 1.00
						
fig: 0.27 	→ strawberry: 0.97	→ strawberry: 1.00	→ strawberry: 1.00	→ pineapple: 0.99	→ pineapple: 1.00	→ pineapple: 1.00
						
	→ custard apple: 0.95	→ custard apple: 0.99	→ custard apple: 1.00	→ pomegran.: 0.76	→ pomegran.: 0.97	→ pomegran.: 0.99
						

Fig. 4: **ImageNet: $l_{1.5}$ -VCEs.** (First two rows) $l_{1.5}$ -VCEs for Madry[19]+FT with varying radii for a misclassified image of class “coral reef” for the target classes: “coral reef”, “cliff”, “valley” and “volcano” (same wordnet category “geological formation”). (Second two rows) $l_{1.5}$ -VCEs for Madry[19]+FT with varying radii for an image of class “fig” for the target classes: “strawberry”, “pineapple”, “custard apple”, and ‘pomegranate’ (same wordnet category “edible fruit”). The changes are sparse and subtle for radius 50 and get more pronounced for larger budgets.

We found the chosen radii to work well for most images. However, for visualizing the VCEs to a user, the best option is to let the user see how the changes evolve as one changes the radius in an interactive fashion. Rather subtle changes with a small budget are already sufficient for some images, whereas for other images larger budgets are necessary due to a significant change of color or shape. As such an interactive process cannot be shown, we provide panels with different radii of the perturbation model in Fig. 4 and the Appendix.

FID evaluation: For a quantitative evaluation, we compute FID scores for our ImageNet models in Tab. 1, where we use for the ImageNet validation set the second predicted class as target (ID) and for out-of-distribution images (OD) from ImageNet-O/ImageNet-A we generate the VCE for the predicted class. The FID scores indicate that $l_{1.5}$ -VCEs have higher **realism** and **sparsity** than l_1 - and l_2 -VCEs, on both in- and out-of-distribution images.

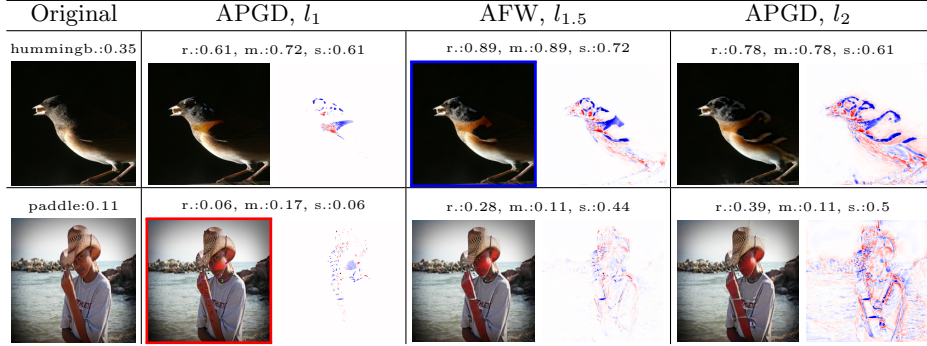


Fig. 5: **Best (blue) and worst (red) rated images from the user study.** l_p -VCEs for $p \in \{1, 1.5, 2\}$ for the change “hummingbird \rightarrow brambling” (top row) and “paddle \rightarrow bearskin” (bottom row) for Madry[19]+FT, with **realism** (r), **meaningful** (m), **subtle** (s) fractions from the user study (r:0.61 means that 61% of the users considered this image to be realistic).

User study: We perform a user study (18 participants) to compare $l_{1.5}$ -, l_1 -, and l_2 -VCEs (the Madry [19]+FT model on ImageNet is used). For each target image (94 in total), we show l_p -VCEs for $p \in \{1, 1.5, 2\}$, to the users, who can choose which ones satisfy the following properties (none or multiple answers are allowed): i) **realism**, ii) “meaningful features in the target class are introduced” (**meaningful**), iii) “subtle, yet understandable changes are introduced” (**subtle**). The percentages for l_1 -, $l_{1.5}$ -, and l_2 -VCEs are: **realism** - 23.5%, **38.2%**, 33.8%; **meaningful** - 37.5%, 63.1%, **64.0%**; **subtle** - 34.7%, **49.1%**, 41.5%. While the difference of $l_{1.5}$ -VCEs compared to l_2 -VCEs is small for meaningfulness, $l_{1.5}$ -VCEs are considered more subtle and realistic. In Fig. 5, we show the best and worst rated images from the user study. Note that for the worst one, the changes into the target class “bearskin” are not achievable in the given budget and thus all methods fail to produce meaningful images.

Details about the user study: Participants are researchers in machine learning (volunteers) not working on VCEs themselves and neither being exposed to the generated images or compared methods before. The p-values, using the two-sample two-sided binomial test, for the best and next best methods are: 0.02 for **realism** of $l_{1.5}$ vs l_2 , 0.6 for **meaningful** of l_2 vs $l_{1.5}$, and $5.3 \cdot 10^{-5}$ for **subtle** of $l_{1.5}$ vs l_2 , that is $l_{1.5}$ outperforms statistically significantly l_2 in realism/subtle (significance level 0.05).

3 Auto-Frank-Wolfe for l_p -VCEs

For deep models, the optimization problem for l_p -VCEs

$$\max_{x \in B_p(x_0, \epsilon) \cap [0, 1]^d} \log \hat{p}(y|x), \quad (3)$$




Fig. 6: **Comparison of our adaptive Frank-Wolfe scheme (AFW) vs non-adaptive FW.** We plot avg. log probability of the $l_{1.5}$ -VCEs (ImageNet: $\epsilon = 50$, CIFAR10: $\epsilon = 6$) with the second predicted class as target for ImageNet (left) and CIFAR10 (right). Our novel AFW (blue curve) is adaptive to the budget, without any need for additional hyperparameter tuning. We show the result of AFW for a budget of 25, 75 and 125. We observe that it matches the best or is better than FW with fixed learning rate (red) or various decaying stepsize schemes (green) for the given number of iterations for both datasets.

is non-convex and related to targeted adversarial attacks, for which AutoPGD (APGD) [14] has been shown to be very effective. APGD requires projections onto l_p -balls which are available either in closed form for l_2 and l_∞ or can be computed efficiently for l_1 . However, for $p \notin \{1, 2, \infty\}$, there is no such projection available and one cannot use APGD. Thus, in order to generate l_p -VCEs for $p > 1$, we propose an adaptive version of the Frank-Wolfe (FW) algorithm [42, 34], which we call Auto-Frank-Wolfe (AFW). FW has the advantage that it is projection-free and thus allows to use more complex constraint sets. In particular, we can use arbitrary l_p norm balls for $p > 1$.

Auto-Frank-Wolfe: At each iteration k , FW maximizes the first-order Taylor expansion at the iterate x^k of the objective in the feasible set, i.e.

$$s^k = \arg \max_{s \in B_p(x_0, \epsilon) \cap [0, 1]^d} \langle s, \nabla_{x^k} \log \hat{p}(y|x^k) \rangle, \quad (4)$$

and the next iterate is the convex combination

$$x^{k+1} = (1 - \gamma^k)x^k + \gamma^k s^k. \quad (5)$$

The choice of the learning rate $\gamma^k \in (0, 1)$ is crucial for the success of the algorithm: in the context of adversarial attacks, [11] use a fixed value γ_0 for every k , while [66, 60] decrease it as $\frac{\gamma_0}{\gamma_0 + k}$. In both cases the schedule is agnostic of the total budget of iterations, and γ_0 needs to be tuned. Thus, we propose to use an adaptive scheme for choosing γ^k at each iteration as $\gamma^k = \frac{M}{2 + \sqrt{k}}$ where $M \leq 2$ is adapted during the optimization. This yields our AFW attack which automatically adapts to different budgets (details of AFW in App. C.1).

Considering box-constraints: Prior FW-based attacks [11,60] do not consider the image domain constraints $[0, 1]^d$ but rather solve Eq. (4) for l_p -ball constraints only (which has a closed form solution) and clip it to $[0, 1]^d$. This is suboptimal, especially when p is close to 1, see [15]. The following proposition shows that it is possible to solve Eq. (4) efficiently in the intersection $B_p(x_0, \epsilon) \cap [0, 1]^d$ for $p > 1$ (proof in App. C.2, $p = 1$ is more simple, see [15]).

Proposition 1. *Let $w \in \mathbb{R}^d$, $x \in [0, 1]^d$, $\epsilon > 0$ and $p > 1$. The solution δ^* of the optimization problem*

$$\arg \max_{\delta \in \mathbb{R}^d} \langle w, \delta \rangle \quad \text{s.th.} \quad \|\delta\|_p \leq \epsilon, \quad x + \delta \in [0, 1]^d \quad (6)$$

is given, with the convention $\text{sign } 0 = 0$, by

$$\delta_i^* = \min \left\{ \gamma_i, \left(\frac{|w_i|}{p\mu^*} \right)^{\frac{1}{p-1}} \right\} \text{sign } w_i, \quad i = 1, \dots, d,$$

where $\gamma_i = \max\{-x_i \text{sign } w_i, (1 - x_i) \text{sign } w_i\}$ and $\mu^* > 0$ can be computed in $O(d \log d)$ time.

Experiments: To evaluate the effectiveness of AFW, we compare its performance when optimizing Eq. (3) in the $l_{1.5}$ -ball of radius 50. We use different budgets of 25, 75, 125 iterations (75 is used for generating all VCEs), and test a variety of fixed parameters for the FW attacks of [11] ($\gamma_0 \in \{0.1, 0.2, \dots, 0.9\}$ as constant stepsize and $\gamma_0 \in \{1, 5, \dots, 25, 50, 75, 100\}$ and $M \in \{2, \dots\}$ for stepsize decaying with k). Figure 6 shows that AFW achieves the maximal objective (log probability of the target class) for 75 and is second best for the budget of 25 and 125 iterations. Thus AFW adapts to the given budget and outperforms FW with fixed stepsize schemes. In App. C.3, we provide additional experiments.




Fig. 7: **Spurious feature: watermark.** The $l_{1.5}$ -VCE with target class “granny smith” (GS) for Madry[19]+FT shows that the model has associated a spurious “text” feature with this class. This is likely due to “iStockphoto” watermarked images in its training set (right). Adding the watermark changes the decision of the classifier to GS.

4 Finding spurious features with $l_{1.5}$ -VCEs

Neural networks are excellent feature extractors and very good at finding correlations in the data. This makes them susceptible to learn spurious features which are not task-related [33,7,9,23,67], to the extent that the actual object is not recognized, *e.g.* a cow on a beach is not recognized without the appearance of the spurious feature “pasture/grass” [6,56]. We show how VCEs can be used for finding such spurious features. While an automatic pipeline is beyond the scope of this paper, we believe that this can be done with minimal human supervision.

Failure A, Watermark text as spurious feature for “granny smith”: We detected this failure when creating VCEs for the target class “granny smith”. We consistently observed text-like features on the generated $l_{1.5}$ -VCEs which are obviously not related to this class. In Fig. 7 we illustrate the $l_{1.5}$ -VCE for an image from the class “bell pepper”. More examples are in App. E.1. Since almost none of the “granny smith” training images contains text, we came up with the hypothesis that the reason is a high percentage of watermarked images in the training set. Manual inspection showed that 90 out of 1300 training images contain a watermark, out of which 53 contain the one of “iStockphoto” (see the rightmost image in Fig. 7). While watermarked images appear in several classes in ImageNet, this significant fraction of one type of watermark in the training set seems to only be present in the class “granny smith”. We tested this hypothesis by simulating the watermark of “iStockphoto” on the test set, for an example see the second image from the right in Fig. 7 and App. E.1 for more details. In Tab. 2 we show that adding the “iStockphoto”-watermark has a negative influence on top-1 accuracy not only for the adv. robust model Madry[19]+FT but also other non-robust models. The reason is that this particular watermark contains features (lines) which are characteristic for classes such as bow, safety pin, reel. However, even though the “iStockphoto”-watermark contains no characteristic features of “granny smith”, adding the watermark leads to significantly worse precision of the granny smith class (basically an increase in false positives while false negatives stay the same). Interestingly, even an accurate model such as NS-B7 [65] shows this effect although trained using the much larger non-public JFT-300M dataset, suggesting that JFT-300M contains these artefacts as well.

Failure B, Cages as spurious feature for “white shark”: The next failure was detected using $l_{1.5}$ -VCEs for the shark classes where very frequently grid-like structures appear - but only for VCEs with target class “white shark” not for “tiger shark” or “hammerhead”. A typical situation is shown in Fig. 8, where the original image is from class “coral reef”. The VCE for “tiger shark” shows a shark coming from the left whereas the VCE for “white shark” shows just a grid-like structure, see App. E.2 for more such VCEs. An inspection of the “white shark” training set reveals that many of the images contain parts of cages protecting the photographing diver, see the rightmost image in Fig. 8 for an illustration. The model has picked up on this feature as a fairly dominant one for white sharks, which is clearly spurious and undesirable. Interestingly, the VCEs allow us to find such artefacts even without images of white-sharks which is an advantage over saliency or feature attribution methods.

Table 2: We show top-1 accuracy for the test set, and precision/recall for the class “granny smith” (GS) on the test set before and after adding the “iStockphoto”-watermark. Note that the watermark has a quite significant impact on accuracy and in particular on the precision of GS (more false positives), which confirms our hypothesis on the bias induced by the GS training set.

Model	Original			w. Watermark		
	Top-1	GS vs Rest		Top-1	GS vs Rest	
	Acc.	Prec.	Rec.	Acc.	Prec.	Rec.
Madry[19]+FT	57.5	61.1	73.3	50.4	43.8	70.0
ResNet50[19]	76.0	90.3	93.3	62.3	53.2	83.3
NS-Eff. B7 [65]	86.6	90.3	93.3	84.1	68.3	93.3




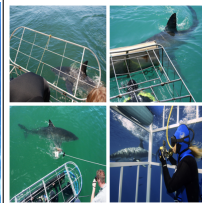
Orig.	$l_{1.5}$ -VCE, $\epsilon = 100$		Train set
coral reef: 0.58	\rightarrow t. shark: 0.96	\rightarrow w. shark: 0.99	white shark
			

Fig. 8: **Spurious feature: cage.** The $l_{1.5}$ -VCE for an image from class coral reef with target “tiger shark” shows a tiger shark, but with target “white shark” grid-like structures as spurious feature. The training set of white shark (right) contains many images with cages.

5 Discussion and Limitations

We have shown that our $l_{1.5}$ -VCEs are sparse and do subtle changes located on the object of interest even for ImageNet resolution. We have shown that $l_{1.5}$ -VCEs are a useful debugging tool for detecting spurious features which the classifier has picked up. Not all VCEs are meaningful which can have different reasons: i) the perturbation budget is too small to be able to change to more distinct classes, ii) the model has not picked up the right features, or iii) VCEs show spurious features, as discussed in Sec. 4. However, these “limitations” just reflect that our classifiers are not yet perfect and are not a failure of VCEs. In the future it will be interesting to generate an automatic pipeline for the detection of spurious features with minimal human supervision.

Acknowledgement

M.H., P.B., and V.B. acknowledge support by the the DFG Excellence Cluster Machine Learning - New Perspectives for Science, EXC 2064/1, Project number 390727645.

References

1. Augustin, M., Meinke, A., Hein, M.: Adversarial robustness on in- and out-distribution improves explainability. In: ECCV (2020) [2](#), [3](#), [24](#)
2. Avrahami, O., Lischinski, D., Fried, O.: Blended diffusion for text-driven editing of natural images (2021) [17](#), [37](#), [38](#), [39](#), [40](#)
3. Bach, S., Binder, A., Gregoire Montavon, F.K., Müller, K.R., Samek, W.: On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. PLoS One **10**(7), e0130140 (2015) [1](#)
4. Baehrens, D., Schroeter, T., Harmeling, S., Kawanabe, M., Hansen, K., Müller, K.R.: How to explain individual classification decisions. Journal of Machine Learning Research (JMLR) **11**, 1803–1831 (2010) [1](#)
5. Barocas, S., Selbst, A.D., Raghavan, M.: The hidden assumptions behind counterfactual explanations and principal reasons. In: FAccT. p. 80–89 (2020) [1](#)
6. Beery, S., van Horn, G., Perona, P.: Recognition in terra incognita. In: ECCV (2018) [11](#)
7. Brendel, W., Bethge, M.: Approximating cnns with bag-of-local-features models works surprisingly well on imagenet. In: ICLR (2019) [11](#)
8. Carmon, Y., Ragunathan, A., Schmidt, L., Duchi, J.C., Liang, P.: Unlabeled data improves adversarial robustness. In: NeurIPS (2019) [24](#)
9. Carter, S., Armstrong, Z., Schubert, L., Johnson, I., Olah, C.: Exploring neural networks with activation atlases. Distill. (2019) [11](#)
10. Chang, C.H., Creager, E., Goldenberg, A., Duvenaud, D.: Explaining image classifiers by counterfactual generation. In: ICLR (2019) [2](#)
11. Chen, J., Yi, J., Gu, Q.: A Frank-Wolfe framework for efficient and effective adversarial attacks. In: AAAI (2019) [9](#), [10](#)
12. Commission, E.: Regulation for laying down harmonised rules on AI. European Commission (2021), <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52021PC0206&from=EN> [1](#)
13. Croce, F., Andriushchenko, M., Sehwal, V., Debenedetti, E., Flammarion, N., Chiang, M., Mittal, P., Hein, M.: Robustbench: a standardized adversarial robustness benchmark. In: NeurIPS Track on Benchmark and Datasets (2021) [24](#), [31](#)
14. Croce, F., Hein, M.: Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In: ICML (2020) [5](#), [9](#), [27](#)
15. Croce, F., Hein, M.: Mind the box: l_1 -apgd for sparse adversarial attacks on image classifiers. In: ICML (2021) [10](#)
16. Croce, F., Hein, M.: Adversarial robustness against multiple l_p -threat models at the price of one and how to quickly fine-tune robust models to another threat model. In: ICML (2022) [2](#), [4](#), [5](#), [17](#), [21](#), [24](#), [42](#)
17. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. arXiv preprint arXiv:2105.05233 (2021) [37](#)
18. Dhurandhar, A., Chen, P.Y., Luss, R., Tu, C.C., Ting, P., Shanmugam, K., Das, P.: Explanations based on the missing: Towards contrastive explanations with pertinent negatives. In: NeurIPS (2018) [1](#)
19. Engstrom, L., Ilyas, A., Salman, H., Santurkar, S., Tsipras, D.: Robustness (python library) (2019), <https://github.com/MadryLab/robustness> [4](#), [5](#), [6](#), [7](#), [8](#), [10](#), [11](#), [12](#), [17](#), [18](#), [19](#), [20](#), [31](#), [32](#), [33](#), [34](#), [35](#), [36](#), [38](#), [39](#), [40](#), [41](#), [43](#)
20. Engstrom, L., Ilyas, A., Santurkar, S., Tsipras, D., Tran, B., Madry, A.: Adversarial robustness as a prior for learned representations (2019) [2](#)

21. Etmann, C., Lunz, S., Maass, P., Schönlieb, C.B.: On the connection between adversarial robustness and saliency map interpretability. In: ICML (2019) [1](#)
22. Gao, S., Li, Z.Y., Yang, M.H., Cheng, M.M., Han, J., Torr, P.: Large-scale unsupervised semantic segmentation. arXiv preprint arXiv:2106.03149 (2021) [6](#), [17](#), [25](#), [26](#)
23. Goh, G., Cammarata, N., Voss, C., Carter, S., Petrov, M., Schubert, L., Radford, A., Olah, C.: Multimodal neurons in artificial neural networks. Distill (2021) [11](#)
24. Gowal, S., Qin, C., Uesato, J., Mann, T., Kohli, P.: Uncovering the limits of adversarial training against norm-bounded adversarial examples. arXiv preprint arXiv:2010.03593v2 (2020) [24](#), [42](#)
25. Goyal, Y., Wu, Z., Ernst, J., Batra, D., Parikh, D., Lee, S.: Counterfactual visual explanations. In: ICML (2019) [2](#)
26. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: ICML (2017) [5](#)
27. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: ECCV (2016) [18](#)
28. Hendricks, L.A., Akata, Z., Rohrbach, M., Donahue, J., Schiele, B., Darrell, T.: Generating visual explanations. In: ECCV (2016) [2](#)
29. Hendricks, L.A., Hu, R., Darrell, T., Akata, Z.: Grounding visual explanations. In: ECCV (2018) [2](#)
30. Hendrycks, D., Mu, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: AugMix: A simple data processing method to improve robustness and uncertainty. ICLR (2020) [24](#)
31. Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., Song, D.: Natural adversarial examples. CVPR (2021) [5](#)
32. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: NeurIPS (2017) [5](#)
33. Hohman, F., Park, H., Robinson, C., Chau, D.H.: Summit: Scaling deep learning interpretability by visualizing activation and attribution summarizations. IEEE Transactions on Visualization and Computer Graphics (TVCG) (2020) [11](#)
34. Jaggi, M.: Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In: ICML (2013) [9](#)
35. Kolesnikov, A., Beyer, L., Zhai, X., Puigcerver, J., Yung, J., Gelly, S., Houlsby, N.: Big transfer (bit): General visual representation learning. In: ECCV (2020) [24](#), [31](#)
36. Laidlaw, C., Singla, S., Feizi, S.: Perceptual adversarial robustness: Defense against unseen threat models. In: ICLR (2021) [24](#)
37. Lang, O., Gandselman, Y., Yarom, M., Wald, Y., Elidan, G., Hassidim, A., Freeman, W.T., Isola, P., Globerson, A., Irani, M., Mosseri, I.: Explaining in style: Training a gan to explain a classifier in stylespace. arXiv preprint arXiv:2104.13369 (2021) [2](#)
38. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: NeurIPS (2017) [1](#)
39. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: ICLR (2018) [4](#)
40. Marcinkevičs, R., Vogt, J.E.: Interpretability and explainability: A machine learning zoo mini-tour. arXiv:2012.01805 (2020) [1](#)
41. Miller, T.: Explanation in artificial intelligence: Insights from the social sciences. Artificial Intelligence **267**, 1 – 38 (2019) [1](#)
42. Moraru, V.: An algorithm for solving quadratic programming problems. Computer Science Journal of Moldova (1997) [9](#)

43. Mothilal, R.K., Sharma, A., Tan, C.: Explaining machine learning classifiers through diverse counterfactual explanations. In: FAccT (2020) [1](#), [3](#), [5](#)
44. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models (2021) [37](#)
45. Pawlowski, N., Coelho de Castro, D., Glocker, B.: Deep structural causal models for tractable counterfactual inference. In: NeurIPS (2020) [1](#)
46. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021) [37](#)
47. Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do cifar-10 classifiers generalize to cifar-10? arXiv preprint arXiv:1806.00451 (2018) [29](#)
48. Ribeiro, M.T., Singh, S., Guestrin, C.: "why should i trust you?": Explaining the predictions of any classifier. In: KDD. p. 1135–1144 (2016) [1](#)
49. Samangouei, P., Saeedi, A., Nakagawa, L., Silberman, N.: Explaining: Model explanation via decision boundary crossing transformations. In: ECCV (2018) [2](#)
50. Sanchez, P., Tsafaris, S.A.: Diffusion causal models for counterfactual estimation. In: First Conference on Causal Learning and Reasoning (2022) [2](#), [37](#)
51. Santurkar, S., Tsipras, D., Tran, B., Ilyas, A., Engstrom, L., Madry, A.: Image synthesis with a single (robust) classifier. In: NeurIPS (2019) [2](#), [4](#)
52. Schut, L., Key, O., McGrath, R., Costabello, L., Sacaleanu, B., Corcoran, M., Gal, Y.: Generating interpretable counterfactual explanations by implicit minimisation of epistemic and aleatoric uncertainties. In: AISTATS (2021) [1](#)
53. Schutte, K., Moindrot, O., Hérent, P., Schiratti, J.B., Jégou, S.: Using stylegan for visual interpretability of deep learning models on medical images. In: NeurIPS Workshop "Medical Imaging Meets NeurIPS" (2020) [2](#)
54. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. International Journal of Computer Vision **128**(2), 336–359 (2019) [1](#)
55. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: Visualising image classification models and saliency maps. In: ICLR (2014) [1](#)
56. Singla, S., Nushi, B., Shah, S., Kamar, E., Horvitz, E.: Understanding failures of deep networks via robust feature extraction. In: CVPR (2021) [11](#)
57. Srinivas, S., Fleuret, F.: Full-gradient representation for neural network visualization. In: NeurIPS (2019) [1](#)
58. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. In: ICLR. pp. 2503–2511 (2014) [2](#), [4](#)
59. Torralba, A., Fergus, R., Freeman, W.T.: 80 million tiny images: A large data set for nonparametric object and scene recognition. IEEE PAMI **30**(11), 1958–1970 (2008) [24](#), [29](#)
60. Tsiligkaridis, T., Roberts, J.: Understanding frank-wolfe adversarial training. In: CVPR (2022) [9](#), [10](#)
61. Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., Madry, A.: Robustness may be at odds with accuracy. In: ICLR (2019) [4](#)
62. Verma, S., Dickerson, J.P., Hines, K.: Counterfactual explanations for machine learning: A review. arXiv preprint, arXiv:2010.10596 (2020) [1](#), [2](#), [3](#), [5](#)
63. Wachter, S., Mittelstadt, B., Russell, C.: Counterfactual explanations without opening the black box: Automated decisions and the GDPR. Harvard Journal of Law & Technology **31**, 841–887 (2018) [1](#), [3](#), [4](#)

64. Wang, Z., Wang, H., Ramkumar, S., Fredrikson, M., Mardziel, P., Datta, A.: Smoothed geometry for robust attribution. In: NeurIPS (2020) [1](#)
65. Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: CVPR (2020) [11](#), [12](#)
66. Yu, Y., Zhang, X., Schuurmans, D.: Generalized conditional gradient for sparse estimation. *Journal of Machine Learning Research* **18**(144), 1–46 (2017) [9](#)
67. Zech, J.R., Badgeley, M.A., Liu, M., Costa, A.B., Titano, J.J., Oermann, E.K.: Confounding variables can degrade generalization performance of radiological deep learning models. *arXiv preprint arXiv:1807.00431* (2018) [11](#)

Overview of Appendix

In the following we present several additional results and experimental details.

- In App. A we present more examples of $l_{1.5}$ -VCEs in the same wordnet categories and examine qualitatively and quantitatively the influence of different threat models used during adversarial training (AT) on the quality of $l_{1.5}$ - and l_2 -VCEs.
- In App. B we discuss additional evaluation of l_p -VCEs using FID scores on CIFAR10 and based on pixel-level segmentations in ImageNet-S [22].
- In App. C we discuss details of AFW and provide the proof of Proposition 1. In this section, we compare the performance of AFW, APGD and FW with different hyper-parameters on CIFAR10 models.
- In App. D we discuss experimental details.
- In App. E we provide more details on the spurious features we discovered using VCEs and present in App. E.3 an additional example of a spurious feature where our VCEs show that human features are associated with the class “tench”.
- In App. F we show how VCEs using guided and regularized diffusion process following [2] can be generated. Currently, this is the only method which uses a generative model, works on ImageNet, provides code and does something similar to the generation of VCEs.
- In App. G we show the best and the worst examples from the user study.
- In App. H we show randomly selected VCEs for both ImageNet and CIFAR10.
- In App. I we compare 3 threat models for different radii.

A VCEs with robust models

We study several aspects of the generation of VCEs with (adversarially) robust models.

A.1 $l_{1.5}$ -VCEs for ImageNet and CIFAR10.

For ImageNet we use as in the main paper the Madry l_2 -robust model [19] plus finetuning for multiple-norm robustness [16] to generate more $l_{1.5}$ -VCEs for classes belonging to a related WordNet synset in Figure 9 and 10. In detail, we traverse the WordNet tree starting at the root node and compute for each inner node the number of ILSVRC2012 leaf-nodes below it. If a node has between 3 and 10 ILSVRC2012 leafs below it, we add it and do not further process any child nodes. If it has more than 10 leafs below it, we process each child node the same way. By doing this, we can create clusters of related classes that can act as meaningful targets for our VCEs. For each cluster, we then randomly sample an image from one of the ILSVRC2012 leafs and create VCEs into all classes

in that cluster. This is particularly important on ImageNet, as it contains very different classes and for some pairs changing one into the other is not feasible via a subtle change and would require large budgets to modify the entire image content. One can see that our $l_{1.5}$ -VCEs realize quite subtle class-specific changes of the original image.

A.2 Ablation of threat models during AT.

We have seen in Tab. 5 that in order to generate realistic l_2 VCEs one needs (adversarially) robust models. In this subsection we want to investigate, which threat model used for adversarial training VCEs is sufficient to produce realistic VCEs. We do this study for models trained both on CIFAR10 and ImageNet and check the corresponding $l_{1.5}$ - and l_2 -VCEs.

CIFAR10. For CIFAR10 we train seven models with adversarial training all with PreActResNet-18 architecture [27]: five using a l_2 -threat model with radii $\epsilon_2 \in \{0.1, 0.25, 0.5, 0.75, 1\}$, one with l_1 -threat model with radius $\epsilon_1 = 12$, and one with l_∞ -threat model with radius $\epsilon_\infty = 8/255$.

As it can be seen from the Figs. 13 and 14 and Tab. 3, for l_2 -AT a certain sufficiently large radius ϵ_2 is required. Whereas 0.1 seems not sufficient, 0.25, 0.5, 0.75 look visually similar, while for 1.0 one observes also more artefacts again. The same is true for the l_∞ -robust model. The best model is actually the l_1 -robust model which produces at least in Figs. 13 and 14 the best visual quality apart from the multiple-norm fine-tuned model GU+FT which we have used throughout the paper and which achieves the best visual quality. Thus one take away message is that multiple-norm robust models are working for all l_p -VCEs with $p \in \{1.5, 2\}$ due to their simultaneous adversarial robustness in all threat models. These observations are also supported by our quantitative evaluation via FID-scores in Tab. 3.

Finally, by comparing Figs. 13 and 14, we can see again that $l_{1.5}$ -VCEs produce more sparse and object-related changes.

ImageNet. For ImageNet-1k we compare six models with ResNet50 architecture [27] which are either taken from [19] or are fine-tuned versions of these models. The first one is non-robust, which we denote by *RN-50*. Two are pre-trained models from [19]: *Madry l_2* , which corresponds to *Madry* model in the main paper and *Madry l_∞* . First one is l_2 -adversarially trained with $\epsilon_2 = 3$. Second one is l_∞ -adversarially trained with $\epsilon_\infty = 4/255$. The other three models we obtained by fine-tuning *Madry l_2* and *Madry l_∞* . One is *Madry l_2 +FT*, which corresponds to the *Madry+FT* model in the main paper. It is obtained by using multiple-norm robust fine-tuning for 3 epochs with $\epsilon_1 = 255$ and $\epsilon_\infty = 4/255$. The other one is *Madry l_∞ +FT*, which is obtained by the same fine-tuning, but for 1 epoch. The last one, *Madry l_2 + l_1 FT*, is obtained by fine-tuning for 1 epoch only wrt l_1 . From Figs. 11 and 12 and Tab. 4 similar observations can be made, that is multiple-norm fine-tuning can significantly improve both the image quality, and ID/OD FIDs by increasing l_1 -, $l_{1.5}$ - and l_2 -RA, and that $l_{1.5}$ produce more object-related changes.


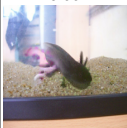

Original	$\epsilon_{1.5} = 50$	$\epsilon_{1.5} = 75$	$\epsilon_{1.5} = 100$	$\epsilon_{1.5} = 50$	$\epsilon_{1.5} = 75$	$\epsilon_{1.5} = 100$
trium. arch: 0.79	→trium. arch: 1.00	→trium. arch: 1.00	→trium. arch: 1.00	→megalith: 0.90	→megalith: 1.00	→megalith: 1.00
						
WordNet: memorial monument - GT class: triumph arch						
pizza: 0.12 	→ potpie: 0.58	→ potpie: 0.86	→ potpie: 0.96	→ pizza: 0.93	→ pizza: 0.99	→ pizza: 1.00
						
	→ hotdog: 0.55	→ hotdog: 0.97	→ hotdog: 1.00	→ cheesb.: 0.82	→ cheesb.: 0.98	→ cheesb.: 1.00
						
WordNet: dish - GT class: potpie						
axolotl: 0.18 	→ axolotl: 0.99	→ axolotl: 1.00	→ axolotl: 1.00	→ fire salam.: 0.47	→ fire salam.: 0.95	→ fire salam.: 0.99
						
	→ c. newt: 0.47	→ c. newt: 0.91	→ c. newt: 0.98	→ t. frog: 0.08	→ t. frog: 0.39	→ t. frog: 0.73
						
WordNet: amphibian - GT class: axolotl						
ladybug: 0.88 	→ rhinoceros beetle: 0.28	→ rhinoceros beetle: 0.96	→ rhinoceros beetle: 1.00	→ ground beetle: 0.15	→ ground beetle: 0.70	→ ground beetle: 0.94
						
	→ long-horned beetle: 0.15	→ long-horned beetle: 0.89	→ long-horned beetle: 0.99	→ tiger beetle: 0.45	→ tiger beetle: 0.98	→ tiger beetle: 1.00
						
WordNet: beetle - GT class: ladybug						

Fig. 9: Further $l_{1.5}$ -VCEs for WordNet synset classes for the Madry[19]+FT model with varying radii. ILSVRC2012 images are transformed into different classes from the same WordNet synset as the original.




























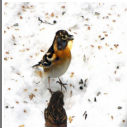





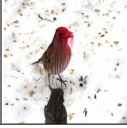





Original	$\epsilon_{1.5} = 50$	$\epsilon_{1.5} = 75$	$\epsilon_{1.5} = 100$	$\epsilon_{1.5} = 50$	$\epsilon_{1.5} = 75$	$\epsilon_{1.5} = 100$
 fig: 0.27	→ strawberry: 0.97	→ strawberry: 1.00	→ strawberry: 1.00	→ pineapple: 0.99	→ pineapple: 1.00	→ pineapple: 1.00
						
	→ custard apple: 0.95	→ custard apple: 0.99	→ custard apple: 1.00	→ pomegran.: 0.76	→ pomegran.: 0.97	→ pomegran.: 0.99
						
WordNet: edible fruit - GT class: fig						
 lifeboat: 0.79	→ gondola: 1.00	→ gondola: 1.00	→ gondola: 1.00	→ fireboat: 0.83	→ fireboat: 1.00	→ fireboat: 1.00
						
	→ lifeboat: 1.00	→ lifeboat: 1.00	→ lifeboat: 1.00	→ speedboat: 0.21	→ speedboat: 0.85	→ speedboat: 0.99
						
WordNet: boat - GT class: lifeboat						
 junco: 0.97	→ brambling: 1.00	→ brambling: 1.00	→ brambling: 1.00	→ goldfinch: 0.81	→ goldfinch: 0.99	→ goldfinch: 1.00
						
	→ house finch: 0.94	→ house finch: 0.99	→ house finch: 1.00	→ indigo bunting: 0.34	→ indigo bunting: 0.99	→ indigo bunting: 1.00
						
WordNet: finch - GT class: brambling						

Fig.10: **Further $l_{1.5}$ -VCEs for WordNet synset classes** for the Madry[19]+FT model with varying radii. ILSVRC2012 images are transformed into different classes from the same WordNet synset as the original.

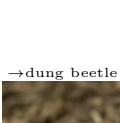
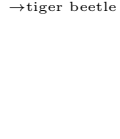
Orig.	RN-50	Madry l_2	Madry l_2 +FT	Madry l_∞	Madry l_∞ +FT	Madry l_2+l_1 FT
 →dung beetle  →tiger beetle	$p_i:0.85,$ $p_e:1.00$	$p_i:0.03,$ $p_e:1.00$	$p_i:0.04,$ $p_e:0.99$	$p_i:0.45,$ $p_e:1.00$	$p_i:0.13,$ $p_e:1.00$	$p_i:0.05,$ $p_e:0.99$
	$p_i:0.02,$ $p_e:1.00$	$p_i:0.93,$ $p_e:1.00$	$p_i:0.85,$ $p_e:1.00$	$p_i:0.37,$ $p_e:1.00$	$p_i:0.77,$ $p_e:1.00$	$p_i:0.85,$ $p_e:1.00$

Fig. 11: **ImageNet: Ablation of the threat model during training.** $l_{1.5}$ -VCEs at $\epsilon_{1.5} = 50$ for ImageNet models trained adversarially with varying threat models.



Orig.	RN-50	Madry l_2	Madry l_2 +FT	Madry l_∞	Madry l_∞ +FT	Madry l_2+l_1 FT
 →dung beetle  →tiger beetle	$p_i:0.85,$ $p_e:1.00$	$p_i:0.03,$ $p_e:1.00$	$p_i:0.04,$ $p_e:1.00$	$p_i:0.45,$ $p_e:1.00$	$p_i:0.13,$ $p_e:1.00$	$p_i:0.05,$ $p_e:1.00$
	$p_i:0.02,$ $p_e:1.00$	$p_i:0.93,$ $p_e:1.00$	$p_i:0.85,$ $p_e:1.00$	$p_i:0.37,$ $p_e:1.00$	$p_i:0.77,$ $p_e:1.00$	$p_i:0.85,$ $p_e:1.00$

Fig. 12: **ImageNet: Ablation of the threat model during training.** l_2 -VCEs at $\epsilon_2 = 12$ for ImageNet models trained to be adversarially robust with respect to different threat models. Here l_1 -finetuning [16] (outermost right column) is fine-tuned to be l_1 -robust.



























Original GU+FT		$l_2, \epsilon=0.1$	$l_2, \epsilon=0.25$	$l_2, \epsilon=0.5$	$l_2, \epsilon=0.75$	$l_2, \epsilon=1$	$l_\infty, \epsilon=8/255$	$l_1, \epsilon=12$
 →car  →truck	$p_i:0.17, p_e:1.00$	$p_i:0.54, p_e:1.00$	$p_i:0.67, p_e:1.00$	$p_i:0.67, p_e:1.00$	$p_i:0.28, p_e:1.00$	$p_i:0.56, p_e:0.99$	$p_i:0.84, p_e:1.00$	$p_i:0.56, p_e:1.00$
								
								
								

Fig. 13: **CIFAR-10: Ablation of the threat model during training.** $l_{1.5}$ -VCEs at $\epsilon_{1.5} = 6$ for CIFAR-10 models trained adversarially with varying threat models.



























Original GU+FT		$l_2, \epsilon=0.1$	$l_2, \epsilon=0.25$	$l_2, \epsilon=0.5$	$l_2, \epsilon=0.75$	$l_2, \epsilon=1$	$l_\infty, \epsilon=8/255$	$l_1, \epsilon=12$
 →car  →truck	$p_i:0.17, p_e:1.00$	$p_i:0.54, p_e:1.00$	$p_i:0.67, p_e:1.00$	$p_i:0.67, p_e:1.00$	$p_i:0.28, p_e:1.00$	$p_i:0.56, p_e:1.00$	$p_i:0.84, p_e:1.00$	$p_i:0.56, p_e:1.00$
								
								
								

Fig. 14: **CIFAR-10: Ablation of the threat model during training.** l_2 -VCEs at $\epsilon_2 = 2.5$ for CIFAR-10 models trained to be adversarially robust with respect to different threat models.

Table 3: **CIFAR10**: Evaluation of employed (robust) classifiers trained with different threat models for standard accuracy, l_1 -, $l_{1.5}$ - and l_2 -robust accuracy (RA) evaluated at $\epsilon_1 = 12$, $\epsilon_{1.5} = 1.5$, and $\epsilon_2 = 0.5$ respectively (first 1k test points). Further, FID scores for l_2 - and $l_{1.5}$ -VCEs for in-and out-of-distribution inputs and their average is shown. The threat model is indicated in the headers.

FID scores for $l_{1.5}$ -VCE	WRN-70-16		PreActResNet-18						
	GU+FT		l_2 , $\epsilon = 0.1$	l_2 , $\epsilon = 0.25$	l_2 , $\epsilon = 0.5$	l_2 , $\epsilon = 0.75$	l_2 , $\epsilon = 1$	l_∞ , $\epsilon = 8/255$	l_1 , $\epsilon = 12$
	Acc.	90.8	91.6	90.8	88.8	84.8	80.6	82.8	87.1
	l_1 -RA	58.0	9.8	15.3	25.7	35.2	44.0	7.1	60.1
	$l_{1.5}$ -RA	76.7	54.8	62.2	66.2	67.1	64.9	44.6	66.8
	l_2 -RA	79.2	61.8	66.5	68.6	67.7	65.3	59.9	64.9
	ID	11.4	16.3	14.1	13.7	12.9	11.9	22.9	13.6
	OD	46.2	46.9	45.8	48.8	51.6	52.2	65.2	48.9
	Avg.	28.8	31.6	29.9	31.2	32.2	32.1	44.0	31.2
	FID scores for l_2 -VCE	ID	11.9	21.3	18.2	16.3	15.0	13.7	20.2
OD		41.2	47.2	44.7	47.1	49.6	50.1	60.7	48.7
Avg.		26.7	34.3	31.4	31.7	32.3	31.9	40.5	34.4

Table 4: **ImageNet**: Evaluation of employed (robust) classifiers trained with different threat models for standard accuracy, l_1 -, $l_{1.5}$ - and l_2 -robust accuracy (RA) evaluated at $\epsilon_1 = 255$, $\epsilon_{1.5} = 12.5$, and $\epsilon_2 = 2$ respectively (1k test points). Further, FID scores for l_2 - and $l_{1.5}$ -VCEs for in-and out-of-distribution inputs and their average is shown. The threat model is indicated in the headers.

ResNet50						
	RN-50	Madry l_2	Madry l_2 +FT	Madry l_∞	Madry l_∞ +FT	Madry l_2 + l_1 FT
Acc.	76.0	57.9	57.5	62.4	57.6	56.9
l_1 -RA	0.0	13.0	25.5	0.0	22.6	29.2
$l_{1.5}$ -RA	0.0	37.4	40.1	2.8	39.4	37.9
l_2 -RA	0.0	45.7	44.6	18.3	44.5	40.3
FID scores for $l_{1.5}$-VCE	ID	9.4	8.4	6.9	8.3	7.1
	OD	50.3	24.3	22.6	39.3	23.3
	Avg.	29.8	16.4	14.8	23.8	15.1
FID scores for l_2-VCE	ID	20.0	8.4	7.9	9.1	7.8
	OD	60.3	22.8	23.1	34.9	23.3
	Avg.	40.1	15.6	15.5	22.0	15.6

B Further evaluation

First, we show the qualitative and quantitative (with FIDs) evaluation of the different models on CIFAR10.












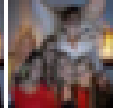
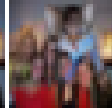
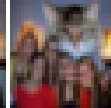
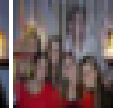
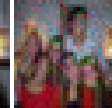
Orig.	BiT-M[35]	RST-S[8]	RATIO[1]	GU[24]	GU+FT[16]	PAT[36]	HenC[30]
car → truck	$p_i:0.01,$ $p_e:1.00$	$p_i:0.51,$ $p_e:1.00$	$p_i:0.78,$ $p_e:1.00$	$p_i:0.44,$ $p_e:1.00$	$p_i:0.83,$ $p_e:1.00$	$p_i:0.57,$ $p_e:0.99$	$p_i:0.04,$ $p_e:1.00$
							
→ cat	$p_i:0.05,$ $p_e:1.00$	$p_i:0.26,$ $p_e:0.95$	$p_i:0.10,$ $p_e:1.00$	$p_i:0.54,$ $p_e:1.00$	$p_i:0.35,$ $p_e:1.00$	$p_i:0.45,$ $p_e:0.95$	$p_i:0.07,$ $p_e:1.00$
							

Fig. 15: **CIFAR10**: l_2 -VCEs of radius $\epsilon = 2.5$ of different classifiers for the change “car → truck” (top) and “frog → cat” (middle) and “→ cat” for an OOD image from 80M TM [59] not belonging to any of the CIFAR10 classes (bottom). We denote by p_i resp. p_e the confidence in the target class for the original image and the generated VCE. All generated VCEs are valid as high confidence in the target class is achieved but only adversarially robust models, see Tab. 5, show class-specific changes.

For this purpose, we qualitatively compare different CIFAR10 classifiers in Fig. 15 and quantitatively in Tab. 5: *BiT-M*: a non-robust BiT-M-R50x1 model [35] with high accuracy, *RST-s*: a WRN-28-10 trained only with additive noise using stability training [8], *RATIO*: a WRN-34-10 using l_2 -adv. training on the in-and out-distribution, *GU*: the SOTA l_2 -robust WRN-70-16 from [24] trained with extra data, *GU+FT*: we fine-tune the GU model [16] to get multiple-norm-robust, *PAT*: the PAT-self ResNet50 from [36] which uses adversarial training with respect to a perceptual distance, *HenC*: a ResNext29 trained to be robust against common corruptions (no adversarial training). Note that all l_2 -adversarially trained models were trained using a radius of 0.5. All models are the original models from the authors, most of them are available in RobustBench [13], and we report their standard and robust accuracy against adversarial l_1 - and l_2 -perturbations in Tab. 5. The perturbation model for the generation of the VCEs in Eq. (2) is a l_2 -ball of radius 2.5 following [1]. From Fig. 15 one can observe that the two non-robust models BiT-M and HenC do not produce any meaningful counterfactuals.

Surprisingly, the RST-s model has some adversarial robustness but its l_2 -VCEs do minimal changes to the image which show little class-specific features of the target class. Thus the FID score for the in-distribution is low, but the

FID score of the out-distribution is high. Moreover, the PAT-model, which uses a threat model with respect to a perceptual distance but also has non-trivial l_1 - and l_2 -adversarial robustness, produces VCEs that show strong artefacts. The best VCEs are generated by RATIO, GU and GU+FT, which also have the highest l_2 -adversarial robustness. Among them, RATIO and GU+FT produce the most visually realistic VCEs and also have the best FID scores for in- and out-distribution. In particular, the multiple-norm finetuning of the GU model seems to significantly boost the generative properties, both for l_2 -VCEs and the $l_{1.5}$ -VCEs (see Sec. 2.3). High l_2 -adversarial robustness alone, as for the SOTA GU model, is not the only factor which determines the quality of the generated VCEs. In App. A we provide a more detailed analysis, *e.g.* which radius for the threat model for adversarial training is required for good VCEs, and also repeat a similar experiment for ImageNet even though there are fewer adversarially robust models available.

Table 5: **CIFAR10**: Evaluation of (robust) classifiers for standard accuracy, l_1 -, $l_{1.5}$ - and l_2 -robust accuracy (RA) evaluated at $\epsilon_1 = 12$, $\epsilon_{1.5} = 1.5$, and $\epsilon_2 = 0.5$ respectively (first 1k test points). Further, FID scores for l_1 -, $l_{1.5}$ -, and l_2 -VCEs for in-and out-of-distribution inputs and their average are shown. For all classifiers except RATIO $l_{1.5}$ -VCEs attain the best average FID-score.






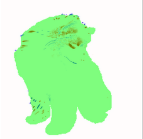
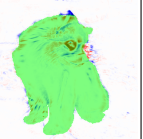
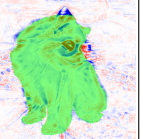
	BiT-M	RST-s	RATIO	GU	GU+FT	PAT	HenC
Acc.	97.4	87.9	94.0	94.7	90.8	82.4	95.8
l_1 -RA	0.0	36.5	34.3	33.4	58.0	32.9	0.0
$l_{1.5}$ -RA	0.0	70.4	75.4	76.8	76.7	59.2	0.3
l_2 -RA	0.0	71.4	79.9	81.7	79.2	62.4	0.1
FID scores for l_1-VCE	ID	25.1	26.0	24.4	31.1	10.2	22.7
	OD	79.5	72.6	57.8	71.4	52.7	79.5
	Avg.	52.3	49.8	41.1	51.3	31.5	51.1
FID scores for $l_{1.5}$-VCE	ID	12.2	8.5	11.7	12.3	9.2	18.8
	OD	62.7	51.6	30.4	52.5	43.4	62.4
	Avg.	42.5	30.1	19.5	32.4	26.3	40.6
FID scores for l_2-VCE	ID	55.4	10.3	12.2	15.8	11.9	37.9
	OD	83.9	50.7	26.0	53.9	41.2	67.2
	Avg.	69.7	30.5	19.1	34.9	26.7	52.6

We additionally evaluate the l_p -VCEs in a user study and evaluate their object-relatedness using ImageNet-S dataset [22].

ImageNet-S Evaluation We use the recently introduced dataset ImageNet-S [22] with pixel-level segmentations of ImageNet images to evaluate how well the l_p -VCEs are located on the object. We limit the evaluation to a subset of 2048

images that only contain a single class and where the segmented pixels are connected (one object). The target class is chosen using the WordNet hierarchy. For each VCE we compute the absolute difference to the original image and sum up over color channels. We normalize this to get a distribution p of pixel changes over the image.

Given p , we evaluate three metrics: a) the expected distance $\mathbb{E}[d]$ calculates the distance to the closest mask pixel for each pixel in the image and then calculates a weighted average using the distribution of changes p . The larger the expected distance $\mathbb{E}[d]$ the more far away are most changes from the object. However, note that a VCE might need to change also pixels outside the object so that an “ideal” VCE need not have zero expected distance; b) the probability mass $\int p$ of changes located in the segmentation; c) the Intersection over Union (IOU). As this is a measure between two binary images, we have to discretize p . To do this, we sort p in a descending fashion and activate pixels until their cumulative probability is at least 0.95, thus they explain most of the changes of the image. While the previous metrics are optimal if all changes are inside the mask, the IOU requires the changes to also cover most of the object. Thus if the changes are too sparse, the IOU will be small even if all of them are located on the object.

	$l_1, \epsilon = 400$	$l_{1.5}, \epsilon = 50$	$l_2, \epsilon = 12$
	\rightarrow	\rightarrow	\rightarrow
GT: Samoyed	Pomeranian: 0.33	Pomeranian: 0.94	Pomeranian: 0.99
			
Mask	$\mathbb{E}[d] = 0.04$	$\mathbb{E}[d] = 2.85$	$\mathbb{E}[d] = 77.82$
			

	l_1	$l_{1.5}$	l_2
$\mathbb{E}[d]$	7.63	14.38	21.20
$\int p$	0.69	0.57	0.45
IOU@0.95	0.05	0.38	0.35

Fig. 16: Evaluation of changes of l_p -VCE based on pixel-level segmentations in the ImageNet-S [22] dataset. l_1 -VCEs are too sparse, while l_2 -VCEs induce a lot of changes in the background. The $l_{1.5}$ -VCE induce more subtle changes which are more concentrated on the object.

Results and an illustration in Fig. 16. The l_1 -VCEs are well located on the object but too sparse which results in a very small IOU. The l_2 -VCEs have a large expected distance $\mathbb{E}[d]$ which means that a lot of pixels outside of the object are changed: this can also be verified from the given example in Fig. 16 where most background pixels are changed. $l_{1.5}$ -VCEs offer a good balance between the

two extremes. They are largely located on or close to the object without being too sparse, which results in more realistic images than the l_1 -VCEs.

C Auto-Frank-Wolfe

C.1 Details of AFW algorithm

We extend the scheme of APGD for selecting the step size to the FW scheme to make the choice of γ^k adaptive (see Eq. (5)). The algorithm follows Algorithm 1 from [14] where the step size at each iteration is indicated by η^k : however, the update in lines 7, 8 becomes

$$\hat{x}^{k+1} = (1 - \eta^k)\hat{x}^k + \eta^k s^k, \quad (7)$$

where s^k is the solution of Eq. (4), computed according to the Proposition 1, and $\eta^k = \frac{M}{\sqrt{k+2}}$, where in the beginning $M = 2$, and then at steps that satisfy the same *Condition 1* and *Condition 2* as in [14], we update $M \leftarrow M \cdot 0.75$.

C.2 Omitted proof

Proof 1 (Proof of Proposition 1) *Introducing*

$$v = |w|, \quad \gamma = \max\{-x \operatorname{sign} w, (1 - x) \operatorname{sign} w\}$$

we have an equivalent problem

$$\begin{aligned} \arg \min_{\eta \in \mathbb{R}^d} -\langle v, \eta \rangle \quad \text{s.t.} \quad & \sum_i \eta_i^p \leq \epsilon^p, \\ & \eta_i \in [0, \gamma_i] \quad i = 1, \dots, d, \end{aligned}$$

with solution η^ for which $\delta^* = \eta^* \operatorname{sign} w$. Note that $\gamma_i = 0$ implies $\eta_i = 0$, and similarly, if $w_i = v_i = 0$ then $\eta_i = 0$ since the i -th coordinate does not contribute to the objective function. Thus in the following we can assume $\gamma_i, v_i > 0$. The Lagrangian is given by*

$$L(\eta, \mu, \alpha, \beta) = -\langle v, \eta \rangle + \mu(\langle \mathbf{1}, \eta^p \rangle - \epsilon^p) - \langle \alpha, \eta \rangle + \langle \beta, \eta - \gamma \rangle,$$

with $\mu \geq 0$, $\alpha, \beta \in \mathbb{R}_+^d$, and the power of a vector is meant componentwise. It has gradient wrt η

$$\nabla_\eta L(\eta, \mu, \alpha, \beta) = -v + p\mu\eta^{p-1} - \alpha + \beta,$$

which yields the optimality conditions

$$\begin{aligned} -v_i + p\mu\eta_i^{p-1} - \alpha_i + \beta_i &= 0 & i = 1, \dots, d \\ \alpha_i \eta_i &= 0 & i = 1, \dots, d \\ \beta_i(\gamma_i - \eta_i) &= 0 & i = 1, \dots, d \\ \mu(\langle \mathbf{1}, \eta^p \rangle - \epsilon^p) &= 0 \\ \alpha_i, \beta_i, \mu &\geq 0 & i = 1, \dots, d. \end{aligned}$$

We can distinguish three cases:

- $\alpha_i = 0, \beta_i > 0 \Rightarrow \eta_i = \gamma_i, \quad p\mu\gamma_i^{p-1} < v_i,$
- $\alpha_i > 0, \beta_i = 0 \Rightarrow \eta_i = 0, \quad v_i = -\alpha_i < 0,$
- $\alpha_i = 0, \beta_i = 0 \Rightarrow \eta_i \in [0, \gamma_i], \quad p\mu\eta_i^{p-1} = v_i.$

Note that the second case is not possible since we assume $v_i > 0$, while in the third case we have $p\mu\gamma_i^{p-1} \geq v_i$ because of the interval η_i belongs to. Thus, we have

$$\mu < \frac{v_i}{p\gamma_i^{p-1}} \Rightarrow \eta_i = \gamma_i, \quad \mu \geq \frac{v_i}{p\gamma_i^{p-1}} \Rightarrow \eta_i = \left(\frac{v_i}{p\mu}\right)^{\frac{1}{p-1}}$$

If we have $\mu > 0$, then

$$\sum_i \eta_i^p = \sum_{i \in I^-(\mu)} \gamma_i^p + \sum_{i \in I^+(\mu)} \left(\frac{v_i}{p\mu}\right)^{\frac{p}{p-1}} = \epsilon^p,$$

with $I^-(\mu) = \{i : \mu < \frac{v_i}{p\gamma_i^{p-1}}\}$ and $I^+(\mu) = \{i : \mu \geq \frac{v_i}{p\gamma_i^{p-1}}\}$.

This is equivalent to finding the solution μ^* of the equation

$$f(\mu) := \sum_{i \in I^-(\mu)} \gamma_i^p + \sum_{i \in I^+(\mu)} \left(\frac{v_i}{p\mu}\right)^{\frac{p}{p-1}} - \epsilon^p = 0. \quad (8)$$

which exists if $f(0) \geq 0$ since f is continuous on $[0, +\infty)$ and converges to $-\epsilon^p$ for $\mu \rightarrow +\infty$. The solution is also unique as f is decreasing. To find μ^* it is possible to sort the set $M = \{\frac{v_i}{p\gamma_i^{p-1}}, i = 1, \dots, d\} \cup \{0, +\infty\}$. There exists $m_j \in M$ is such that $f(m_j) \geq 0$ and $f(m_{j+1}) < 0$, and μ^* solves

$$\sum_{i \in I^-(m_j)} \gamma_i^p + \sum_{i \in I^+(m_j)} \left(\frac{v_i}{p\mu}\right)^{\frac{p}{p-1}} - \epsilon^p = 0,$$

that is

$$(p\mu^*)^{\frac{p}{p-1}} = \sum_{i \in I^+(m_j)} v_i^{\frac{p}{p-1}} \left(\epsilon^p - \sum_{i \in I^-(m_j)} \gamma_i^p \right)^{-1},$$

and $\eta_i^* = \min\{\gamma_i, \left(\frac{v_i}{p\mu^*}\right)^{\frac{1}{p-1}}\}$. Finally note that if $f(0) < 0$, then $\mu = 0$ which implies $\eta_i = \gamma_i$ if $v_i > 0$, $\eta_i = 0$ else. Note the most complex operation involved is the sorting of the set M , which gives the complexity of the algorithm.

C.3 Additional experiments

First, we repeat the experimental comparison of AFW to APGD and the existing variants of FW reported in Sec. 3 on CIFAR10, using the robust models, in the l_2 -threat model with $\epsilon = 2.5$. Figure 17 shows that even in this case AFW is competitive with the best methods and then effective optimizer. Moreover, while the highest loss is attained by a version of FW with decaying schedule, this is




Fig. 17: **APGD and AFW vs FW.** APGD and AFW outperform most of the FW versions except for 2 in terms of final mean (over 1000 images) log probability.




Fig. 18: **AFW, FW with vs without clipping.** Applying Proposition 1 to do AFW in the intersection of l_p -ball and $[0, 1]^d$ consistently outperforms constrained optimization in the l_p -ball followed by clipping.

achieved only by tuning γ_0 which needs to be done for each classifier and threat model, unlike for AFW and APGD.

Second, we study the effect of including the box-constraints in the FW scheme rather than clipping after each iteration as done by prior works. Then, we compare in Fig. 18 the performance of AFW and FW (same setup as above) with either the linear optimization suggested by Proposition 1 (solid lines) or the standard optimization in the l_p -ball followed by clipping (dashed lines). We observe that considering the true threat models, i.e. the intersection of l_p -ball and $[0, 1]^d$, yields consistently a small improvement over the baseline.

C.4 Effect of AFW versus APGD on the resulting l_2 -VCEs

In this section we compare l_2 -VCEs for AFW and APGD quantitatively. We use for CIFAR10 the GU+FT model and for ImageNet the Madry+FT models to compare the end probabilities for 10.000 VCEs used for the FID computation in Fig. 19. One sees that the two methods achieve similar results, with APGD producing only slightly higher probabilities.

D Experimental Details

D.1 Calibration

In order to have comparable confidence values, we use CIFAR10.1 [47, 59] for CIFAR10 models, respectively 20.000 images from the ImageNet test set for ImageNet models to do temperature scaling. In Tab. 6 we show the computed values of temperature T , by which the output of each classifier is divided before computing applying the softmax function. We observe that models with standard

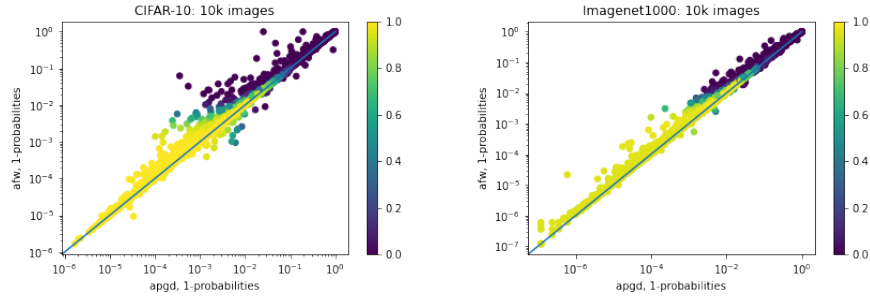


Fig. 19: **End probabilities for APGD vs AFW.** APGD performs similarly to AFW in terms of end probabilities for 10.000 VCEs used for the FID computation. For CIFAR10 GU+FT model is used, and for ImageNet - Madry+FT. Color represents the density of the datapoints with yellow indicating the highest density.

Table 6: Temperature obtained after the temperature scaling for CIFAR10 models.

Temperature for CIFAR10 models introduced in Sec. 2							
	BiT-M	RST-s	RATIO	GU	GU+FT	PAT	HenC
Temp.	1.02	1.12	0.42	0.69	0.59	0.52	1.31
Temperature for CIFAR10 models introduced in Tab. 3							
	$l_2, \epsilon=0.1$	$l_2, \epsilon=0.25$	$l_2, \epsilon=0.5$	$l_2, \epsilon=0.75$	$l_2, \epsilon=1$	$l_\infty, \epsilon=8/255$	$l_1, \epsilon=12$
Temp.	1.49	1.44	1.31	0.90	0.79	0.8	1.03
Temperature for ImageNet models introduced in Tab. 4							
	RN50	Madry	Madry + FT	Madry l_∞	Madry l_∞ + FT	Madry + l_1 FT	
Temp.	1.16	0.84	0.72	0.82	0.73	0.74	

training or with adversarial training with small radii are in general overconfident ($T > 1$), while the most robust ones, especially trained for multiple norm robustness, are underconfident ($T < 1$).

D.2 VCEs generation

For generating the VCEs with AFW and APGD we run the algorithms for 75 restarts and 5 random restarts. The final output is selected as the one attaining the highest value of the objective function, i.e. the log-probability of the target class.

D.3 Reproducibility

Most of the classifiers we use are publicly available. In particular, RATIO, GU and Hen can be found in RobustBench [13], RST-s and PAT are provided by the original papers, BiT-M is a BiT-M-R50x1 model that can be fine-tuned following instructions from the authors of [35], Madry l_2 and l_∞ are part of the Robustness library [19], while the remaining ones are obtained via personal communications with the authors. Moreover, we train the classifiers for the ablation study in App. A.

E Detection of spurious features via VCEs

We revisit in each subsection the spurious feature which we detected with our VCEs. Additionally, we illustrate a new spurious feature which is known, that certain classes of fish are associated with humans. For each case we show the VCEs and samples from the training set illustrating the origin of the spurious feature.

E.1 Failure Mode A: Watermarks

We show in Fig. 27 more examples of the spurious text like feature appearing in the $l_{1.5}$ -VCEs for the target class “granny smith” which are obviously unrelated to the class. In Section 4 we came up with the hypothesis that the reason for this spurious feature which the model has picked up is the large number of images containing a watermark, in particular one from “iStockphoto”, in the training set of “granny smith”, see Fig. 21 for other examples from the training set.

E.2 Failure Mode B: Cages in shark images

In Fig. 22 we show more examples of the spurious “cage” feature the model has picked up for the white shark class. It is interesting that this is even a very dominating feature, in the sense that it is easier to paint in some “cage”-like structures in the image to change the class to “white shark” rather than changing some parts of the image into a white shark. In Fig. 23 we show a sample of training images from the class “white shark” which shows this spurious feature of “cages” which the model has picked up.

E.3 Failure Mode C: Human features in fish images

It is well known that some classes of fish, e.g. “tench”, appear very often with humans in the training set, as this is a popular fish for anglers which they proudly present to the camera. In some of the $l_{1.5}$ -VCEs for the class “tench”, see examples in Fig. 24, we see human fingers or even human faces appearing, which shows that the model has picked up this artefact of the training set. We show in Fig. 25 examples from the training set of the class “tench” in ImageNet.













Original	$\epsilon_{1.5} = 50$	$\epsilon_{1.5} = 75$	$\epsilon_{1.5} = 100$
granny sm.: 0.92	→granny sm.: 1.00	→granny sm.: 1.00	→granny sm.: 1.00
			
lemon: 0.36	→granny sm.: 0.96	→granny sm.: 0.99	→granny sm.: 1.00
			
bell pepper: 1.00	→granny sm.: 0.01	→granny sm.: 0.09	→granny sm.: 0.75
			

Fig. 20: **Spurious feature of watermarks for the target class “granny smith”**. $l_{1.5}$ -VCEs for Madry[19]+FT with varying radii for 3 images for the target class “granny smith”. They show again that classifier has learned to associate the spurious “text” feature with this class.

Watermarks in the training set of “granny smith” in ImageNet



Fig. 21: More examples from the training set of the class “granny smith” in ImageNet. In total 90 out of 1300 training images contain a watermark, out of which 53 contain the one of “iStockphoto”. The fraction of watermarked images seems to be significantly larger than for other classes.




Fig. 22: **Spurious feature of “cages” for the target class “white shark”.** $l_{1.5}$ -VCEs for Madry[19]+FT with varying radii for 3 images for the target class “white shark”. They show again that classifier has learned to associate the spurious “cage” feature with this class.

“Cages” in the training set of “white shark” in ImageNet



Fig. 23: A large fraction of training images of the class “white shark” in ImageNet contains cages as the divers who take the photographs need to be protected.




Fig. 24: **Spurious feature of “human hands” for the target class “tench”.** $l_{1.5}$ -VCEs for Madry[19]+FT with varying radii for 3 images for the target class “tench”. They show again that classifier has learned to associate the spurious features of “human hands” or “human faces” with this class.

“Human hands” in the training set of “tench” in ImageNet



Fig. 25: A large fraction of training images of the class “tench” contains images of anglers presenting their “tench” into the camera. The model learns humans in particular human fingers as a spurious feature associated to “tench”.

E.4 Spurious features using other threat models

Here, we show, how the previously found spurious features look using l_2 - and l_1 -VCEs. l_2 -VCEs allow as well to see indicated spurious features in all three classes, while l_1 -VCEs have many artefacts and do not seem to be useful for this task.




Fig. 26: Spurious feature of watermarks for the target classes “granny smith”, “white shark”, and “tench”. l_2 -VCEs for Madry[19]+FT with varying radii for 3 images for the target classes “granny smith”, “white shark”, and “tench”. As $l_{1.5}$ -VCEs, l_2 -VCEs display spurious features of the classifier.




Fig. 27: **Spurious feature of watermarks for the target classes “granny smith”, “white shark”, and “tench”.** l_1 -VCEs for Madry[19]+FT with varying radii for 3 images for the target classes “granny smith”, “white shark”, and “tench”. Unlike $l_{1.5}$ -VCEs, l_1 -VCEs do not show spurious features of the classifier and have many artefacts.

F VCEs using guided diffusion process with regularization

Diffusion processes [17] have been used to generate VCEs on natural images when a causal structure is available [50]. We extend such approach to generic classifiers by leveraging the method of [2]: they show that using a diffusion process, guided by the CLIP model [46], with regularization is effective in a variety of text-driven image generation tasks like object or background replacement, object editing [44]. In particular, in [2], sampling exploits a hand-crafted mask, and is guided by image-text consistency according to the CLIP model and an l_2 -regularization term which ensures consistency with the background (the strength of such terms is controlled by the parameters λ_g and λ_2 respectively).

To apply such approach to our setting, we modify the following components:

(i) we use $l_{1.5}$ regularization instead of l_2 to have sparser changes and to be consistent with our proposed approach, (ii) apply regularization on the whole image, while [2] did so outside the mask only, and (iii) do not use the mask and operate on the whole image. Moreover, to generate VCEs for arbitrary classifiers which do not take text prompt as CLIP, we guide the sampling process by setting

$$\mathcal{L} \leftarrow \lambda_g \log \hat{p}(k|\hat{x}_0) - \lambda_{1.5} \|\hat{x}_0 - x\|_{1.5}^{1.5} - \lambda_{\text{LPIPS}} \text{LPIPS}(\hat{x}_0, x), \quad (9)$$

as loss function in **Algorithm 1** of [2], where k represents the target class, \hat{x}_0 the output of the denoising step and x the original image (following the notation of [2]).

In Fig. 30 we show the VCEs obtained with this scheme for two images (which we already used in Fig. 1 and Fig. 4). In particular we test several combination parameters and repeat for two random seeds. For convenience of the reader for direct comparison to our $l_{1.5}$ -VCEs we show again the corresponding images of the main paper.

We can see that (i) even though for some seeds and images the resulting VCEs can look realistic, they are perceivable less sparse than $l_{1.5}$ -VCEs, oftentimes changing the image completely, or are not valid, in the sense that no features of the target class appear (ii) parameters are more difficult to tune as $\lambda_{1.5}$ and λ_g are both model but even worse also image dependent, e.g. for the Madry+FT model $\lambda_g = 100$ and $\lambda_{1.5} = 0.2$ works best for the “cougar \rightarrow cheetah”-VCE but does not yield any meaningful VCE for “fig \rightarrow pineapple” where now $\lambda_{1.5} = 0.025$ works best and for “coral reef \rightarrow valley” none of these parameters works well, and are thus much more difficult to control than the radius of our l_p -VCEs, and (iii) as their algorithm works not in the image space, and only projects to $[0, 1]^d$ in the end, some images have visible artefacts (black or red dots). Moreover, when zooming in one notices that their generated images are often slightly blurred. In particular, we would argue that our $l_{1.5}$ -VCE (shown at the top) induce much more subtle changes of the original image clearly visible for “fig \rightarrow pineapple”, where their VCEs do not preserve the background.








Fig. 28: $l_{1.5}$ -VCEs and VCEs generated using diffusion processes approach of [2] for CLIP and Madry[19]+FT model. Here for VCEs generated using diffusion processes we use $l_{1.5}$ regularization, to have sparser changes for CLIP model used in [2] and Madry[19]+FT model. In the columns we vary coefficient for $l_{1.5}$ regularization $\lambda_{1.5}$ and in the rows - coefficient for the CLIP and Madry[19]+FT guidance λ_g for the change “cougar \rightarrow cheetah” for the same seed.

$l_{1.5}$ -VCEs for the Madry[19]+FT model (ours)				
Original	$\epsilon_{1.5} = 50$	$\epsilon_{1.5} = 75$	$\epsilon_{1.5} = 100$	
				

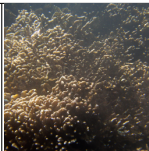
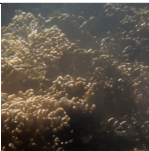



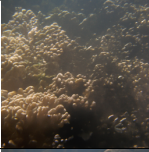


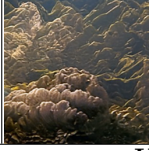


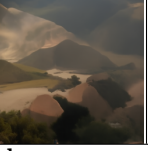
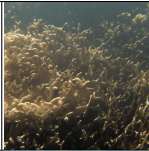
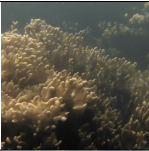
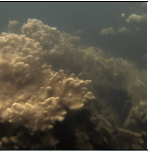
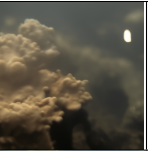
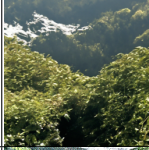
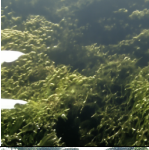

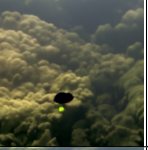




VCEs using diffusion processes approach of [2]				
$\lambda_{1.5}=0.0125$	$\lambda_{1.5}=0.025$	$\lambda_{1.5}=0.05$	$\lambda_{1.5}=0.1$	$\lambda_{1.5}=0.2$
Using CLIP model				
$\lambda_g=10$				
$\lambda_g=100$				
$\lambda_g=1000$				
Using Madry[19]+FT model				
$\lambda_g=10$				
$\lambda_g=100$				
$\lambda_g=1000$				

Fig. 29: $l_{1.5}$ -VCEs and VCEs generated using diffusion processes approach of [2] for CLIP and Madry[19]+FT model. Here for VCEs generated using diffusion processes we use $l_{1.5}$ regularization, to have sparser changes for CLIP model used in [2] and Madry[19]+FT model. In the columns we vary coefficient for $l_{1.5}$ regularization $\lambda_{1.5}$ and in the rows - coefficient for the CLIP and Madry[19]+FT guidance λ_g for the change “coral reef \rightarrow valley” for the same seed.

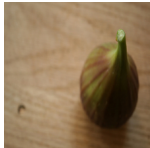
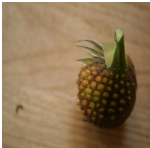




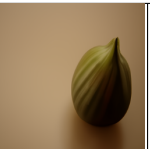



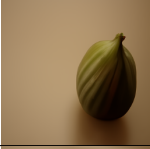
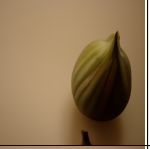





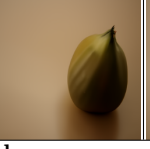









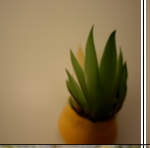






$l_{1.5}$ -VCEs for the Madry[19]+FT model (ours)					
Original	$\epsilon_{1.5} = 50$		$\epsilon_{1.5} = 75$	$\epsilon_{1.5} = 100$	
					
VCEs using diffusion processes approach of [2]					
	$\lambda_{1.5}=0.0125$	$\lambda_{1.5}=0.025$	$\lambda_{1.5}=0.05$	$\lambda_{1.5}=0.1$	$\lambda_{1.5}=0.2$
Using CLIP model					
$\lambda_g=10$					
$\lambda_g=100$					
$\lambda_g=1000$					
Using Madry[19]+FT model					
$\lambda_g=10$					
$\lambda_g=100$					
$\lambda_g=1000$					

Fig. 30: $l_{1.5}$ -VCEs and VCEs generated using diffusion processes approach of [2] for CLIP and Madry[19]+FT model. Here for VCEs generated using diffusion processes we use $l_{1.5}$ regularization, to have sparser changes for CLIP model used in [2] and Madry[19]+FT model. In the columns we vary coefficient for $l_{1.5}$ regularization $\lambda_{1.5}$ and in the rows - coefficient for the CLIP and Madry[19]+FT guidance λ_g for the change “fig \rightarrow pineapple” for the same seed. The empty square means that for this seed and this combination the method has encountered a numerical instability.

G User study examples

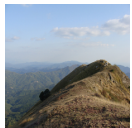



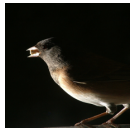


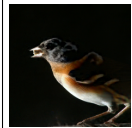
Original	APGD, l_1	AFW, $l_{1.5}$	APGD, l_2
promontory:0.60 	r.:0.00, m.:0.56, s.:0.39 	r.:0.72, m.:1.00, s.:0.78 	r.:0.94, m.:0.94, s.:0.83 
hummingb.:0.35 	r.:0.61, m.:0.72, s.:0.61 	r.:0.89, m.:0.89, s.:0.72 	r.:0.78, m.:0.78, s.:0.61 

Fig. 31: **Best rated images from the user study in blue.** l_p -VCEs for $p \in \{1, 1.5, 2\}$ for the change “promontory \rightarrow volcano” (top row) and “hummingbird \rightarrow brambling” (bottom) for Madry[19]+FT, where the images with blue frame are those for which users answered yes at the same time on **realism** (r), **meaningful** (m), **subtle** (s) questions introduced in App. G most frequently. For each VCE we show the difference to the original image and proportions how often each of the three questions is answered with yes .









Original	APGD, l_1	AFW, $l_{1.5}$	APGD, l_2
paddle:0.11 	r.:0.06, m.:0.17, s.:0.06 	r.:0.28, m.:0.11, s.:0.44 	r.:0.39, m.:0.11, s.:0.5 
lakeside:0.17 	r.:0.06, m.:0.50, s.:0.44 	r.:0.28, m.:0.94, s.:0.17 	r.:0.06, m.:0.00, s.:0.11 

Fig. 32: **Worst rated images from the user study in red.** l_p -VCEs for $p \in \{1, 1.5, 2\}$ for the change “paddle \rightarrow bearskin” (top row) and “lakeside \rightarrow geyser” (bottom) for Madry[19]+FT, where the images with red frame are those for which users have answered least frequently yes to the **realism** (r), **meaningful** (m), nor **subtle** (s) questions introduced in App. G. For each VCE we show the difference to the original image and proportions how often each of the three questions is answered with yes .

Here we provide two examples for the l_p -VCEs ($p \in \{1, 1.5, 2\}$) that according to the user study Sec. 2.3 were the best (users have answered yes to all three questions most frequently) in Fig. 31 and two that were the worst (users have answered yes to all three questions least frequently) in Fig. 32. The two worst examples show qualitatively why l_1 - and l_2 -VCEs might lead to undesired behaviors. In fact, they introduce changes which are either too localized and with intense colors (l_1 -VCEs) or cover the entire image (l_2 -VCEs).

H Randomly selected VCEs

In Fig. 34 we show randomly selected pairs of original images and $l_{1.5}$ -VCEs where the target class is chosen randomly from the same WordNet clusters described previously. Nevertheless several of the chosen target classes are close to impossible to realize with the given budget.

For CIFAR10 this is different as due to the lower image resolution of 32×32 $l_{1.5}$ -VCEs are even possible for quite distinct pairs of classes, see Fig. 33. Note that here the target class is the second most likely class predicted by the classifier for the original image. The model is here again GU+FT, that is the SOTA l_2 -robust model of [24] fine-tuned [16] for multiple-norm robustness.




Fig. 33: **Random CIFAR10 images together with their $l_{1.5}$ -VCE for $\epsilon = 6$** Target class is the second most probable class predicted by the classifier for this image. We see that for CIFAR10 due to the smaller image resolution the budget is sufficient to produce realistic counterfactuals even for pairs of true and target class which are distinct e.g. “deer” and “bird” or “ship and “cat”. However, there are cases where with the given budget the target class cannot be reached respectively there are artefacts remaining from the original image. On the other hand one can note that the changes are often quite subtle and most of the time inserted at the correct position in the image.

Original	$\epsilon_{1.5}=50$	$\epsilon_{1.5}=75$	$\epsilon_{1.5}=100$	Original	$\epsilon_{1.5}=50$	$\epsilon_{1.5}=75$	$\epsilon_{1.5}=100$
magnetic compass: 0.58	→ digital watch: 0.15	→ digital watch: 0.84	→ digital watch: 0.99	sock: 0.54	→ sock: 1.00	→ sock: 1.00	→ sock: 1.00
bearskin: 1.00	→ sombrero: 0.00	→ sombrero: 0.58	→ sombrero: 0.98	flagpole: 0.72	→ crutch: 0.02	→ crutch: 0.80	→ crutch: 0.99
eel: 0.05	→ hourglass: 0.97	→ hourglass: 1.00	→ hourglass: 1.00	soccer ball: 1.00	→ tennis ball: 0.00	→ tennis ball: 0.07	→ tennis ball: 0.57
goldfish: 0.67	→ lionfish: 0.92	→ lionfish: 1.00	→ lionfish: 1.00	spindle: 0.48	→ flute: 0.44	→ flute: 0.80	→ flute: 0.95

Fig. 34: **Random selection of $l_{1.5}$ -VCEs on ILSVRC2012 validation images for Madry[19]+FT.** The target class is randomly chosen from other classes that are contained in the same WordNet cluster. We note that sometimes it is not possible to transform an image into an image from the target in class with the given budget. This can for example be observed from the soccer ball which mostly contains either black or white pixels. To turn this image into a tennis ball, the attack has to turn those pixel values that lie in the corners of the RGB-cube into yellow, which is not possible with an $l_{1.5}$ budget of 100. We note that even though the soccer ball to tennis ball example can be regarded as failure cases, as it does not visually transform the image into the target class, it is important to note that the end confidence is relatively small. After calibration, it is typically possible to achieve a confidence of 97.0% or higher on valid VCEs whereas the confidence on those two images stays below 80.0% even for the largest radius. For other images, like the compass to digital watch, we can see that class-specific features like the numbers from a digital watch appear and integrate well with the remaining image, even though the budget is not sufficiently large to completely change the image subject. The flagpole to crutch VCE shows a ghosting artifact that can appear if the attack is not able to integrate the target object into the image. In this case, it paints the crutch on top of the flagpole as there is no reasonable way to integrate it with either the lighthouse or the sea in the background.

I Comparing threat models

For the paper, the radii were chosen such that on average the confidence of the classifier in the target class is high and meaningful changes are visible and the perturbation budget is comparable (the avg ℓ_2 -radius of the $\ell_{1.5}$ -VCE (we use $\epsilon_{1.5} = 50$) is 11.0 which is very close to $\epsilon_2 = 12$ used for ℓ_2 -VCEs). As confirmed by the user study ℓ_2 -VCEs are inferior to $\ell_{1.5}$ -VCE. An example VCE is shown below for different radii for $\ell_1/\ell_{1.5}/\ell_2$ where the ℓ_2 -VCEs either do not show meaningful changes for the target class “geyser” on the background or are far away from the original image. The ℓ_1 -VCEs show color artefacts and changes are too sparse.

