

Diffusion-based Visual Counterfactual Explanations - Towards Systematic Quantitative Evaluation [★]

Philipp Väh^{1,2}[0000-0002-8247-7907],
Alexander M. Frühwald¹[0009-0007-5347-6239],
Benjamin Paassen²[0000-0002-3899-2450], and
Magda Gregorova¹[0000-0002-1285-8130]

¹ Center for Artificial Intelligence and Robotics, Technical University of Applied Sciences Würzburg-Schweinfurt, Franz-Horn-Straße 2, Würzburg, Germany

{philipp.vaeth,alexander.fruehwald,magda.gregorova}@thws.de

² Bielefeld University, Universitätsstraße 25, Bielefeld, Germany

{bpaassen}@techfak.uni-bielefeld.de

Abstract. Latest methods for visual counterfactual explanations (VCE) harness the power of deep generative models to synthesize new examples of high-dimensional images of impressive quality. However, it is currently difficult to compare the performance of these VCE methods as the evaluation procedures largely vary and often boil down to visual inspection of individual examples and small scale user studies. In this work, we propose a framework for systematic, quantitative evaluation of the VCE methods and a minimal set of metrics to be used. We use this framework to explore the effects of certain crucial design choices in the latest diffusion-based generative models for VCEs of natural image classification (ImageNet). We conduct a battery of ablation-like experiments, generating thousands of VCEs for a suite of classifiers of various complexity, accuracy and robustness. Our findings suggest multiple directions for future advancements and improvements of VCE methods. By sharing our methodology and our approach to tackle the computational challenges of such a study on a limited hardware setup (including the complete code base), we offer a valuable guidance for researchers in the field fostering consistency and transparency in the assessment of counterfactual explanations.

Keywords: Explainability · XAI · Visual Counterfactual Explanations · VCE · Generative Modeling · Diffusion

1 Introduction

One of the greatest challenges of modern machine learning, especially deep neural networks, is maintaining the trust of the public and application experts in the

[★] Supported by the Bavarian HighTech agenda and the Würzburg Center for Artificial Intelligence and Robotics (CAIRO).

results and recommendations of these powerful technologies [6]. This is essential in sensitive and high-risk application domains such as medicine [21, 32], law enforcement [22], and finance [4]. To increase confidence in the reliability of these models and to enhance their transparency, a new research field called Explainable AI (XAI) has been established in recent years [2].

The goal of early XAI approaches has been to explain the internal workings of models, such as the gradient backpropagation dynamics [20, 27, 28]. Later the focus has shifted towards the explanation of model outputs rather than that of the internal structures of the model [5, 11, 23, 24]. For example, multiple popular explanation methods highlight features of the input data that contribute the most to the current prediction [20, 24]. By contrast, counterfactual explanations highlight necessary changes in the input to achieve a different output [31]. We focus in particular on counterfactual explanations for image data (visual counterfactual explanations; VCE) [11].

To create such VCEs, modern deep generative models are particularly promising as they are designed to synthesize new high-dimensional examples by sampling. The best quality in generative models to date has been achieved by the so-called diffusion models. Quite naturally, it has recently been also proposed to harness the power of the diffusion models for generating counterfactuals [1, 15].




Fig. 1: Hand-picked examples of pairs of original images and VCE attempts from our experiments. First row: desirable visual counterfactual explanations. Second row: examples with changes not sufficient to flip to the target class. Third row: method failures generating nonsensical outputs.

One of the difficulties of research in counterfactual methods is the lack of broadly accepted evaluation procedures. The state-of-the-art are visual inspection of individual explanations and small-scale user studies. While human expert judgments are crucial for model evaluation in domain context, they remain subjective and do not scale to typical large-scale image data sets. Further, the showcased examples are subject to cherry-picking (such as row one in Table 1).

Instead, we recommend a comprehensive evaluation on all data, including failure cases (such as rows two and three in the same table).

In this paper we contribute to the discussion on best evaluation practices for VCE methods. We review a number of previously proposed metrics and identify the most useful ones for assessing the properties of natural image VCEs (Section 3). Using these, we systematically quantitatively evaluate recent, diffusion-based, generative VCE methods [1, 15] across a range of robust and non-robust classifiers (Section 5). We complement this by an ablation of critical design choices that we identified as promising candidates for future improvements (Section 2.3). We provide access to our streamlined implementation of the VCE methods containing the ablation options as well as additional classifiers not present in the original code³. In line with the best reproducibility practices, this codebase also includes all our scripts to execute the evaluation, replicate our results and recreate all the summary tables presented in this paper. To foster transparency, we also provide access to our experimental tracking⁴. Finally, we highlight the challenges related to setting up such an evaluation exercise on a conservative hardware configuration and, based on our experience, recommend an approach to tackle these (Section 4.1). We hope fellow researchers will benefit from these when setting up their own evaluation protocols.

2 Visual Counterfactual Explanations through Diffusion Models

In the following section we introduce the basic idea of diffusion models, specifically Denoising Diffusion Probabilistic Models (DDPM) [13], explain the concept of classifier guidance, and draw the connection to counterfactual generation.

2.1 Denoising Diffusion Probabilistic Models

The DDPM model is an encoder-decoder architecture, where the encoder is defined as a Markov chain linear Gaussian model $q(x_t|x_{t-1})$ progressively adding random Gaussian noise to an image x_0 with a pre-defined variance schedule β_1, \dots, β_T over the diffusion steps $x_1, \dots, x_t, \dots, x_T$ so that $x_T \sim \mathcal{N}(0, I)$. The model learns to revert the diffusion process via the decoder $p_\theta(x_{t-1} | x_t)$, gradually removing noise from the signal x_T in order to produce samples matching the original data x_0 . It can be shown [19], that the reverse process is also Gaussian so that the model boils down to learning the mean and covariance of the reverse process $p_\theta(x_{t-1} | x_t) = \mathcal{N}(\mu_\theta(x_t, t), \Sigma_\theta(x_t, t))$.

2.2 Classifier Guidance

The vanilla DDPM introduced in Section 2.1 enables synthesizing random samples from the learned underlying data distribution. However, to generate class-specific data points, such as VCEs, the reverse transition has to be adapted.

³ https://github.com/cairo-thws/DBVCE_eval

⁴ https://wandb.ai/fhws_cairo/DBVCE_eval

Given a classifier $p_\phi(y | x_t)$, the mean of the denoising transition can be shifted by the gradients of the classifier w.r.t x_t to yield class-conditional samples [8].

$$p_{\theta,\phi}(x_{t-1} | x_t, y) = \mathcal{N}(\mu_t, \Sigma_\theta(x_t, t))$$

$$\mu_t = \mu_\theta(x_t, t) + s \cdot \Sigma_\theta(x_t, t) \nabla_{x_t} \log [p_\phi(y | x_t)] \quad , \quad (1)$$

where s is a gradient scaling factor controlling the trade-off between sample diversity and class consistency.

2.3 Diffusion-based Counterfactual Explanations

We introduce two recent diffusion-based generative models for counterfactual explanations, DVCE [1] and DiME [15], focusing on the particular design choices related to adapting the diffusion models for XAI purposes. The methods assume the availability of a DDPM pre-trained over the same images as the classifier p_ϕ to be explained. Both of the methods propose to guide the learned decoder of the process via the classifier gradients similar to equation (1) to produce a realistic example classified by p_ϕ as the target counterfactual class. However, to remain visually near the original image, both DVCE and DiME begin the denoising not from x_T , but from $x_{T/2}$, such that much of the rough structure of the original image is maintained but finer details can still be adapted to the target class.

One difficulty in using the classifier guidance of [8] for VCEs of an arbitrary classifier p_ϕ is that p_ϕ cannot be expected to have been trained on noisy images. Accordingly, p_ϕ cannot be expected to perform well on noisy data x_t generated during the reverse diffusion process and provide sensible gradients $\nabla_{x_t} \log p_\phi(y | x_t)$ for the guidance. To tackle this issue, the authors in both the DVCE and DiME paper propose to replace the noisy example x_t in the input of the classifier by an approximation of the denoised image $\hat{x}_0 = x_\theta(x_t, t)$, where x_θ is the learned DDPM denoising function. We call this important deviation from the original formulation of the classifier guidance [8] the *x_0 -prediction* approach and explore its effects more closely in Section 5.

While the *x_0 -prediction* approach solves the mismatch between the classifier and the noised diffusion data distributions, it presents an additional challenge for the gradient-based guidance of the sampling process. In principle, it requires to back-propagate not only through the classifier p_ϕ but also through the diffusion process p_θ (from \hat{x}_0 to x_t) to obtain the gradients with respect to the noised x_t . This greatly increases the computational as well as the memory requirements for generating VCEs. We discuss the related technical challenges and our approach to tackle them with a fairly conservative hardware configuration in Section 4.1.

Another important design choice introduced in the DVCE paper is the *cone-projection* approach. The authors argue that gradients of non-robust classifiers are noisy, making the guided generation of semantically meaningful images particularly difficult. Therefore, they propose to use the gradients of an independent, pre-trained, robust classifier p_r , which are expected to be less noisy, to help in the guidance. This is achieved by replacing the gradients of the classifier to be explained $\nabla_{x_t} \log p_\phi$ in equation (1) with those of the robust classifier $\nabla_{x_t} \log p_r$,

projected onto a cone of 30° around the gradients $\nabla_{x_t} \log p_\phi$. This reduces the noise of the non-robust classifier gradients and steers the counterfactual generation towards more semantically meaningful changes. We discuss the effects of the *cone-projection* in more detail in Section 5.

3 Metrics

The desirable properties of counterfactual explanations can be broadly categorized into four axes: validity, closeness, realism and diversity [1, 14, 15, 25]. Though these are generally alluded to in the VCE literature, they are rarely systematically quantitatively evaluated. We describe these properties here below together with relevant metrics that can be used for their assessment. We employ these metrics in Section 5 to systematically and quantitatively evaluate sets of VCEs generated under varying conditions to document their usefulness and provide some novel insights into the functioning of the explored methods.

For the following, let $f(\cdot)$ be the classifier that shall be explained, X the set of images of class y for which we wish to create counterfactual explanations, y_t the target class label, $c_f(\cdot)$ the counterfactual generation method trying to flip the label to y_t , and $\mathbb{1}[\cdot]$ the indicator function.

3.1 Validity

A counterexample generated by $c_f(\cdot)$ is considered a valid counterfactual if the classifier assigns it to the counterfactual target class y_t .

A concept closely related to counterfactuals is that of adversarial examples [10]. These are both perturbations of the original data example in such a way, that the classifier changes its decision. There is, however, a significant difference between the two. While the goal of an adversarial attack is to fool the classifier without actually changing the class of the example, a counterfactual shall truly be a realistic example of the target class. In this paper we use the *misclassification* criterion [9] as the distinguishing property between the two concepts and illustrate the difference in Figure 2.

To explore the validity of the generated counterexamples, we propose to use a set of metrics based on the classification accuracy.

Target Accuracy, $TA = \frac{1}{|X|} \sum_{x \in X} \mathbb{1}[f(c_f(x)) = y_t]$, is the fraction of generated counterexamples classified by f to be within the target class y_t .




Fig. 2: Classifier class boundary versus true (oracle) class boundary and the interpretation thereof.

Original Accuracy, $OA = \frac{1}{|X|} \sum_{x \in X} \mathbb{1}[f(c_f(x)) = y]$, is the fraction of generated counterexamples classified by f to be within the original class y . For these examples, the image changes are not sufficient to flip the classification result.

For a multi-class data set, the generator may produce examples that the classifier attributes to neither the original y nor the target class y_t . The percentage of these is clearly $1 - TA - OA$.

To confirm that the generated counterexamples classified by f as y_t are really counterfactuals rather than just adversarial examples, we employ a committee of alternative classification models f_o , trained on the same data set, to serve as an oracle replacing human-expert assessment.

Oracle score, $OS = \frac{1}{|X|} \sum_{x \in X} \mathbb{1}[f(c_f(x)) = f_o(c_f(x))]$, is the percentage of generated examples that are classified the same by both the classifier f and the oracle f_o [14], as a cheap though not perfect proxy for human validity assessment.

Oracle Target Accuracy, $OTA = \frac{1}{|X|} \sum_{x \in X} \mathbb{1}[f_o(c_f(x)) = y_t]$, is the percentage of counterexamples that are classified by the oracle f_o to be within the target class y_t . This might include examples for which the classifier f does not flip the class to y_t but still show some changes that convince the oracle f_o .

3.2 Closeness

For a generated counterexample to be considered a counterfactual, it shall be as close as possible to the original example. Various forms of Minkowski metric $D_p = \frac{1}{|X|} \sum_{x \in X} \|x - c_f(x)\|_p$ (e.g. $p \in 1, 1.5, 2$) can be utilized to measure the distance.

A metric corresponding better to human visual perception is the Learned Perceptual Image Patch Similarity (LPIPS) [33]. Instead of comparing the image examples in the original pixel space, it compares activations of image patches in a pre-trained network. In our study we use the PyTorch implementation of TorchMetrics⁵.

3.3 Realism

VCE shall not just contain class-discriminative artifacts flipping the predicted class, but they shall be realistic examples of the target class. We propose to use the state-of-the-art metric for image generation quality, the Fréchet Inception Distance (FID) [12], suggested in some previous XAI studies [1, 15, 25, 29]. The FID compares the distribution of the Inception V3 model [30] image embeddings between a set of original $x \in X$ images and their generated counterparts $\hat{x} = c_f(x)$. We use the TorchMetrics implementation with 64 features.

3.4 Diversity

A property often disregarded by VCE methods [25] is the ability to generate multiple valid counterfactuals for a single analyzed image. Differences within the set

⁵ <https://torchmetrics.readthedocs.io>

of these examples can provide richer information to the user and improve the understanding of the results. To measure diversity, we propose to use pairwise LPIPS [14, 33] between the set of generated VCEs. Since the diffusion-based methods we analyze in this paper do not natively generate sets of counterfactuals, there is no efficient way to produce these. Rerunning the experiments multiple times with different random seeds should produce some variations in the data and it would be interesting to evaluate their diversity. Such large-scale experimentation is, however, beyond the capacity of our limited hardware setup and we leave it for future work.

4 Experimental Setup

In our experiments, we systematically and quantitatively evaluate diffusion-generated VCEs (see Section 2.3) across a suite of classifiers. We analyze similar data as those presented in the DVCE paper but, rather than manually selecting a handful of examples for visual inspection, we calculate the relevant metrics (Section 3) across a representative set of generations.

The generative diffusion model has been pre-trained on the unlabeled ImageNet (2012) data set [26]. Same as the DVCE paper we use the OpenAI implementation of guided diffusion⁶. We choose six classifiers pre-trained on the labeled version of the data set summarized in Table 1 complemented by a RandomNet, which is an AlexNet architecture with randomly initialized weights.

classifier	Madry	AlexNet	UnetE	ConvNeXt	SwinTFL	SIMCLR	RandomNet
Name	MNR-RN50	AlexNet	UNet	ConvNeXt-L	Swin-L	SIMCLR	AlexNet
Parameters	25M	62M	54M	350M	197M	94M	62M
Accuracy	57.9%	63.3%	NA	87.8%	87.3%	74.2%	0%
Robustness	✓	✗	(✓)	✗	✗	✗	✗
Used in DVCE	✓	✗	✗	✓	✓	✗	✗
Published	[3]	[16]	[8]	[18]	[17]	[7]	-

Table 1: Classifiers included in our study. NA: Not reported.

In our analysis, we focus on the six source and eleven target classes presented in the DVCE paper: source - *cheetah*, *chesapeake bay retriever*, *mashed potato*, *coral reef*, *alp*, *burrrito*, target - *leopard*, *tiger*, *golden retriever*, *labrador retriever*, *guacamole*, *carbonara*, *alp*, *cliff*, *volcano*, *pizza*, *burrrito*. For each of the source classes we generate counterfactuals for all 50 original images from the validation data set (6 x 50 = 300 source images). For each of the 50 source images per source class, we generate a counterfactual towards each of the 11 target classes.

⁶ <https://github.com/openai/guided-diffusion>

We thus obtain a total of 3300 (= 6 x 50 x 11) synthesized examples split into 66 (= 6 x 11) source-target groups.

For each of the source classes there is a class within the target set identified by the DVCE paper as the most relevant for counterfactuals. For example, for cheetah it is the leopard, for mashed potato it is the carbonara. These have been identified based on the WordNet similarity of the class labels. In our view, these choices of “ideal” source-target pairs are rather arbitrary and other source-target pairs may be required in practice. Accordingly, while most of our evaluation in Section 5 is performed over these “ideal” pairs, we also explore the performance when other classes are chosen for the counterfactual generations.

We explore the effects of two critical design choices explained in Section 2.3. Firstly, we let the gradient guidance be based on classifying the point x_t either with the x_0 -prediction (using $\hat{x}_0(x_t, t)$) or without (as in equation (1)). Secondly, we use the *cone-projection* of the robust classifier gradients or we drop it and rely solely on the gradient directions of the classifier that shall be explained. To cover these four architectural combinations, we need to repeat all our experiments four times leading to a total of 13’200 generated counterexamples, organized into the 66 source-target groups (split by the ablation choices). All these experiments have to be repeated for each classifier in our study, totaling 92’400 (= 13’200 x 7) diffusion-based generations.

4.1 Challenges

All experiments are conducted on a single compute node with four Nvidia A100 GPUs with 80 GB of video memory each. We re-implemented the x_0 -prediction and the *cone-projection* to fit our evaluation procedure and make multiple adjustments (e.g. CPU offloading, deleting gradients, clearing CUDA cache) to streamline the code and its execution (see the released code base).

A major challenge for our experiments are the GPU/VRAM requirements. The classifiers as well as the diffusion model are large neural networks with millions of parameters⁷ (Table 1) requiring several gigabytes of VRAM. The highest memory trace is, however, due to the gradient calculation and especially that of the diffusion model due to the x_0 -prediction approach. This is further exacerbated by the *cone-projection* which needs to use two classifiers at a time.

We rely on the DataParallel functionality of PyTorch and spread the generation of the counterfactuals for a single batch of 50 examples (from a single source class) across all four GPUs. Very large models such as UnetE and ConvNeXt had to be placed onto a dedicated GPU while the remaining three GPUs were used in parallel for the diffusion model. This setup required further adaptations in the code (available within our released code base).

Given the high memory requirements, further parallelization was not possible and we had to run the 28 experiments sequentially (4 ablation options x 7 classifiers). The total runtime for these was approximately two days (wall-clock time). Some evaluation metrics could be calculated right after the generation step as

⁷ According to [13], the DDPM model has 554M parameters.

averages across the batch of generated images. Many more (oracles, LPIPS, FID) had to be calculated in a post-processing step as they required loading additional models into memory.

5 Experimental Results

In this section, we evaluate the generated VCEs using the metrics defined in Section 3. We calculate the metrics for each source-target group (indicated in Section 3 as the set X) and, for the final presentation, calculate averages and standard deviations across the seven source classes.

5.1 Minimal Set of Metrics

In Table 2 we report the complete set of metrics in the base setup of the experiments: “ideal” source-target pairs (as defined by [1]), with x_0 -prediction, and with cone-projection. We first review these metrics and reduce the list to a recommended minimal set for further evaluation.

classifier metric	Madry	AlexNet	UnetE	ConvNeXt	SwinTFL	SIMCLR	RandomNet
TA \uparrow	0.84 \pm 0.1	0.99 \pm 0.01	0.99 \pm 0.01	0.95 \pm 0.04	0.9 \pm 0.08	0.45 \pm 0.29	0.0 \pm 0.0
OA \downarrow	0.07 \pm 0.07	0.0 \pm 0.0	0.01 \pm 0.01	0.02 \pm 0.02	0.05 \pm 0.03	0.0 \pm 0.0	0.01 \pm 0.01
OS Madry \uparrow	-	0.32 \pm 0.14	0.36 \pm 0.23	0.29 \pm 0.17	0.32 \pm 0.15	0.05 \pm 0.03	0.0 \pm 0.0
OTA Madry \uparrow	-	0.32 \pm 0.14	0.36 \pm 0.23	0.28 \pm 0.16	0.3 \pm 0.16	0.2 \pm 0.13	0.21 \pm 0.16
OS AlexNet \uparrow	0.72 \pm 0.13	-	0.48 \pm 0.21	0.36 \pm 0.13	0.42 \pm 0.13	0.02 \pm 0.03	0.0 \pm 0.01
OTA AlexNet \uparrow	0.72 \pm 0.14	-	0.48 \pm 0.21	0.36 \pm 0.13	0.41 \pm 0.15	0.15 \pm 0.05	0.23 \pm 0.11
OS RandomNet \downarrow	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	-
OTA RandomNet \downarrow	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	-
OS UnetE \uparrow	0.72 \pm 0.08	0.43 \pm 0.15	-	0.43 \pm 0.19	0.48 \pm 0.19	0.06 \pm 0.04	0.0 \pm 0.01
OTA UnetE \uparrow	0.76 \pm 0.07	0.43 \pm 0.15	-	0.4 \pm 0.19	0.45 \pm 0.2	0.17 \pm 0.09	0.22 \pm 0.13
OS ConvNeXt \uparrow	0.72 \pm 0.1	0.38 \pm 0.14	0.57 \pm 0.22	-	0.69 \pm 0.13	0.03 \pm 0.04	0.0 \pm 0.0
OTA ConvNeXt \uparrow	0.69 \pm 0.13	0.38 \pm 0.14	0.57 \pm 0.22	-	0.64 \pm 0.17	0.08 \pm 0.05	0.17 \pm 0.12
OS SwinTFL \uparrow	0.7 \pm 0.08	0.39 \pm 0.23	0.53 \pm 0.23	0.56 \pm 0.22	-	0.01 \pm 0.02	0.0 \pm 0.0
OTA SwinTFL \uparrow	0.66 \pm 0.12	0.39 \pm 0.23	0.53 \pm 0.23	0.54 \pm 0.23	-	0.07 \pm 0.06	0.14 \pm 0.13
D ₁ (m) \downarrow	1.49 \pm 0.21	1.23 \pm 0.22	1.28 \pm 0.21	1.2 \pm 0.22	1.22 \pm 0.22	1.16 \pm 0.24	1.14 \pm 0.23
D _{1.5} (k) \downarrow	9.08 \pm 1.16	7.55 \pm 1.22	7.81 \pm 1.12	7.31 \pm 1.24	7.43 \pm 1.24	7.08 \pm 1.33	6.97 \pm 1.27
D ₂ (k) \downarrow	0.77 \pm 0.09	0.64 \pm 0.1	0.66 \pm 0.08	0.62 \pm 0.1	0.63 \pm 0.1	0.6 \pm 0.1	0.59 \pm 0.1
LPIPS \downarrow	0.38 \pm 0.04	0.32 \pm 0.05	0.32 \pm 0.03	0.29 \pm 0.04	0.29 \pm 0.05	0.34 \pm 0.07	0.26 \pm 0.03
FID \downarrow	1.87 \pm 0.77	1.16 \pm 0.83	0.93 \pm 0.45	0.8 \pm 0.4	0.94 \pm 0.39	3.16 \pm 2.95	1.18 \pm 0.52

Table 2: Comprehensive quantitative evaluation of the VCEs in the base setup.

The critical metric for counterfactual validity is the TA – a generated example is not a counterfactual if it has not flipped the classifier decision. OA can provide additional information for failure cases but we deem it more useful for early development stages of novel methods since reasonably well-tuned counterfactual methods shall be able to move out of the original class. We therefore propose to drop OA from the minimal list.

We argued in Section 3 for the use of oracle classifiers f_o to discriminate counterfactuals from adversarial attacks. Since each single classifier may be prone

to error, we further extend the concept into a set of classifiers as a committee of oracles. A possible alternative would be merging all these into a single oracle in an ensemble-like manner. The previously suggested OS [14] is in our view not fit for purpose as it measures agreements and disagreements irrespective of the class. We therefore recommend the simpler OTA which focuses directly on the target class.

The Minkowski distance metrics D_p are, on occasion, used as regularization terms for both counterfactual and adversarial generations. Such per-pixel distances, however, do not well align with human perception [33] and are therefore of little use for measuring the quality of *visual* counterfactual explanations. We propose to exclude them from the shortlist and use the LPIPS [33] instead.

FID [12] is the only metric in our list for measuring realism, hence we recommend its inclusion.

Overall, we recommend the TA and OTA to evaluate validity, the LPIPS for closeness, and the FID measure for realism of VCEs. Accordingly, we report these metrics in the remainder of this section.

5.2 Evaluation of Diffusion-based VCEs

We structure our analysis around five hypotheses about the functioning of the method.

classifier metric	Madry	AlexNet	UnetE	ConvNeXt	SwinTFL	SIMCLR	RandomNet
TA \uparrow	0.84 ± 0.1	0.99 ± 0.01	0.99 ± 0.01	0.95 ± 0.04	0.9 ± 0.08	0.45 ± 0.29	0.0 ± 0.0
OTA Madry \uparrow	-	0.32 ± 0.14	0.36 ± 0.23	0.28 ± 0.16	0.3 ± 0.16	0.2 ± 0.13	0.21 ± 0.16
OTA AlexNet \uparrow	0.72 ± 0.14	-	0.48 ± 0.21	0.36 ± 0.13	0.41 ± 0.15	0.15 ± 0.05	0.23 ± 0.11
OTA UnetE \uparrow	0.76 ± 0.07	0.43 ± 0.15	-	0.4 ± 0.19	0.45 ± 0.2	0.17 ± 0.09	0.22 ± 0.13
OTA ConvNeXt \uparrow	0.69 ± 0.13	0.38 ± 0.14	0.57 ± 0.22	-	0.64 ± 0.17	0.08 ± 0.05	0.17 ± 0.12
OTA SwinTFL \uparrow	0.66 ± 0.12	0.39 ± 0.23	0.53 ± 0.23	0.54 ± 0.23	-	0.07 ± 0.06	0.14 ± 0.13
LPIPS \downarrow	0.38 ± 0.04	0.32 ± 0.05	0.32 ± 0.03	0.29 ± 0.04	0.29 ± 0.05	0.34 ± 0.07	0.26 ± 0.03
FID \downarrow	1.87 ± 0.77	1.16 ± 0.83	0.93 ± 0.45	0.8 ± 0.4	0.94 ± 0.39	3.16 ± 2.95	1.18 ± 0.52

Table 3: Quantitative evaluation of the VCEs in the base setup.

Hypothesis 1: The diffusion-based counterfactual method in its base setup (for “ideal” source-target pairs, with *cone-projection* and *x₀-prediction*) produces valid, close and realistic VCEs even for non-robust classifiers of diverse model capacity. Table 3 summarizes the performance for the seven classification algorithms in our study. As documented by high TA, most classifiers achieve high proportions of correctly classified counterexamples. Zero validity of RandomNet is in line with the inability of this untrained network to classify the images correctly. Somewhat surprisingly, the validity of SIMCLR counterfactuals is also fairly low, around 45% with very high standard deviation between the 6 original classes. This suggests that, for some of the original classes, SIMCLR manages to generate high proportions of

counterexamples within the target class while for others it dramatically fails. We speculate that this failure is related to the self-supervised nature of the SIMCLR. Unlike the other models, its hidden features are not attuned to the specific classification task that is being explained by the counterfactual method. The effects of this difference in the feature learning on the counterfactual generation, however, require more detailed investigation.

On the other hand, the oracle accuracy OTA for most classifiers and oracles is well below 50%. We can conclude that the generative process is successfully guided by the classifier to be explained but, rather than valid counterfactuals, produces adversarial examples which are not of the expected target class (see examples in Table 1). The exception here is Madry, the only robust classifier in our set, which achieves around 70% oracle accuracy, actually flipping the images to the target class.

While these conclusions contradict our Hypotheses 1, the low OTA scores shall not be interpreted unfavorably for the counterfactual method. The method shall help the users establish trust in the classifiers by explaining their behavior. From these metrics, we can see that most classifiers in our set rely on features that are not truly consequential for the class. By modifying these within the images, the classifiers flip their decisions but fail to truly shift the image conceptually to the target class. This shall lead into reducing the trust of the user in the classifier results (despite possibly high classification accuracy).

The closeness measured by LPIPS ranges between 0.29 to 0.32 for all classifiers except Madry, which leads to rather higher distance (0.38) between the origin-target pairs on average. This indicates that smaller perceptual changes may be sufficient to fool the classifier under study, but not sufficient to genuinely flip the class (which only Madry achieves).

We use the FID scores to measure the realism of the counterexamples. It is the ConvNeXt guidance that generates examples which are the most coherent with the underlying probability distribution of the target class. Since this method also achieves rather small LPIPS, it leads us to belief that only meaningful changes get introduced into the images but these are either too small or not semantically coherent to flip the decision of the oracles. Madry, in contrast, suffers from much higher mismatch between the distributions of the true and generated images. This backs the LPIPS conclusions: Madry guidance introduces more important changes into the images sometimes pushing them further away from the underlying data distribution. SIMCLR has by far the largest FID as it introduces changes that are often not coherent with the target class at all (low TA and OTA). This reconfirms the failure of the counterfactual method on this self-supervised classifier.

Hypothesis 2: No meaningful changes are introduced into the images when generating counterexamples for a random classifier. The RandomNet results in Table 3 confirm that, as expected, the method produces no valid counterfactuals (zero TA). Nevertheless, some small changes are introduced (low LPIPS) keeping the images realistic (low FID) but, surprisingly, switching the class of almost 20% of the images (OTA scores). This contradicts

our Hypotheses 2. Given that the classifier is random and its gradients toward the target class are therefore meaningless, this is a rather undesirable behavior which can be explained by the robust classifier feeding the method through the *cone-projection* with semantically meaningful gradient directions. This could, in practice, be resolved by introducing a rejection step into the counterfactual method, dropping all generated examples that are not re-classified by the classifier to be explained into the target class. We recommend to adopt this step as a simple fix into the VCE methods in the future. Based on the above analysis, we advocate for systematic inclusion of similar edge-case experiments into the evaluation of XAI methods as these can lead to useful unexpected insights.

classifier metric	Madry	AlexNet	UnetE	ConvNeXt	SwinTFL	SIMCLR	RandomNet
TA \uparrow	0.00 \pm 0.00	-0.01 \pm 0.02	-0.04 \pm 0.05	-0.12 \pm 0.06	-0.19 \pm 0.12	0.12 \pm -0.03	0.00 \pm 0.00
OTA Madry \uparrow	-	-0.29 \pm -0.11	-0.31 \pm -0.16	-0.27 \pm -0.14	-0.29 \pm -0.14	-0.19 \pm -0.11	-0.20 \pm -0.15
OTA AlexNet \uparrow	0.00 \pm 0.01	-	-0.27 \pm -0.05	-0.30 \pm -0.07	-0.30 \pm -0.07	-0.12 \pm -0.02	-0.19 \pm -0.07
OTA UnetE \uparrow	0.00 \pm 0.01	-0.37 \pm -0.13	-	-0.35 \pm -0.10	-0.39 \pm -0.14	-0.15 \pm -0.08	-0.21 \pm -0.11
OTA ConvNeXt \uparrow	0.00 \pm -0.01	-0.30 \pm -0.09	-0.27 \pm -0.07	-	-0.36 \pm -0.04	-0.06 \pm -0.04	-0.16 \pm -0.12
OTA SwinTFL \uparrow	-0.01 \pm 0.00	-0.29 \pm -0.16	-0.26 \pm -0.08	-0.34 \pm -0.12	-	-0.06 \pm -0.05	-0.14 \pm -0.14
LPIPS \downarrow	0.00 \pm 0.00	-0.04 \pm 0.00	-0.02 \pm 0.00	-0.04 \pm 0.00	-0.04 \pm -0.01	-0.01 \pm 0.01	-0.04 \pm 0.00
FID \downarrow	0.01 \pm 0.01	-0.12 \pm -0.04	0.42 \pm -0.04	0.17 \pm 0.34	-0.07 \pm 0.37	1.18 \pm 1.66	0.36 \pm -0.06

Table 4: Change in metrics as compared to Table 3: no *cone-projection*.

Hypothesis 3: Cone projection is the critical factor for generating valid counterfactuals rather than just adversarial examples. In Table 4, we report the changes in the metrics (increases and decreases) as compared to Table 3 when the guidance is based solely on the gradients of the classifier to be explained without any robust classifier *cone-projection*.

The results for the robust classifier Madry remain unchanged in this scenario because *cone-projection* merely projects Madry’s gradients onto its own cone (thus leaving them unchanged). Similarly, the target accuracy of the random classifier also remains at zero. However, the OTA decrease corroborates our conclusions from the previous section on the undesirable effects of the robust classifier *cone-projection* for this edge-case.

For the other classifiers, the target accuracy drops (slightly for the AlexNet and UnetE, rather more for ConvNetXt and SwinTFL). OTA drops significantly across the board meaning that though the counterexamples produced by the method without the *cone-projection* fool the classifiers to be explained, they are not valid counterfactuals as they do not de facto belong to the target class (low OTA). This confirms our Hypotheses 3 about the significance of the *cone-projection* step in preventing adversarial examples.

Similar to prior analysis, the behavior of SIMCLR stands out. It is the only classifier for which the TA improves when not using *cone-projection*. However, given that OTA drops as well, this again signifies that mostly adversarial examples get generated. From the increase in TA we surmise that the initial failure of the counterfactual method for SIMCLR reported in Table 3 may be due to large

disagreements between the SIMCLR and the robust classifier in the desirable gradient directions. This will lead to frequent projections of the gradients on the edge of the cone finding a compromise direction which does not guide the process towards useful and meaningful changes.

classifier metric	Madry	AlexNet	UnetE	ConvNeXt	SwinTFL	SIMCLR	RandomNet
TA \uparrow	0.00 \pm 0.00	-0.08 \pm 0.07	0.00 \pm 0.00	-0.58 \pm 0.15	-0.55 \pm 0.18	-0.45 \pm -0.31	0.00 \pm 0.00
OTA Madry \uparrow	-	-0.06 \pm 0.01	-0.02 \pm 0.00	-0.06 \pm 0.00	-0.08 \pm 0.01	0.02 \pm 0.03	0.00 \pm 0.00
OTA AlexNet \uparrow	-0.01 \pm -0.02	-	-0.02 \pm -0.02	-0.07 \pm 0.00	-0.12 \pm -0.01	0.09 \pm 0.06	0.02 \pm 0.01
OTA UnetE \uparrow	0.00 \pm 0.02	-0.09 \pm -0.01	-	-0.15 \pm -0.06	-0.19 \pm -0.07	0.05 \pm 0.04	0.00 \pm 0.00
OTA ConvNeXt \uparrow	-0.02 \pm 0.00	-0.10 \pm 0.01	-0.01 \pm 0.00	-	-0.38 \pm -0.01	0.07 \pm 0.04	-0.03 \pm -0.02
OTA SwinTFL \uparrow	0.01 \pm 0.00	-0.10 \pm -0.04	0.04 \pm 0.00	-0.30 \pm -0.06	-	0.07 \pm 0.08	-0.01 \pm -0.01
LPIPS \downarrow	0.00 \pm 0.00	-0.03 \pm -0.01	0.00 \pm 0.00	-0.03 \pm 0.00	-0.03 \pm -0.01	-0.08 \pm -0.04	0.00 \pm 0.00
FID \downarrow	-0.13 \pm -0.03	-0.18 \pm -0.19	-0.03 \pm 0.00	0.37 \pm 0.17	0.27 \pm 0.21	-1.96 \pm -2.66	-0.03 \pm -0.01

Table 5: Change in metrics as compared to Table 3: no x_0 -prediction.

Hypothesis 4: The x_0 -prediction enables the counterfactual generation on non-robust classifiers. We argued in Section 2.3 that the x_0 -prediction enables the counterfactual generation for classifiers not trained on noisy images and therefore not able to correctly classify the noisy samples of the intermediate diffusion generative process.

In Table 5 we see that the x_0 -prediction has a negligible effect for Madry and UnetE classifiers which were both trained on noisy images. For the random classifier, the effect is also almost zero, which is as expected and verifies our previous conclusion that the image changes in the standard setting are in fact caused by the *cone-projection*. The three non-robust classifiers ConvNeXt, SwinTFL and SIMCLR not trained on noisy images loose substantial target accuracy and some oracle target accuracy. This confirms our Hypotheses 4. Interestingly, the TA and OTA scores of AlexNet decrease only slightly. We hypothesize that this is due to the low model capacity of AlexNet which forces it to learn only a low number of class-discriminative features. This makes the model, perhaps somewhat surprisingly, robust to noised images.

Hypothesis 5: Counterfactual generation for badly chosen target classes fails or produces adversarial attacks.

The evaluation in Table 6 is conducted over the 10 non-“ideal” source-target pairs for each of the 6 original classes. TA, but mainly OTA, decrease rather dramatically while the LPIPS and FID increase. This indicates that the method struggles to produce valid counterfactuals for these target classes. This is in line with our intuition, which suggests that it is fundamentally very difficult (and at the same time not very useful) to produce counterfactuals belonging to a class which is semantically very distant from the original class (e.g. from the animal world to food or physical geographical phenomena).

These results trigger the question to what extent VCEs for natural scene classification even make sense (e.g. multiple concepts may be present within single image, the combination of the minor concepts may be inconsistent with major

classifier metric	Madry	AlexNet	UnetE	ConvNeXt	SwinTFL	SIMCLR	RandomNet
TA \uparrow	-0.16 \pm 0.12	-0.03 \pm 0.03	-0.09 \pm 0.09	-0.23 \pm 0.12	-0.35 \pm 0.16	0.08 \pm -0.06	0.00 \pm 0.00
OTA Madry \uparrow	-	-0.14 \pm 0.05	-0.19 \pm -0.05	-0.16 \pm 0.00	-0.17 \pm 0.00	-0.12 \pm -0.01	-0.13 \pm -0.02
OTA AlexNet \uparrow	-0.23 \pm 0.10	-	-0.29 \pm -0.03	-0.23 \pm 0.04	-0.24 \pm 0.02	-0.11 \pm 0.02	-0.17 \pm 0.00
OTA UnetE \uparrow	-0.22 \pm 0.16	-0.22 \pm 0.04	-	-0.25 \pm -0.02	-0.28 \pm -0.03	-0.11 \pm 0.03	-0.15 \pm 0.01
OTA ConvNeXt \uparrow	-0.24 \pm 0.09	-0.16 \pm 0.04	-0.35 \pm -0.05	-	-0.34 \pm 0.03	-0.05 \pm 0.05	-0.12 \pm 0.01
OTA SwinTFL \uparrow	-0.21 \pm 0.11	-0.18 \pm -0.05	-0.33 \pm -0.05	-0.31 \pm -0.01	-	-0.04 \pm 0.03	-0.09 \pm 0.00
LPIPS \downarrow	0.05 \pm 0.00	0.06 \pm 0.00	0.03 \pm 0.00	0.02 \pm 0.00	0.04 \pm -0.01	0.02 \pm -0.01	0.02 \pm 0.00
FID \downarrow	0.46 \pm 0.68	0.41 \pm 0.25	0.61 \pm 0.45	0.46 \pm 0.37	0.56 \pm 0.67	0.34 \pm 1.36	0.44 \pm 0.33

Table 6: Change in metrics as compared to Table 3: non-“ideal” targets.

concept of the target class, etc.). At the same time, they clearly document that some target classes are more suitable for counterfactual generation than others. The choice of these lies in the hands of the analyst and for now there is little guidance on how this shall be achieved (DVCE used the WordNet similarity). In our view both these questions deserve further attention in future research.

6 Conclusion

In this work, we advocate for more rigorous, quantitative assessment of visual counterfactual explanations (VCEs) in order to foster consistency and transparency of evaluation procedures and thus advance the understanding of the strength and weaknesses of different VCE methods. We propose a framework and a set of metrics to be used, and we conduct an extensive study exploring the effects of critical design choices of diffusion-based VCE methods across a suite of classifiers. Our investigation of the results illustrates how such a systematic quantitative analysis can provide valuable insights into the functioning of the methods and suggest directions for future improvements. We hope that by sharing the results of our study and our experience with tackling the computational challenges, together with the complete code to replicate our results, we contribute to betterment of experimental and evaluation protocols in future reports on VCE methods.

An important limitation of our work is the lack of evaluation of the diversity in the counterfactual explanations. We briefly outlined a possible approach to such an evaluation in Section 3.4. However, given that generating multiple diverse counterfactuals for the same original image would require re-running the VCE method multiple times in sequence, the overall execution time proved to be prohibitively long for our study. Nevertheless, we are convinced that having access to a diverse set of counterfactuals should greatly improve the user understanding of the classifier’s decisions (as compared to a single counterfactual or a set of very similar ones). In our future work we will investigate a more efficient experimental setup for our hardware configuration that would allow generation and evaluation of diverse counterfactuals more rapidly. We also believe that the ability to generate diverse sets of counterfactuals is an important aspect of the VCE methods that deserves concentrated effort in future research.

References

1. Augustin, M., Boreiko, V., Croce, F., and Hein, M.: Diffusion Visual Counterfactual Explanations. *NeurIPS* (2022)
2. Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbadó, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., and Herrera, F.: Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion* 58, 82–115 (2020)
3. Boreiko, V., Augustin, M., Croce, F., Berens, P., and Hein, M.: Sparse visual counterfactual explanations in image space. *Pattern Recognition* (2022)
4. Cao, L.: Ai in finance: challenges, techniques, and opportunities. *ACM Computing Surveys* 55(3), 1–38 (2022)
5. Casalicchio, G., Molnar, C., and Bischl, B.: Visualizing the Feature Importance for Black Box Models. *ECML 11051*, 655–670 (2018)
6. Chatila, R., Dignum, V., Fisher, M., Giannotti, F., Morik, K., Russell, S., and Yeung, K.: “Trustworthy AI”. In: *Reflections on Artificial Intelligence for Humanity*. Ed. by B. Braunschweig and M. Ghallab. Cham: Springer International Publishing, 2021, pp. 13–39.
7. Chen, T., Kornblith, S., Norouzi, M., and Hinton, G.: A simple framework for contrastive learning of visual representations. *ICML* (2020)
8. Dhariwal, P., and Nichol, A.: Diffusion models beat gans on image synthesis. *NeurIPS* 34, 8780–8794 (2021)
9. Freiesleben, T.: The Intriguing Relation Between Counterfactual Explanations and Adversarial Examples. *Minds & Machines* 32(1), 77–109 (2022)
10. Goodfellow, I.J., Shlens, J., and Szegedy, C.: Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014)
11. Goyal, Y., Wu, Z., Ernst, J., Batra, D., Parikh, D., and Lee, S.: Counterfactual visual explanations. *ICML* (2019)
12. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS* 30 (2017)
13. Ho, J., Jain, A., and Abbeel, P.: Denoising diffusion probabilistic models. *NeurIPS* 33, 6840–6851 (2020)
14. Hvilshøj, F., Iosifidis, A., and Assent, I.: On quantitative evaluations of counterfactuals. *arXiv* (2021)
15. Jeanneret, G., Simon, L., and Jurie, F.: Diffusion Models for Counterfactual Explanations. *ACCV 13847*, 219–237 (2022)
16. Krizhevsky, A., Sutskever, I., and Hinton, G.E.: ImageNet Classification with Deep Convolutional Neural Networks. *NeurIPS* (2012)
17. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. *ICCV* (2021)
18. Liu, Z., Mao, H., Wu, C.-Y., Feichtenhofer, C., Darrell, T., and Xie, S.: A convnet for the 2020s. *CVPR* (2022)
19. Luo, C.: Understanding Diffusion Models: A Unified Perspective. *CoRR abs/2208.11970* (2022)
20. Montavon, G., Binder, A., Lapuschkin, S., Samek, W., and Müller, K.-R.: Layer-wise relevance propagation: an overview. *Explainable AI: interpreting, explaining and visualizing deep learning* (2019)

21. Panch, T., Mattie, H., and Celi, L.A.: The “inconvenient truth” about AI in healthcare. *NPJ Digital Medicine* 2(1), 77 (2019)
22. Raaijmakers, S.: Artificial intelligence for law enforcement: challenges and opportunities. *IEEE Security & Privacy* 17(5), 74–77 (2019)
23. Ribeiro, M.T., Singh, S., and Guestrin, C.: Anchors: High-Precision Model-Agnostic Explanations. *AAAI* (2018)
24. Ribeiro, M.T., Singh, S., and Guestrin, C.: Why Should I Trust You?: Explaining the Predictions of Any Classifier. *KDD* (2016)
25. Rodriguez, P., Caccia, M., Lacoste, A., Zamparo, L., Laradji, I., Charlin, L., and Vazquez, D.: Beyond trivial counterfactual explanations with diverse valuable explanations. *ICCV* (2021)
26. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., and Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *IJCV* 115(3), 211–252 (2015)
27. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. *ICCV* (2017)
28. Simonyan, K., Vedaldi, A., and Zisserman, A.: Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. *ICLR* (2014)
29. Singla, S., Pollack, B., Chen, J., and Batmanghelich, K.: Explanation by Progressive Exaggeration. *ICLR* (2020)
30. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z.: Rethinking the Inception Architecture for Computer Vision. *CVPR* (2016)
31. Wachter, S., Mittelstadt, B., and Russell, C.: Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harv. JL & Tech.* 31, 841 (2017)
32. Yu, K.-H., Beam, A.L., and Kohane, I.S.: Artificial intelligence in healthcare. *Nature biomedical engineering* 2(10), 719–731 (2018)
33. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., and Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. *CVPR* (2018)