

# Where and How to Attack? A Causality-Inspired Recipe for Generating Counterfactual Adversarial Examples

Ruichu Cai<sup>1,2</sup>, Yuxuan Zhu<sup>1</sup>, Jie Qiao<sup>1\*</sup>, Zefeng Liang<sup>1</sup>, Furui Liu<sup>3</sup>, Zhifeng Hao<sup>4</sup>

<sup>1</sup>School of Computer Science, Guangdong University of Technology, Guangzhou, China

<sup>2</sup>Peng Cheng Laboratory, Shenzhen, China

<sup>3</sup>Zhejiang Lab, Hangzhou, China

<sup>4</sup>College of Science, Shantou University, Shantou, China

{cairuichu, iamyuxuanzhu, qiaojie.chn, lzfeng011021}@gmail.com, liufurui@zhejianglab.com, haozhifeng@stu.edu.cn

## Abstract

Deep neural networks (DNNs) have been demonstrated to be vulnerable to well-crafted *adversarial examples*, which are generated through either well-conceived  $\mathcal{L}_p$ -norm restricted or unrestricted attacks. Nevertheless, the majority of those approaches assume that adversaries can modify any features as they wish, and neglect the causal generating process of the data, which is unreasonable and unpractical. For instance, a modification in income would inevitably impact features like the debt-to-income ratio within a banking system. By considering the underappreciated causal generating process, first, we pinpoint the source of the vulnerability of DNNs via the lens of causality, then give theoretical results to answer *where to attack*. Second, considering the consequences of the attack interventions on the current state of the examples to generate more realistic adversarial examples, we propose CADE, a framework that can generate **C**ounterfactual **A**dversarial **E**xamples to answer *how to attack*. The empirical results demonstrate CADE’s effectiveness, as evidenced by its competitive performance across diverse attack scenarios, including white-box, transfer-based, and random intervention attacks.

## Introduction

Deep Neural Networks (DNNs) have achieved tremendous success in various tasks and have been widely used in critical domains such as facial recognition (Schroff, Kalenichenko, and Philbin 2015), medical diagnostics (Peng et al. 2021), and autonomous driving (Tian et al. 2018). Despite their unprecedented achievements, DNNs remain vulnerable to the well-crafted adversarial examples (Szegedy et al. 2014; Biggio et al. 2013), and there has been a recent thrust on generating adversarial examples through, e.g.,  $\mathcal{L}_p$ -norm restricted attack (Goodfellow, Shlens, and Szegedy 2015; Kurakin, Goodfellow, and Bengio 2017; Madry et al. 2018; Carlini and Wagner 2017; Moosavi-Dezfooli, Fawzi, and Frossard 2016), and unrestricted attack (Brown et al. 2017; Hosseini and Poovendran 2018; Bhattad et al. 2020; Song et al. 2018; Qiu et al. 2020; Yuan et al. 2022).

The  $\mathcal{L}_p$ -norm approaches reveal DNNs’ vulnerability by searching for the perturbation in raw pixel-space within a

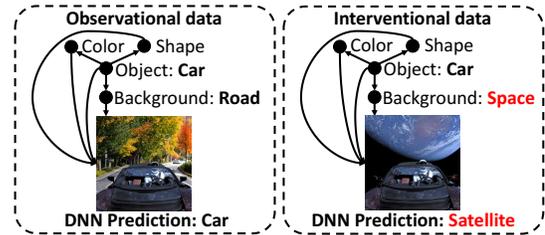


Figure 1: Discriminative DNN’s vulnerability to the interventional data.

bounded norm to preserve the photo-realism, while the unrestricted approaches replace such bounded perturbation with, e.g., geometric distortions (Guo et al. 2018), color/texture changing (Hosseini and Poovendran 2018; Bhattad et al. 2020; Yuan et al. 2022), and semantic changing (Qiu et al. 2020; Song et al. 2018), etc. Nevertheless, the majority of these methods assume that an attacker can modify any features as they wish, which is unreasonable if we aim to generate an adversarial example in real-world, e.g., the intractability of accessing the digital input to an image recognition model renders those methods perturbing the raw pixel-space fail. Moreover, we argue that only altering the alterable features while leaving others unchanged might also be impractical as it ignores the effect caused by the altering features, which has been underappreciated by the majority of the existing methods.

As a motivating example, consider a credit scoring model used by a financial institution to assess the creditworthiness of loan applicants. The model incorporates various features such as income, debt-to-income ratio, and credit history. To produce the adversarial example, it is unreasonable to disturb the income while leaving the debt-to-income ratio unchanged as it is induced by income and debt. This seemingly trivial observation has the underappreciated aspect that a causal generating process should also be involved to produce the adversarial example toward a practical scenario.

In this work, we provide a new perspective view on the adversarial attacks by taking the causal generating process into consideration, and propose a framework, CADE, that can generate **C**ounterfactual **A**dversarial **E**xamples. We introduce our CADE by answering two fundamental ques-

\*Corresponding author.

tions: 1) where to attack: understanding the adversarial example from the causal perspective to select valid disturbed variables; 2) how to attack: leveraging the causal generating process to generate more realistic/reasonable adversarial examples, since naively changing the cause variable without changing the effect variables will result in unrealistic examples. First, to answer where to attack, incorporated with structural information of the data, we give theoretical characterizations of the vulnerability of discriminative DNNs, i.e., the non-robustness to the interventional data, to which human perception is robust thanks to the capability of causal inference. For example, as the data-generating process shown on the left of Figure 1, a car is always on the ground, humans can recognize the car even if it is in space (interventional), while the DNNs recognize it as a satellite since it leverages the background of “space” on decision-making. Addressing this vulnerability, we analyze the effects of interventions explicitly to offer clear guidance for both observable and latent attacks. Second, to answer how to attack, the key problem is to predict the consequences when variables are intervened given the current observation, and the examples obtained are also known as counterfactuals in the literature of causality. For instance, in Figure 1, given the observation (left), when we intervene on background, the consequence is the changed image while other characters of the current observation are preserved (color, shape), where the preserved part is referred to the exogenous representing the current state of the example. To generate counterfactuals, we resort to the generation framework proposed in (Pearl 2009), which requires the causal generating process incorporated in. Thanks to the recent success of causal discovery (Zheng et al. 2018; Yu et al. 2019), generative modeling (Kingma and Welling 2014; Goodfellow et al. 2014; Ho, Jain, and Abbeel 2020), and causal representation learning (Kocaoglu et al. 2018; Yang et al. 2021; Shen et al. 2022), it is plausible to recover the generating process and generate counterfactual examples from interventional distribution practically. By knowing where and how to attack, our CADE offers an executable recipe to generate counterfactual examples. Empirically, our CADE achieves competitive results on white-box and transfer-based black-box attacks, and non-trivial performance with random intervention where no substitute model is involved.

Overall, our contributions are summarized as follows:

- We give a theoretical characterization of the discriminative DNNs’ vulnerability via the lens of causality, which offers clear guidance to answer where to attack.
- To generate more realistic examples, we propose CADE, a framework that can generate Counterfactual ADversarial Examples by considering the consequences of the interventions.
- The experimental results prove the effectiveness of our proposed CADE, by achieving competitive results on white-box, transfer-based, and even random attacks.

## Background

To reason counterfactual, we adopt the structural causal model (SCM) framework (Pearl 2009) which defines a

causal model as a triplet  $M(\mathbf{x}, f, \mathbf{u})$  over variables  $\mathbf{x} = \{x_1, \dots, x_d\}$  as: (i) a collection of structural assignments  $\{x_i := f_i(\text{Pa}_i, u_i)\}_{i=1}^d$ , where  $f_i$  are deterministic functions computing variable  $i$  from its causal parents  $\text{Pa}_i \subseteq \{x_1, \dots, x_d\} \setminus x_i$ ; and (ii) a factorizing joint distribution over the unobserved noise variables  $\mathbf{u} = \{u_1, \dots, u_d\}$ . Together, (i) and (ii) define a causal generative process and imply an observational joint distribution over  $x_1, \dots, x_d$  which factorizes over the causal graph  $\mathcal{G}$  as:  $p(x_1, \dots, x_d) = \prod_{i=1}^d p(x_i | \text{Pa}_i)$ .

## Methodology

In this section, we first explain the existence of adversarial examples from a causal perspective, and give theoretical analyses to answer *where to attack*. Then, we propose CADE, a framework that can generate Counterfactual ADversarial Examples, to answer *how to attack*.

### Motivating Example

We start with a motivating example of a linear model to illustrate the existence of adversarial examples via the model coefficients. We consider a linear data-generating process shown in Eq. (1).

$$\begin{cases} x_1 = u_1 & u_1 \sim \mathcal{N}(0, \sigma_1^2) \\ y = ax_1 + u_y & u_y \sim \mathcal{N}(0, \sigma_y^2) \\ x_2 = bx_3 + cy + u_2 & u_2 \sim \mathcal{N}(0, \sigma_2^2) \\ x_3 = u_3 & u_3 \sim \mathcal{N}(0, \sigma_3^2) \end{cases}, \quad (1)$$

where  $u_1, u_y, u_2$ , and  $u_3$  are the exogenous with zero mean and finite variance. Considering regressing the target variable  $y$  with variables  $x_1, x_2$  and  $x_3$  in a linear fashion, that is,  $\hat{y} = \mathbf{w}^T \mathbf{x}$ , where  $\mathbf{x} = [x_1, x_2, x_3]^T$  and  $\mathbf{w} = [w_1, w_2, w_3]^T$ . The model parameter  $\mathbf{w}$  is obtained by Empirical Risk Minimization (ERM), which is:

$$\mathbf{w} = \left[ \frac{\sigma_2^2 a}{\sigma_2^2 + \sigma_y^2 c^2}, \frac{\sigma_y^2 c}{\sigma_2^2 + \sigma_y^2 c^2}, -\frac{\sigma_y^2 bc}{\sigma_2^2 + \sigma_y^2 c^2} \right]^T, \quad (2)$$

where the proof of Eq. (2) is shown in Appendix A.1<sup>1</sup>. Since  $u_2$  has finite variance, we have non-zero  $w_2$  and  $w_3$ . However, according to Eq. (1), the most robust model should be  $\mathbf{w}^* = [a, 0, 0]^T$ , i.e., conceptualizing  $y$  only by  $x_1$ . We can observe the vulnerability of the model with non-zero  $w_2$  and  $w_3$ , indicating that we can generate adversarial examples by changing the realization of  $x_2$  (child of  $y$ ), or  $x_3$  (co-parent of  $y$ ). Nevertheless, when interventions are conducted, it is necessary to consider the consequence of each intervention, which the majority of the existing methods ignore. Considering the causal generating process, we analyze the interventional effect of both child ( $x_2$ ) and co-parent ( $x_3$ ) as follows.

**Children Intervention** Given an original input example  $\mathbf{x}$ , when intervening on the child of  $y$  (namely,  $x_2$ ), i.e.,  $do(x_2 = x_2 + \eta_2)$ , the generated counterfactual adversarial example  $\mathbf{x}^{adv}$  will be  $\mathbf{x}^{adv} = [x_1, x_2 + \eta_2, x_3]^T$ , which leads to an adversarial output:

$$\mathbf{w}^T \mathbf{x}^{adv} = \mathbf{w}^T \mathbf{x} + w_2 \eta_2,$$

where  $\mathbf{w}$  is shown in Eq. (2). The intervention on the child variable ( $x_2$ ) causes the shifted adversarial output by  $w_2 \eta_2$ .

<sup>1</sup>Appendices are available in (Cai et al. 2023).

**Co-parents Intervention** Regarding the intervention on *co-parents* of  $y$  (namly,  $x_3$ ), that is,  $do(x_3 = x_3 + \eta_3)$ , the counterfactual adversarial will be  $\mathbf{x}^{adv} = [x_1, x_2 + b\eta_3, x_3 + \eta_3]^T$ , rendering the adversarial output to be:

$$\mathbf{w}^T \mathbf{x}^{adv} = \mathbf{w}^T \mathbf{x} + w_2 b\eta_3 + w_3 \eta_3 = \mathbf{w}^T \mathbf{x}.$$

The intervention on the co-parent variables ( $x_3$ ) does not cause any damage to the output, where the shifted term  $w_3 \eta_3$  is canceled out by the interventional effect of  $x_3$  to its child variable  $x_2$ , which is  $w_2 b\eta_3$ . Intuitively, intervening co-parent variables ( $x_3$ ) does not change any structural mechanism as in Eq. (1), which the linear model has already modeled, thus causing no shift of the output.

### Generating Adversarial Example: Where to Attack?

The linear motivating example suggests an intuition of the vulnerability of the discriminative model via the coefficients. Despite the linear model, the theoretical analysis of this vulnerability of non-linear model (e.g., DNNs) is further investigated from a probabilistic view in this section. Specifically, we first give our analysis on observable variables, then extend the results to the latent variables in which the causal variables are not observable (e.g. the objects in the image).

**Observable Variable Intervention** From a probabilistic view, a discriminative DNN aims to approximate the conditional distribution  $p_M(y|\mathbf{x})$  via ERM, where  $M$  denotes the SCM that parameterizes the generating process of  $\mathbf{x}$  and  $y$ . With the great capacity, DNN can well approximate a distribution  $p_\theta(y|\mathbf{x}) \approx p_M(y|\mathbf{x})$  where  $\theta$  denotes the parameter of DNN, but exhibits limitation in generalizing to shifted distribution (Arjovsky et al. 2019; Tan et al. 2023). Here we analyze how  $y$  is predicted given the observable  $\mathbf{x}$  and how to obtain a shifted distribution by intervention to fool the model. First, given the observable  $\mathbf{x}$  to predict the target  $y$ , the conditional distribution can be derived as the following Proposition 1.

**Proposition 1.** *Given the SCM  $M$ , the discriminative conditional distribution:*

$$p_M(y|\mathbf{x}) = p_M(y|\text{Mb}_y^M).$$

where  $\text{Mb}_y^M$  denotes the Markov blanket of  $y$  under SCM  $M$ , including the parents, children, co-parents of  $y$ .

The proof of Proposition 1 is given in Appendix A.2. Proposition 1 suggests that the  $\text{Mb}_y^M$  are only variables needed to predict  $y$  given the observable  $\mathbf{x}$ . However, in the literature of causal inference, the most robust way of conceptualizing the target variable  $y$  is to only use its parent variables. Thus, the dependencies between  $y$  and  $\text{Mb}_y^M$  reveal the vulnerability of discriminative DNNs, offering adversaries an opportunity to attack by leveraging this property. When generating adversarial examples, intervening on  $y$  and its parents do not correspond to the attack since it changes the  $y$ , thus the remaining children and co-parents of  $y$  are only valid variables we can control. To answer *where to attack*, we theoretically analyze the effect of the intervention on children and co-parents separately as follows.

**Proposition 2.** *Given a SCM  $M(\{\mathbf{x}, y\}, f, \mathbf{u})$  where the underlying conditional distribution is  $p_M(y|\mathbf{x})$ , if an intervention is conducted on a  $y$ 's child  $x_j$  resulting in a new SCM  $M'(\{\mathbf{x}, y\}, f', \mathbf{u})$  with  $f' = \{f_i | i \neq j\} \cup \{f'_j\}$ , then we get nonequivalent interventional distribution  $p_{M'}(y|\mathbf{x}) \neq p_M(y|\mathbf{x})$ ; and if an intervention is conducted on a  $y$ 's co-parent but not child variable  $x_k$  resulting in a new SCM  $M'(\{\mathbf{x}, y\}, f', \mathbf{u})$  with  $f' = \{f_i : |i \neq k\} \cup \{f'_k\}$ , then the interventional distribution  $p_{M'}(y|\mathbf{x}) = p_M(y|\mathbf{x})$ .*

The proof of Proposition 2 is given in Appendix A.3. Proposition 2 states the nonequivalence in distributions between  $p_{M'}(y|\mathbf{x})$  and  $p_M(y|\mathbf{x})$  when *children* are intervened, and the equivalence between  $p_{M'}(y|\mathbf{x})$  and  $p_M(y|\mathbf{x})$  when *co-parents but not children* are intervened. Intuitively, the underlying rationale for the inequality lies in the disruption caused to the internal structure of  $\text{Mb}_y^M$ . Since the distribution  $p_\theta(y|\mathbf{x})$  learned by DNN is to approximate  $p_M(y|\mathbf{x})$  under the SCM  $M$ , it exhibits limitations in generalizing to the shifted interventional distribution  $p_{M'}(y|\mathbf{x})$ . This incapability of generalization offers a clear answer to where to attack, i.e., crafting adversarial examples drawn from a shifted interventional distribution with  $y$  preserved. This can be done by, such as children intervention, or both children and co-parents interventions that can damage the inner mechanisms within  $\text{Mb}_y^M$ , suggested by Proposition 2.

**Latent Variable Intervention** When facing images, most of the existing methods modify them in the raw pixel space, which is, however, impractical and highly-cost in the real-world. To mitigate this issue and keep the realism of the generated examples, it is plausible to attack the latent variables  $\mathbf{z}$  with semantics that determines  $\mathbf{x}$ , where each variable of  $\mathbf{z}$  can be causal-related, and the target  $y$  is included in  $\mathbf{z}$  that determines  $\mathbf{x}$  (Zhang et al. 2013). Since the image  $\mathbf{x}$  is a child of  $\mathbf{z}$ , each variable of  $\mathbf{z}$  except  $y$  becomes co-parent of  $y$ , indicating that when an image  $\mathbf{x}$  is given to predict  $y$ ,  $y$  is correlated with every other variable of  $\mathbf{z}$ . To answer *where to attack*, we investigate the connection between  $\mathbf{z}$  and the conditional distribution  $p_M(y|\mathbf{x})$  through the following proposition.

**Proposition 3.** *Given a SCM  $M(\{\mathbf{x}, \mathbf{z}\}, \{f, g\}, \{\mathbf{u}_x, \mathbf{u}_z\})$  where  $z_i = f_i(\text{Pa}_i^M, \mathbf{u}_i)$ ,  $\mathbf{x} = g(\mathbf{z}, \mathbf{u}_x)$ , and target  $y$  is included in  $\mathbf{z}$ , when interventions are conducted on the latent  $\mathbf{z}$  to obtain a new SCM  $M'(\{\mathbf{x}, \mathbf{z}\}, \{f', g\}, \{\mathbf{u}_x, \mathbf{u}_z\})$ , if the conditional distribution  $p_{M'}(y|\mathbf{x}) \neq p_M(y|\mathbf{x})$ , then  $p_{f'}(\mathbf{z}) \neq p_f(\mathbf{z})$ .*

Though  $p_{f'}(\mathbf{z}) \neq p_f(\mathbf{z})$  is a necessary condition for  $p_{M'}(y|\mathbf{x}) \neq p_M(y|\mathbf{x})$  according to Proposition 3, we analyze that at most cases,  $p_{f'}(\mathbf{z}) \neq p_f(\mathbf{z})$  can yield  $p_{M'}(y|\mathbf{x}) \neq p_M(y|\mathbf{x})$ , and the detailed proof and analysis of Proposition 3 is given in Appendix A.4. Since the DNN only fits the  $p_M(y|\mathbf{x})$ , we can generate the adversarial example drawn from a shifted  $p_{f'}(\mathbf{z})$ . Since the joint can be factorized as  $p_f(\mathbf{z}) = \prod_{i=1}^n p(z_i | \text{Pa}_i)$ , a new mechanism  $f'$  obtained by interventions that cause the structural change of  $\mathbf{z}$  can result in a shifted  $p_{f'}(\mathbf{z})$ . To preserve  $y$ , interventions on  $y$  and its parents do not correspond to the attack we consider, and some possible choices can be variables that

changed the structure of  $\mathbf{z}$  except for those two, such as the children of  $y$ .

### Generating Adversarial Example: How to Attack?

Knowing where to attack, the next step is to generate the adversarial example by considering the consequence of each intervention on the current state, since intervening on one variable will inevitably cause its descendants to change. The generated example, also called counterfactual, is the consequences under a hypothetical scenario where interventions are conducted, given the original example. To generate it, we resort to the framework proposed in (Pearl 2009; Pearl and Mackenzie 2018), which requires three fundamental steps: 1) abduction, 2) action, and 3) prediction. Here, we introduce the process for observable  $\mathbf{x}$ , and adapt it for latent  $\mathbf{z}$  in a similar fashion. First, to properly parameterize the causal generating process of an SCM  $M$  and compute the interventional effect efficiently, we can adopt the general non-linear generating process proposed in (Yu et al. 2019), which is:

$$f(\mathbf{x}) = \mathbf{A}^T f(\mathbf{x}) + \mathbf{u}, \quad (3)$$

where  $f$  denotes an invertible non-linear function, and  $\mathbf{A}$  denotes the weighted adjacency matrix of the causal DAG.

**Abduction** The abduction step aims to maintain the characters of the current state of the given observation, by recovering the exogenous  $\mathbf{u}$ . From Eq. (3), the exogenous  $\mathbf{u}$  can be recovered by:

$$\mathbf{u} = (\mathbf{I} - \mathbf{A}^T) f(\mathbf{x}). \quad (4)$$

**Action (Intervention) and Prediction** Formally, these two processes are as follows. First, intervene on the desired variables  $\mathbf{x}_S$  to obtain  $\mathbf{x}'$ , where  $\mathbf{x}_S$  is obtained from the variable selection process depicted in Figure 2 and  $\mathbf{x}'$  denotes the full input vector with  $\mathbf{x}_S$  changed. Then, predict the consequence of each intervention to obtain the corresponding adversarial example  $\mathbf{x}^{adv}$ . One way to incorporate these two processes is through the inversion version of Eq. (3), i.e.,  $\mathbf{x}^{adv} = f^{-1}((\mathbf{I} - \mathbf{A}^T)^{-1} \mathbf{u}')$ , where  $\mathbf{x}^{adv}$  is obtained by  $\mathbf{u}'$  with  $\mathbf{u}_S$  intervened. However, the matrix inversion process of  $(\mathbf{I} - \mathbf{A}^T)^{-1}$  will introduce extra error and computational cost. To remedy this issue, we directly intervene on  $\mathbf{x}_S$  to obtain  $\mathbf{x}^{adv}$  with the help of the mask  $\mathbf{m}$ :

$$\mathbf{x}^{adv} = f^{-1}(\mathbf{A}^T f(\mathbf{x}') \odot (1 - \mathbf{m}) + f(\mathbf{x}') \odot \mathbf{m} + \mathbf{u} \odot (1 - \mathbf{m})), \quad (5)$$

where  $\mathbf{m}$  is a binary mask with  $\mathbf{m}_S = 1$  indicating variable  $\mathbf{x}_S$  is intervened. Intuitively, the first term in the first bracket of r.h.s. of Eq. (5) aims to update the effect of each intervention and set the realization of the intervened variables to 0. The second term aims to add the realization of the intervened variables back. The third term aims to set the exogenous of the intervened variables to 0.

**White-Box Attack** Under a white-box setting, we can leverage the gradient information of the target model to guide the interventions. Attacking a discriminative DNN, from the probabilistic perspective, our goal is to generate an adversarial example  $\mathbf{x}^{adv}$  that can shift the conditional

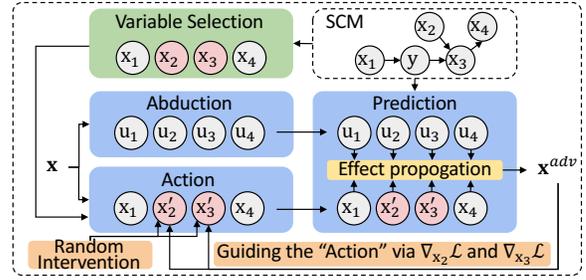


Figure 2: Framework of CADE.

probability  $p(y|\mathbf{x}^{adv})$ , which can be accomplished by maximizing the prediction loss:

$$\max_{\mathbf{x}^{adv} \in \mathcal{X}'} \mathcal{L}_{pred}(f_{\theta}(\mathbf{x}^{adv}), y), \quad (6)$$

where  $\mathcal{X}'$  denotes the sample space under a specific interventional distribution, and  $f_{\theta}(\cdot)$  denotes the output of DNN. To search on  $\mathcal{X}'$  to obtain  $\mathbf{x}^{adv}$ , we only update the selected variables  $\mathbf{x}_S$  by freezing the gradients of others except for  $\mathbf{x}_S$  in the *action* step, then update the consequence of each intervention according to (5) in the *prediction* step, as depicted in Figure 2. The algorithm is given in Appendix B.3.

**Black-Box Attack** The obtained white-box adversarial examples then can be transferred to other victim models thanks to the transferability property (Szegedy et al. 2014). Further, our CADE can be applied to a more strict scenario where the substitute model is absent. Benefit from the causal information, it is plausible to generate query-free adversarial examples without any white-box gradient information, by addressing the variables that can shift the conditional probability  $p(y|\mathbf{x})$ , according to Propositions 2, and 3. To achieve this, one way is to add random noises to those effective variables or reassign random realizations to them.

**Latent Attack for Image** For image data, we assume the images are generated according to a two-level "latent variable-image" model and we attack on the latent variables. When attacking, we first recover the causally related latent variables  $\mathbf{z} = \text{Encoder}(\mathbf{x})$  from the images using an encoder step. Then, following (4) and (5), we obtain the adversarial latent code  $\mathbf{z}^{adv}$  by replacing the observable  $\mathbf{x}$  with the latent  $\mathbf{z}$ . We finally obtain the adversarial example through  $\mathbf{x}^{adv} = \text{Decoder}(\mathbf{z}^{adv})$ . The detailed implementation and algorithm are given in Appendix B.2 and B.3.

## Experiment

### Experimental Setup

**Dataset** We evaluate our approach on three datasets, *Pendulum* (Yang et al. 2021), *CelebA* (Liu et al. 2015), and a synthetic measurement dataset, denoted as *SynMeasurement*. *Pendulum*, a synthetic image dataset generated by four causally related continuous variables, which follow the physical mechanism depicted in Figure 3 (a). *CelebA*, a real-world human face dataset with 40 labeled binary attributes has been well investigated in generative modeling (Choi

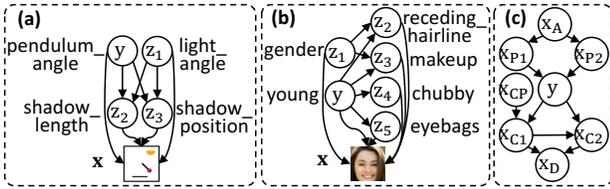


Figure 3: The causal graphs for each dataset: (a) *Pendulum*, (b) *CelebA(Attractive)*, and (c) *SynMeasurement*.

et al. 2018, 2020), including causal generative modeling (Kocaoglu et al. 2018; Yang et al. 2021; Shen et al. 2022). In our experiments, we deploy our CADE using a causal graph proposed in (Shen et al. 2022), which is depicted in Figure 3 (b). To further validate our theoretical characterization of CADE, we introduce a synthetic measurement dataset, *SynMeasurement*, whose corresponding causal graph is depicted in Figure 3 (c). More details about the generating process of those are given in Appendix B.1.

**Setup for Pendulum and CelebA** We select Res-50 (He et al. 2016), VGG-16 (Simonyan and Zisserman 2015), and their respective defense variants (adversarial trained with PGD (Madry et al. 2018), shortly denoted as Res-50(D) and VGG-16(D)) as our classification model. For *Pendulum*, our task is to predict the *pendulum\_angle* where the discrete categorical target labels are converted from the continuous *pendulum\_angle* based on to which angle intervals an image belongs. To make the dataset more realistic, we introduce random noises on *pendulum\_angle* on 15 % of data when generating images, representing the measurement error. For *CelebA*, our task is to predict if a person is young or not. To recover the causal latent representations depicted in Figure 3 (a), (b), we leverage a state-of-the-art generative model (Shen et al. 2022), and interventions on such representation are performed to generate counterfactual adversarial examples.

We evaluate the effectiveness of the attacks by reporting the attack success rates (ASR) of white-box, transfer-based, and random black-box attacks. For comparison, we select various state-of-the-art attack methods including PGD (Madry et al. 2018), C&W (Carlini and Wagner 2017), SAE (Hosseini and Poovendran 2018), ACE (Zhao, Liu, and Larson 2020), APGD( $\mathcal{L}_{inf}$ ) (Croce and Hein 2020), and APGD( $\mathcal{L}_1$ ) (Croce and Hein 2021) as our baselines. For *Pendulum*, we compare the results of our CADE intervening on different variables, i.e., CADE(1) on *light\_angle* (co-parent of target), CADE(2) on *shadow\_length* (child of target), CADE(3) on *shadow\_position* (child of target), and CADE(Mb) (Markov blanket of target). For *CelebA*, we implement our CADE by intervening on  $z_{1:5}$  depicted in Figure 3 (b). Details of such implementation are shown in Appendix B.

**Setup for SynMeasurement** We select Linear, MLP, and their respective adversarial trained (PGD) defense variants, denoted as Linear(D) and MLP(D) as our regression model to predict  $y$ . We test our CADE in a finer-grained setting where the CP (co-parent), C1 (child 1), and C1+C2 (child

1+2) interventions are conducted separately. Further, we investigate our CADE w/ and w/o the prediction process in counterfactual generation. Specifically, we use the term “intervention” (denoted as (i) shortly) to refer to the attacks w/ the prediction process, and “perturbation” (denoted as (p)) to refer to the attacks w/o the prediction process.

## Attacks on Pendulum

**Quantitative Analysis** We compare our proposed CADE with various state-of-the-art baselines, where the results of attack success rate are reported in Table 1.

Regarding the results compared with the baselines, we can observe that the adversarial examples obtained by CADE(Mb) achieve the highest transfer-based ASR among all competitors. Specifically, our CADE(Mb), CADE(2), CADE(3) not only achieve high ASR on the standard trained model but also achieve consistently high scores on the adversarial trained defense model, which indicates the effectiveness of our resulting adversarial examples to reveal the vulnerability of both standard and adversarial trained models. Besides, our CADE(Mb) also achieves competitive ASR on the white-box scenario, which is 99.4 and 99.7 under Res-50 and VGG-16, respectively. Further, our CADE with random intervention and no substitute model achieves competitive results compared to baselines, and even becomes the first or second winner on the transfer-based results.

Regarding the results of our CADE intervening on different variables, i.e., CADE(1), CADE(2), CADE(3), and CADE(Mb), first, we observe that the ASR of CADE(1) which intervenes on co-parent of  $y$  is lower than those of CADE(2) and CADE(3) which intervenes on the child of  $y$ . The lowest result of CADE(1) can be suggested by Proposition 3 with no structural change in  $z$ . However, the vulnerability of models to the CADE(1) can be due to, the white-box gradient exploiting the weakness of the model, and the models that overfit to the perfectly generated data tend to make mistakes when the imperfect legitimate examples are generated from the generative model with noises, which is also suggested in (Song et al. 2018). Further, we provide the result of the attack using *Pendulum* simulator in Appendix C.2.

**Qualitative Analysis** The adversarial examples obtained by different approaches are shown in Figure 4. Regarding the examples obtained by our CADE, we make the following observations. First, the example obtained by CADE(2) with shadow length shorter, the example obtained by CADE(3) with shadow position shifted, and the example obtained by CADE(Mb) with light angle, shadow length, shadow position changed can successfully fool the DNNs. Second, the example obtained by CADE(1) with light angle intervened, has caused the shadow changed based on the generative mechanism, rendering it fail to fool the DNNs.

## Attacks on CelebA

**Quantitative Analysis** We compare our CADE with several baselines where the ASR are reported in Table 2. First, we observe that our CADE guided by the causal graph in Figure 3 (b) outperforms the baselines under the transfer-based black-box setting on both standard and defense mod-

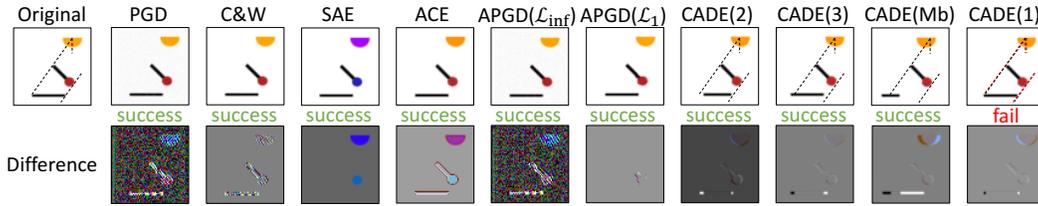


Figure 4: Visualization of adversarial examples against Res-50 on Pendulum obtained by different approaches. The black dash-line highlights the original projection trajectory, while the red dash-line highlights the intervened projection trajectory.

	Attacks	R50	R50(D)	V16	V16(D)
Res-50	PGD	<b>100.0*</b>	0.3	0.8	0.2
	C&W	<b>100.0*</b>	0.3	0.0	0.0
	SAE	<b>100.0*</b>	37.2	83.5	46.0
	ACE	<b>100.0*</b>	0.3	22.1	1.0
	APGD( $\mathcal{L}_{inf}$ )	<b>100.0*</b>	0.3	0.3	0.2
	APGD( $\mathcal{L}_1$ )	<b>100.0*</b>	57.2	55.8	47.3
	CADE(1)	75.3*	72.8	76.8	75.6
	CADE(2)	98.6*	96.3	<b>98.9</b>	98.9
	CADE(3)	95.0*	93.9	94.7	92.6
	CADE(Mb)	99.4*	<b>98.6</b>	<b>98.9</b>	<b>99.2</b>
VGG-16	PGD	5.8	0.5	<b>100.0*</b>	0.0
	C&W	2.9	0.3	99.8*	0.0
	SAE	24.4	1.8	<b>100.0*</b>	10.4
	ACE	26.4	2.1	<b>100.0*</b>	1.3
	APGD( $\mathcal{L}_{inf}$ )	5.4	0.5	<b>100.0*</b>	0.0
	APGD( $\mathcal{L}_1$ )	78.7	27.9	<b>100.0*</b>	63.9
	CADE(1)	70.7	74.2	75.5*	72.6
	CADE(2)	97.4	97.4	97.9*	98.2
	CADE(3)	93.1	95.2	95.2*	92.0
	CADE(Mb)	<b>98.9</b>	<b>99.7</b>	99.7*	<b>99.4</b>
×	CADE(1)	57.7	59.0	58.2	58.9
	CADE(2)	84.2	83.2	81.9	82.4
	CADE(3)	72.2	74.6	72.1	70.2
	CADE(Mb)	87.6	87.3	86.8	86.6

Table 1: ASR (%) on Pendulum. \* denotes white-box results, × denotes results without substitute models. R50 and V16 refer to Res-50 and VGG-16, respectively.

els. Second, we observe our CADE does not achieve the best on white-box, this can be due to, the causal graph provided in Figure 3 (b) is incomplete, and examples drawn from distribution obtained from intervening on such incomplete SCM have limited capacity to flip the prediction. Further, the random interventions with no substitute achieve the highest results on the defense model compared with baselines.

**Qualitative Analysis** Figure 5 shows the generated adversarial examples against Res-50 of different methods. We observe that compared with baselines, the examples generated by our CADE have reasonably-looking appearances with latent semantics intervened instead of naively adding noises in the pixel-space. Also, we can get some intuitions of why the attack success from the visualizations, e.g., on the left of Figure 5, examples generated by CADE look more feminist (gender) but still bald (receding hairline), suggesting that it could be drawn from an interventional distribution (that females can be bald) to which the DNNs cannot generalize.

	Attacks	R50	R50(D)	V16	V16(D)	
Res-50	PGD	<b>100.0*</b>	0.5	51.2	0.2	
	C&W	<b>100.0*</b>	0.5	0.8	0.2	
	SAE	77.0*	24.1	16.9	19.5	
	ACE	99.8*	21.7	15.3	18.3	
	APGD( $\mathcal{L}_{inf}$ )	<b>100.0*</b>	0.6	49.7	0.2	
	APGD( $\mathcal{L}_1$ )	92.5*	0.5	14.4	0.2	
	CADE	75.8*	<b>47.0</b>	<b>52.7</b>	<b>47.2</b>	
	VGG-16	PGD	41.1	0.6	97.3*	0.2
		C&W	0.9	0.5	<b>100.0*</b>	0.2
		SAE	23.4	23.7	90.2*	23.6
ACE		18.8	16.7	<b>100.0*</b>	14.4	
APGD( $\mathcal{L}_{inf}$ )		33.9	0.5	<b>100.0*</b>	0.2	
APGD( $\mathcal{L}_1$ )		3.0	0.5	99.8*	0.2	
×	CADE	<b>47.0</b>	<b>41.6</b>	80.6*	<b>44.7</b>	
×	CADE	25.9	25.8	26.3	25.0	

Table 2: ASR (%) on CelebA. \* denotes white-box results, × denotes results without substitute models. R50 and V16 refer to Res-50 and VGG-16, respectively.

Further, we showcase additional case study in Appendix C.

### Attacks on SynMeasurement

We evaluate our CADE in a finer-grained setting where the CP (co-parent), C1 (child 1), and C1+C2 (child 1+2) interventions are conducted separately, and further compare the results of both intervention (w/ prediction process) and perturbation (w/o prediction process), where the perturbation can be seen as a degradation to the majority of the existing methods, i.e, naively adding noises, and the RMSE results are reported in Figure 6.

**Child and Co-parent Interventions** Regarding the intervention, CP(i), C1(i), C1+C2(i), from Figure 6 we can make the following observations. First, the intervention on co-parent, CP(i), makes no effect on attacking the threat models under both black-box and white-box scenarios, which can be well suggested by Proposition 2. Second, both child interventions, C1(i) and C1+C2(i), achieve remarkable results with high RMSE scores. More specifically, intervention on one child C1(i) achieves a competitive result with the intervention on two children, C1+C2(i) on both standardly trained and adversarially trained models, which suggests a way of efficient attacks with lesser costs.

**Intervention and Perturbation** Comparing the intervention and perturbation results shown in Figure 6, we observe an interesting result where the interventions, C1(i) and

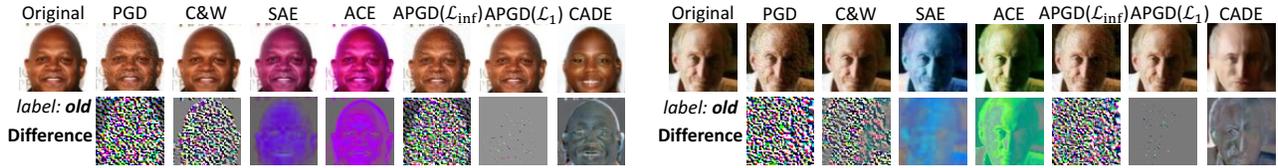


Figure 5: Visualization of adversarial examples against Res-50 on CelebA obtained by different approaches.

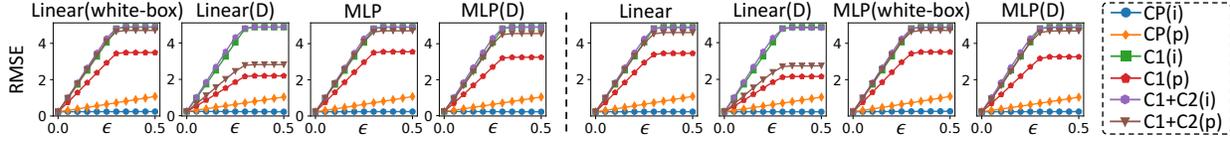


Figure 6: RMSE w.r.t. budget  $\epsilon$  of various interventions/perturbations.

C1+C2(i), achieve higher RMSE than their respective perturbations, C1(p), and C1+C2(p), suggesting that it is possible to achieve better attacks with lesser costs. This can be due to, the unchanged consequences of the perturbation, which contribute negative effects to the attacks.

### Related Work

**Adversarial Robustness** Here we briefly summarize the existing adversarial attack and defense approaches. For adversarial attack, one line of research lies on adding imperceptible noises within a bounded norm, for example, (Goodfellow, Shlens, and Szegedy 2015) searches for the perturbation direction by one-step gradient, (Moosavi-Dezfooli, Fawzi, and Frossard 2016) produces closer adversarial example by  $\mathcal{L}_2$  minimum perturbation, and (Madry et al. 2018) aims to find “most-adversarial” example by search for the local maximum loss value. Another line of works goes beyond the bounded perturbation, replacing it with, such as, (Sharif et al. 2016) fools a face-recognition model by wearing well-crafted eyeglasses frames, (Zhao, Dua, and Singh 2018) search for examples in the vicinity of latent space, and similarly (Song et al. 2018) generate class-conditioned examples in latent space. As for adversarial defense, one line of works focuses on defense at the training stage, including data augmentation, and adversarial training (Madry et al. 2018; Tramèr et al. 2018; Xie et al. 2019; Ilyas et al. 2019). Also, another line of works relies on input pre-processing, including JPEG compression, cropping, rescaling, etc (Guo et al. 2018; Dziugaite, Ghahramani, and Roy 2016; Xu, Evans, and Qi 2018). Further, a new line of research focuses on causality-inspired defense, e.g., (Zhang, Zhang, and Li 2020) improve the robustness of DNNs by test-time fine-tuning on unseen perturbation, which is explicitly modeled from a causal view, and (Zhang et al. 2022) improve the robustness by penalizing the DNNs incorporated with the restricted attack, where such perturbation is modeled from a causal perspective. However, they do not propose a proper way to generate adversarial examples, which is one of the main differences between ours and those methods.

**Causal Learning** Causal inference has a long history in statistics (Pearl 2009), and there is now increasing interest

in solving crucial problems of machine learning that benefit from causality (Schölkopf et al. 2021). For instance, causal discovery methods (Zheng et al. 2018; Yu et al. 2019; Cai et al. 2018; Qiao et al. 2023) identify the underlying causal structure from observations, causal representation learning (Kocaoglu et al. 2018; Yang et al. 2021; Shen et al. 2022) connect causal factors to high-dimensional observations, and several works show the benefits of causality in various applications, including domain adaptation (Zhang et al. 2013; Zhao et al. 2019; Cai et al. 2019), adversarial robustness, and reinforcement learning (Bareinboim, Forney, and Pearl 2015; Buesing et al. 2019), etc.

### Conclusion and Discussion

In this work, we propose CADE, a causality-inspired framework that can generate Counterfactual Adversarial Examples, by considering the underappreciated causal generating process. First, we reveal the vulnerability of discriminative DNNs to examples drawn from a nonequivalent interventional distribution with the equivalent target  $y$ , then we provide an explicit recipe to answer where to attack: intervening on variables that renders distribution shift while preserving the consistent target  $y$ . Second, to generate more realistic examples, we consider the consequence of each intervention on the current state, then propose CADE to generate counterfactual adversarial examples, answering how to attack. Our experiments demonstrate the effectiveness of CADE, by achieving competitive results on white-box and transfer-based attacks, even non-trivial results on random intervention where the substitute model is not required.

Our CADE is effective, however, obtaining the required causal generating process remains an open challenge. Our framework should be adapted to more realistic settings where the full generating process can not be obtained. In future work, we will explore such settings with limited causal knowledge, in which only partial causal knowledge can be obtained from causal discovery algorithms or expert knowledge. Further, we will verify our framework beyond the digital simulation, and adapt it to the physical world. Overall, we believe this work would open up new research opportunities and challenges in the field of adversarial learning, potentially inspiring new designs of defense mechanisms.

## Acknowledgments

This research was supported in part by National Key R&D Program of China (2021ZD0111501), National Science Fund for Excellent Young Scholars (62122022), Natural Science Foundation of China (61876043, 61976052), the major key project of PCL (PCL2021A12).

## References

- Arjovsky, M.; Bottou, L.; Gulrajani, I.; and Lopez-Paz, D. 2019. Invariant Risk Minimization. arXiv:1907.02893.
- Bareinboim, E.; Forney, A.; and Pearl, J. 2015. Bandits with Unobserved Confounders: A Causal Approach. In *Advances in Neural Information Processing Systems*, 1342–1350.
- Bhattad, A.; Chong, M. J.; Liang, K.; Li, B.; and Forsyth, D. A. 2020. Unrestricted Adversarial Examples via Semantic Manipulation. In *8th International Conference on Learning Representations, ICLR*. OpenReview.net.
- Biggio, B.; Corona, I.; Maiorca, D.; Nelson, B.; Srndic, N.; Laskov, P.; Giacinto, G.; and Roli, F. 2013. Evasion Attacks against Machine Learning at Test Time. In *ECML-PKDD*, volume 8190 of *Lecture Notes in Computer Science*, 387–402. Springer.
- Brown, T. B.; Mané, D.; Roy, A.; Abadi, M.; and Gilmer, J. 2017. Adversarial Patch. arXiv:1712.09665.
- Buesing, L.; Weber, T.; Zwols, Y.; Heess, N.; Racanière, S.; Guez, A.; and Lespiau, J. 2019. Woulda, Coulda, Shoulda: Counterfactually-Guided Policy Search. In *7th International Conference on Learning Representations, ICLR 2019*. OpenReview.net.
- Cai, R.; Li, Z.; Wei, P.; Qiao, J.; Zhang, K.; and Hao, Z. 2019. Learning Disentangled Semantic Representation for Domain Adaptation. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*.
- Cai, R.; Qiao, J.; Zhang, Z.; and Hao, Z. 2018. Self: structural equational likelihood framework for causal discovery. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Cai, R.; Zhu, Y.; Qiao, J.; Liang, Z.; Liu, F.; and Hao, Z. 2023. Where and How to Attack? A Causality-Inspired Recipe for Generating Counterfactual Adversarial Examples. arXiv:2312.13628.
- Carlini, N.; and Wagner, D. A. 2017. Towards Evaluating the Robustness of Neural Networks. In *2017 IEEE Symposium on Security and Privacy, SP*, 39–57. IEEE Computer Society.
- Choi, Y.; Choi, M.; Kim, M.; Ha, J.; Kim, S.; and Choo, J. 2018. StarGAN: Unified Generative Adversarial Networks for Multi-Domain Image-to-Image Translation. In *CVPR 2018*, 8789–8797. Computer Vision Foundation / IEEE Computer Society.
- Choi, Y.; Uh, Y.; Yoo, J.; and Ha, J. 2020. StarGAN v2: Diverse Image Synthesis for Multiple Domains. In *CVPR 2020*, 8185–8194. Computer Vision Foundation / IEEE.
- Croce, F.; and Hein, M. 2020. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020*, volume 119 of *Proceedings of Machine Learning Research*, 2206–2216. PMLR.
- Croce, F.; and Hein, M. 2021. Mind the Box:  $l_1$ -APGD for Sparse Adversarial Attacks on Image Classifiers. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021*, volume 139 of *Proceedings of Machine Learning Research*, 2201–2211. PMLR.
- Dziugaite, G. K.; Ghahramani, Z.; and Roy, D. M. 2016. A study of the effect of JPG compression on adversarial images. arXiv:1608.00853.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A. C.; and Bengio, Y. 2014. Generative Adversarial Networks. arXiv:1406.2661.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2015. Explaining and Harnessing Adversarial Examples. In *3rd International Conference on Learning Representations, ICLR*.
- Guo, C.; Rana, M.; Cissé, M.; and van der Maaten, L. 2018. Countering Adversarial Images using Input Transformations. In *6th International Conference on Learning Representations, ICLR 2018*. OpenReview.net.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *CVPR 2016*, 770–778. IEEE Computer Society.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems, NeurIPS 2020*.
- Hosseini, H.; and Poovendran, R. 2018. Semantic Adversarial Examples. In *2018 IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 1614–1619. Computer Vision Foundation / IEEE Computer Society.
- Ilyas, A.; Santurkar, S.; Tsipras, D.; Engstrom, L.; Tran, B.; and Madry, A. 2019. Adversarial Examples Are Not Bugs, They Are Features. In *Advances in Neural Information Processing Systems, NeurIPS 2019*, 125–136.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014*.
- Kocaoglu, M.; Snyder, C.; Dimakis, A. G.; and Vishwanath, S. 2018. CausalGAN: Learning Causal Implicit Generative Models with Adversarial Training. In *6th International Conference on Learning Representations, ICLR 2018*. OpenReview.net.
- Kurakin, A.; Goodfellow, I. J.; and Bengio, S. 2017. Adversarial examples in the physical world. In *5th International Conference on Learning Representations, ICLR 2017*. OpenReview.net.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *6th International Conference on Learning Representations, ICLR*. OpenReview.net.
- Moosavi-Dezfooli, S.; Fawzi, A.; and Frossard, P. 2016. DeepFool: A Simple and Accurate Method to Fool Deep Neural Networks. In *2016 IEEE Conference on Computer*

- Vision and Pattern Recognition, CVPR*, 2574–2582. IEEE Computer Society.
- Pearl, J. 2009. *Causality: Models, Reasoning and Inference*. USA: Cambridge University Press, 2nd edition. ISBN 052189560X.
- Pearl, J.; and Mackenzie, D. 2018. *The Book of Why: The New Science of Cause and Effect*. USA: Basic Books, Inc., 1st edition. ISBN 046509760X.
- Peng, J.; Wang, P.; Desrosiers, C.; and Pedersoli, M. 2021. Self-Paced Contrastive Learning for Semi-supervised Medical Image Segmentation with Meta-labels. In *Advances in Neural Information Processing Systems, NeurIPS 2021*, 16686–16699.
- Qiao, J.; Cai, R.; Wu, S.; Xiang, Y.; Zhang, K.; and Hao, Z. 2023. Structural Hawkes Processes for Learning Causal Structure from Discrete-Time Event Sequences. In *IJCAI-23*, 5702–5710. IJCAI Organization.
- Qiu, H.; Xiao, C.; Yang, L.; Yan, X.; Lee, H.; and Li, B. 2020. SemanticAdv: Generating Adversarial Examples via Attribute-Conditioned Image Editing. In *Computer Vision - ECCV 2020 - 16th European Conference*, volume 12359, 19–37. Springer.
- Schölkopf, B.; Locatello, F.; Bauer, S.; Ke, N. R.; Kalchbrenner, N.; Goyal, A.; and Bengio, Y. 2021. Towards Causal Representation Learning. arXiv:2102.11107.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. FaceNet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, 815–823. IEEE Computer Society.
- Sharif, M.; Bhagavatula, S.; Bauer, L.; and Reiter, M. K. 2016. Accessorize to a Crime: Real and Stealthy Attacks on State-of-the-Art Face Recognition. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 1528–1540. Association for Computing Machinery. ISBN 9781450341394.
- Shen, X.; Liu, F.; Dong, H.; Lian, Q.; Chen, Z.; and Zhang, T. 2022. Weakly Supervised Disentangled Generative Causal Representation Learning. *J. Mach. Learn. Res.*, 23: 241:1–241:55.
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *3rd International Conference on Learning Representations, ICLR 2015*.
- Song, Y.; Shu, R.; Kushman, N.; and Ermon, S. 2018. Constructing Unrestricted Adversarial Examples with Generative Models. In *Advances in Neural Information Processing Systems, NeurIPS 2018*, 8322–8333.
- Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I. J.; and Fergus, R. 2014. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations, ICLR*.
- Tan, X.; Yong, L.; Zhu, S.; Qu, C.; Qiu, X.; Xu, Y.; Cui, P.; and Qi, Y. 2023. Provably Invariant Learning without Domain Information. In *ICML 2023*, volume 202 of *Proceedings of Machine Learning Research*, 33563–33580. PMLR.
- Tian, Y.; Pei, K.; Jana, S.; and Ray, B. 2018. DeepTest: automated testing of deep-neural-network-driven autonomous cars. In *Proceedings of the 40th International Conference on Software Engineering, ICSE 2018*, 303–314. ACM.
- Tramèr, F.; Kurakin, A.; Papernot, N.; Goodfellow, I. J.; Boneh, D.; and McDaniel, P. D. 2018. Ensemble Adversarial Training: Attacks and Defenses. In *6th International Conference on Learning Representations, ICLR 2018*. OpenReview.net.
- Xie, C.; Wu, Y.; van der Maaten, L.; Yuille, A. L.; and He, K. 2019. Feature Denoising for Improving Adversarial Robustness. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019*, 501–509. Computer Vision Foundation / IEEE.
- Xu, W.; Evans, D.; and Qi, Y. 2018. Feature Squeezing: Detecting Adversarial Examples in Deep Neural Networks. In *NDSS 2018*. The Internet Society.
- Yang, M.; Liu, F.; Chen, Z.; Shen, X.; Hao, J.; and Wang, J. 2021. CausalVAE: Disentangled Representation Learning via Neural Structural Causal Models. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021*, 9593–9602. Computer Vision Foundation / IEEE.
- Yu, Y.; Chen, J.; Gao, T.; and Yu, M. 2019. DAG-GNN: DAG Structure Learning with Graph Neural Networks. In *ICML 2019*, volume 97, 7154–7163. PMLR.
- Yuan, S.; Zhang, Q.; Gao, L.; Cheng, Y.; and Song, J. 2022. Natural Color Fool: Towards Boosting Black-box Unrestricted Attacks. In *NeurIPS*.
- Zhang, C.; Zhang, K.; and Li, Y. 2020. A Causal View on Robustness of Neural Networks. In *NeurIPS 2020*.
- Zhang, K.; Schölkopf, B.; Muandet, K.; and Wang, Z. 2013. Domain Adaptation under Target and Conditional Shift. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013*, volume 28 of *JMLR Workshop and Conference Proceedings*, 819–827. JMLR.org.
- Zhang, Y.; Gong, M.; Liu, T.; Niu, G.; Tian, X.; Han, B.; Schölkopf, B.; and Zhang, K. 2022. Adversarial Robustness Through the Lens of Causality. In *The Tenth International Conference on Learning Representations, ICLR 2022*. OpenReview.net.
- Zhao, H.; des Combes, R. T.; Zhang, K.; and Gordon, G. J. 2019. On Learning Invariant Representations for Domain Adaptation. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, volume 97 of *Proceedings of Machine Learning Research*, 7523–7532. PMLR.
- Zhao, Z.; Dua, D.; and Singh, S. 2018. Generating Natural Adversarial Examples. In *6th International Conference on Learning Representations, ICLR 2018*. OpenReview.net.
- Zhao, Z.; Liu, Z.; and Larson, M. A. 2020. Adversarial Color Enhancement: Generating Unrestricted Adversarial Images by Optimizing a Color Filter. In *31st British Machine Vision Conference 2020, BMVC*. BMVA Press.
- Zheng, X.; Aragam, B.; Ravikumar, P.; and Xing, E. P. 2018. DAGs with NO TEARS: Continuous Optimization for Structure Learning. In *Advances in Neural Information Processing Systems, NeurIPS 2018*, 9492–9503.