Training on Plausible Counterfactuals Removes Spurious Correlations

Shpresim Sadiku^{1,2}, Kartikeya Chitranshi^{1,2}, Hiroshi Kera^{2,3}, and Sebastian Pokutta^{1,2}

¹Technische Universität Berlin, Institute of Mathematics ²Zuse Institute Berlin, Department AIS2T, *lastname*@zib.de ³Chiba University, Institute for Advanced Academic Research, kera@chiba-u.jp

Abstract

Plausible counterfactual explanations (p-CFEs) are perturbations that minimally modify inputs to change classifier decisions while remaining plausible under the data distribution. In this study, we demonstrate that classifiers can be trained on p-CFEs labeled with induced *incorrect* target classes to classify unperturbed inputs with the original labels. While previous studies have shown that such learning is possible with adversarial perturbations, we extend this paradigm to p-CFEs. Interestingly, our experiments reveal that learning from p-CFEs is even more effective: the resulting classifiers achieve not only high in-distribution accuracy but also exhibit significantly reduced bias with respect to spurious correlations.

1 Introduction

Altering a classifier's prediction through minimal input perturbations has yielded valuable insights into the decision-making processes of machine learning models. *Adversarial attacks* (Szegedy et al., 2014), for instance, have demonstrated the unexpected vulnerability of well-trained models to imperceptibly small perturbations, and various forms of such perturbations have been found through extensive studies (Wachter et al., 2017; Madry et al., 2018b). In contrast, *plausible counterfactual explanations* (p-CFEs) are minimal perturbations that alter classifications in a semantically coherent manner. Designed to align with the data manifold and to be interpretable, p-CFEs offer visual explanations of model predictions through "what-if" scenarios (Zhang et al., 2023; Sadiku et al., 2025a). Despite differing in their constraints, adversarial attacks and p-CFEs share a common objective: altering model predictions through minimal input changes-—suggesting potential synergies that remain underexplored in the literature.

Recently, Kumano et al. (2024a) revisited a seminal study by Ilyas et al. (2019) and theoretically justified their observations: adversarial perturbations, although seemingly subtle and meaningless, actually contain generalizable, class-specific features—a model trained on them labeled with their *induced* (incorrect) classes can successfully classify clean images into *original* classes. Given the structural similarity between adversarial examples and p-CFEs as minimal input perturbations, it is natural to ask whether the representational richness observed in adversarial examples also applies to p-CFEs. However, the defining characteristic of p-CFEs—*plausibility*—may lead to fundamentally different outcomes.

This study addresses *learning from p-CFEs* and empirically demonstrates that it is more effective than learning from adversarial perturbations, achieving higher in-distribution and out-of-distribution accuracy in the presence of *spurious correlations*-features that correlate with the label during training but are semantically irrelevant. Our results suggest that p-CFEs not only induce prediction flips but also guide models toward learning features that better reflect the true, semantic structure of the data.



Figure 1: Random samples from our WaterBirds training set variants. The bottom row shows perturbations (magnified 40 times for visibility) applied to the original image on the left by different methods. The true label is *water bird*; the target label is *land bird*.

Contributions

- 1. New instance of learning from perturbations. We demonstrate that the insightful observation of learning from adversarial perturbations by Ilyas et al. (2019) generalizes to p-CFEs, highlighting the broader applicability of the *learning from perturbations* paradigm.
- 2. High classification accuracy. Our experiments show that learning from p-CFEs (Sadiku et al., 2025a) yields high classification accuracy on original samples—comparable to models trained on adversarial perturbations (including ℓ_2 and ℓ_{∞} PGD, and ℓ_2 CFEs).
- 3. Removing spurious correlations. The experiments further reveal that learning from p-CFEs significantly outperforms learning from perturbations in mitigating spurious correlations. On the WaterBirds dataset, where evaluation sets are designed with strong spurious correlations, learning from p-CFEs even surpasses standard (noise-free) training by 12% in worst-group accuracy.

2 Related Work

Small input perturbations have yielded various insightful observations that deepen our understanding of machine learning models. A prominent example is adversarial perturbations (Szegedy et al., 2014), which can easily fool seemingly strong classifiers with imperceptibly small changes, thereby questioning the reliability of machine learning models. Such perturbations exist in various forms (Madry et al., 2018b; Xu et al., 2019; Kazemi et al., 2023), and they are known to transfer across different models (Xiaosen et al., 2023; Sadiku et al., 2025b). Counterfactual explanations (CFEs; (Wachter et al., 2017)) represent another line of research on small input perturbations. While early CFEs resembled adversarial perturbations, recent approaches increasingly emphasize alignment with the data manifold. In particular, Sadiku et al. (2025a) proposed a method for generating plausible counterfactuals (p-CFEs) via proximal gradient optimization, yielding perturbations that contain semantics aligned with the target class. However, such counterfactuals have primarily been used for interpretability purposes rather than as training data. Although a few studies have analyzed the connection between CFEs and adversarial perturbations (Pawelczyk et al., 2022; Freiesleben, 2022), the interaction between these two research directions remains limited.

This study investigates whether observations in the literature of adversarial perturbations, particularly learning from adversarial perturbations (Ilyas et al., 2019; Kumano et al., 2024a,b), extend to p-CFEs, and how the plausibility of p-CFEs gives rise to distinct outcomes.

3 Preliminaries

Assume a binary classification setting where $\mathcal{X} \subseteq \mathbb{R}^d$ denotes the input space, $\mathcal{Y} = \{\pm 1\}$ denotes the set of possible class labels, and $\mathcal{D} = \{(\boldsymbol{x}_i, y_i) \in \mathcal{X} \times \mathcal{Y}\}_{i=1}^n$ is a training dataset consisting of n independent and identically distributed data points generated from a joint density $\psi : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}_+$. Furthermore, we define $q(\boldsymbol{x}, y) := \psi(\boldsymbol{x}|y)$, which is the corresponding density of the inputs conditioned on the given label y.

We let $f_{\theta} : \mathcal{X} \to \mathbb{R}^2$ denote a neural network classifier that takes a *d*-dimensional sample as input and outputs logits of the two classes. The final decision is denoted by $f(\boldsymbol{x}) := \arg \max_i [f_{\theta}(\boldsymbol{x})]_i$. For $d \in \mathbb{N}$ let $[d] = \{1, ..., d\}$.

Standard Training. With the exponential loss $\mathcal{L}(\boldsymbol{x}, y) := \exp(-y \cdot f(\boldsymbol{x}))$ or logistic loss $\mathcal{L}(\boldsymbol{x}, y) := \ln(1 + \exp(-y \cdot f(\boldsymbol{x})))$, training a classifier is performed by minimizing a loss function $\hat{R}(\theta) := \sum_{i=1}^{n} \mathcal{L}(\boldsymbol{x}_i, y_i)/n$ from the training set via *empirical risk minimization* (ERM).

4 Learning from p-CFEs

We now introduce a formal definition of learning from perturbations.

Definition 4.1 (Learning from perturbations). Let $\mathcal{D} := \{(\boldsymbol{x}_i, y_i)\}_{i=1}^n$ be a training dataset, where each \boldsymbol{x}_i is an input and y_i is its corresponding label. Let f be a classifier trained on \mathcal{D} via standard training. For each i, a perturbed example $\tilde{\boldsymbol{x}}_i$ is generated to increase the probability of a target label $\tilde{y}_i \neq y_i$ under f, resulting in a perturbed dataset $\tilde{\mathcal{D}} := \{(\tilde{\boldsymbol{x}}_i, \tilde{y}_i)\}_{i=1}^n$. Training a classifier from scratch on $\tilde{\mathcal{D}}$, is referred to as learning from perturbations.

Prior studies (Ilyas et al., 2019; Kumano et al., 2024a,b) assume the perturbations to be adversarial ones, typically generated using targeted Projected Gradient Descent (PGD; (Madry et al., 2018a)) via the following optimization for input \boldsymbol{x} , target label \tilde{y} , and perturbation budget $\epsilon > 0$

$$\min_{\boldsymbol{x}\in\mathcal{X}} \mathcal{L}(\tilde{\boldsymbol{x}}, \tilde{y}) \quad \text{s.t.} \quad \|\tilde{\boldsymbol{x}} - \boldsymbol{x}\|_p \le \epsilon,$$
(1)

where $\|\cdot\|_p$ denotes the ℓ_p norm. Crucially, this differs from standard *adversarial training*: each training sample \tilde{x}_i is paired with a *target incorrect label* \tilde{y}_i , while evaluation is performed on the clean input x_i with the *original label* y.

We extend this idea to p-CFEs by perturbing each input-label pair (\boldsymbol{x}, y) using the method of Sadiku et al. (2025a), formulated as the following unconstrained optimization problem

$$\tilde{\boldsymbol{x}} := \underset{\boldsymbol{x}' \in \mathcal{A}}{\operatorname{arg\,min}} \|\boldsymbol{x}' - \boldsymbol{x}\|_{2}^{2} + \gamma \mathcal{L}(\boldsymbol{x}', \tilde{y}) - \tau \hat{q}(\boldsymbol{x}', \tilde{y}) + \beta \|\boldsymbol{x}' - \boldsymbol{x}\|_{0},$$
(2)

where $\hat{q}(\cdot, \tilde{y})$ estimates the density of the target class \tilde{y} in \mathcal{X} , and $\mathcal{A} := \bigotimes_{i=1}^{d} [-\mathcal{A}_i, \mathcal{A}_i]$, with $\mathcal{A}_i \in \mathbb{R}$, defines the feature value range, either derived from the dataset or specified by the user. The parameters $\gamma, \tau, \beta > 0$ control tradeoffs for *validity* (flipping the decision), *plausibility* (staying on the data manifold), and *sparsity* (minimizing feature changes), respectively.

Plausibility Term. The plausibility term $\hat{q}(\cdot, \tilde{y})$ only needs to be differentiable, enabling gradient-based optimization. For example, KDEs and GMMs are standard differentiable estimators (Sadiku et al., 2025a).



Figure 2: **Top row**: Original and Grad-CAM (Selvaraju et al., 2017) visualizations for a misclassified *landbird* (with a water background) from the WaterBirds dataset—incorrectly predicted as a *waterbird*. **Bottom row**: Original and Grad-CAM visualizations for a misclassified *big dog* (with an indoor background) from the SpuCoAnimals dataset—incorrectly predicted as a *small dog*.



Figure 3: Saliency maps for different models. Left to right: (1) original image of a land bird (top) and a dog (bottom), (2) saliency map from a standard model, (3-5) maps from models trained on PGD (ℓ_2 , ℓ_∞) and CFE ℓ_2 adversarial examples, (6) maps from models trained on p-CFE ℓ_0 examples. All models use ResNet50.

For higher-dimensional data, more expressive models like VAEs or GANs can be used (Van Looveren and Klaise, 2021). The method is compatible with any modern classifier that supports backpropagation, including large language models (e.g., GPT-n with n > 2) and vision models such as Swin Transformers with GELU activations (Liu et al., 2021).

Adversarial vs. Counterfactual. The key distinction in Eq. (2) compared to adversarial attacks is the term minimizing the negative target-class density estimate, which encourages perturbed instances to lie on the target data manifold. By rewriting Eq. (2) as the sum of a smooth (possibly non-convex) term and a non-smooth term with a closed-form proximal operator, we adopt the solution strategy of Sadiku et al. (2025a), using the accelerated proximal gradient method of Beck and Teboulle (2009) to generate p-CFEs.

5 Experiments

Datasets. We adopt two standard benchmark datasets that involve spurious correlations. The **Water-Birds** (Sagawa et al., 2020a) dataset has two labels, namely landbird and waterbird. The spurious correlation arises from the change in background (like a water bird on land or a landbird above or on water). The **SpuCoAnimals** (Joshi et al., 2023) dataset contains two categories: big dogs and small dogs. Spurious correlation arises on the assumption that big dogs were mostly outside and the small dogs inside the house. Fig. 2 visually illustrates spurious correlations in both datasets, with additional examples given in Appendix A. For a review on spurious correlations, see (Ye et al., 2024).

Setup. We fine-tune a pre-trained ResNet-50 model (He et al., 2016) on original images, PGD (ℓ_2, ℓ_∞) , CFE ℓ_2 -adversarial perturbations (Madry et al., 2018b; Wachter et al., 2017), and p-CFE (ℓ_0) perturbations (Sadiku et al., 2025a).¹ The target classes \tilde{y} are chosen uniformly at random, which makes the features of original images become uncorrelated with the labels. Training examples are shown in Fig. 1. We use the SGD optimizer in PyTorch with a learning rate of 1e-5, momentum of 0.9, weight decay of 5e-2, batch size of 8, and train for 360 epochs.

Metrics. Train and test accuracies are the fractions of correctly classified samples in the training and test sets, respectively. To quantify spurious correlations, we use worst-group accuracy (WGA) as defined in (Sagawa et al., 2020b). For instance, for WaterBirds, groups are defined as *(attribute, label)* pairs. The

 $^{^{1}}$ Kumano et al. (2024a) extended Theorem 4.1 to show that training solely on adversarial perturbations can match the accuracy of models trained on original data. Our experiments with this setup and standard vision datasets are detailed in Appendix B.

Table 1: Train and test accuracies across various methods for various datasets. The model used is ResNet50 and we follow the setup of Yang et al. (2023). The results of standard training (Original) are provided for reference. Learning from p-CFEs suffers less from spurious correlations than learning from adversarial perturbations.

Dataset	Split Set	PGD ℓ_2 (std)	PGD ℓ_{∞} (std)	CFE ℓ_2 (std)	p-CFE ℓ_0 (std)	Original (std)
WaterBirds	Train Test	$\begin{array}{c} 97.04 \ (0.07) \\ 86.08 \ (0.69) \end{array}$	98.00 (0.08) 86.02 (0.56)	97.93 (0.07) 88.58 (0.61)	$\begin{array}{c} 91.50 \ (0.56) \\ 86.54 \ (0.26) \end{array}$	$\begin{array}{ c c c c c c c c c c c c c c c c c c c$
SpuCoAnimals	Train Test	$\begin{array}{c} 96.25 \ (0.03) \\ 78.10 \ (0.92) \end{array}$	$\begin{array}{c} 97.47 \ (0.02) \\ 79.43 \ (0.68) \end{array}$	97.89 (0.08) 79.00 (0.93)	97.10 (0.01) 81.78 (0.59)	$\begin{array}{ c c c c c c c c c c c c c c c c c c c$

Table 2: Worst-group accuracies across various methods for different datasets, following the configuration of Yang et al. (2023). The results of standard training (Original) are provided for reference. Learning from p-CFEs suffers less from spurious correlations and even outperforms standard training on the WaterBirds dataset.

Dataset	Split Set	PGD ℓ_2 (std)	PGD ℓ_{∞} (std)	CFE ℓ_2 (std)	p-CFE ℓ_0 (std)	Original (std)
WaterBirds	Train Test	56.55 (4.20) 56.58 (2.17)	$\begin{array}{c} 72.00 \ (4.69) \\ 61.72 \ (2.50) \end{array}$	$\begin{array}{c} 74.99 \ (2.54) \\ 63.04 \ (2.19) \end{array}$	77.97 (2.22) 76.05 (1.45)	$\begin{array}{ c c c c c c c c c c c c c c c c c c c$
SpuCoAnimals	Train Test	$\begin{array}{c} 62.60 \ (1.33) \\ 56.06 \ (1.99) \end{array}$	$\begin{array}{c} 74.60 \ (1.41) \\ 57.53 \ (1.79) \end{array}$	$\begin{array}{c} 72.86 \ (1.61) \\ 56.60 \ (3.15) \end{array}$	80.20 (0.86) 63.53 (1.55)	$\begin{array}{ c c c c c } 99.70 & (0.14) \\ 65.60 & (0.90) \end{array}$

worst-group—(waterbirds, land)—has only 56 training samples, while other groups (e.g., (landbirds, land), (waterbirds, water), (landbird, water) have up to 20 times more. A WGA of 0 indicates that all (waterbirds, land) examples were misclassified, revealing a strong spurious correlation with the land background. We compute these metrics across five classifiers trained on original data, PGD (ℓ_2 , ℓ_∞), CFE (ℓ_2), and p-CFE ℓ_0 examples.

Results. From Tab. 1, learning from p-CFEs matches learning from perturbations (with ℓ_2 and ℓ_{∞} PGD) as well as CFE ℓ_2 on achieving comparable training accuracy to standard training, thus extending empirical findings of Kumano et al. (2024a) to p-CFEs. Note that except for standard training, the models are trained on perturbed images with target incorrect labels (cf. Theorem 4.1), and the training accuracy was evaluated on the original training images and labels. Moreover, Tab. 2 shows that on WaterBirds and SpuCoAnimals, training with p-CFEs substantially boosts worst-group accuracy—surpassing even standard (noise-free) training by 12 % on WaterBirds. This indicates that p-CFE-trained models rely less on spurious background features. Fig. 3 confirms that models trained on other perturbations focus heavily on spurious backgrounds, while p-CFE training shifts attention to relevant features (e.g., the bird and the dog), effectively mitigating spurious correlations.

6 Conclusion and Discussion

We showed that training with p-CFEs provides a compelling alternative to adversarial perturbations, guiding models toward semantically meaningful features and reducing reliance on spurious correlations. Our approach is data-efficient, model-agnostic, and requires no group labels. Future work includes scaling to higher-dimensional datasets using more expressive density estimators, and extending to large models such as LLMs and vision-language models (VLMs). Further, combining the theoretical framework of learning from adversarial perturbations in (Kumano et al., 2024a,b) with the connection between adversarial perturbations and p-CFEs (Pawelczyk et al., 2022; Freiesleben, 2022) can justify our observations.

Acknowledgement

This research was partially supported by the DFG Cluster of Excellence MATH+ (EXC-2046/1, project id 390685689) funded by the Deutsche Forschungsgemeinschaft (DFG) as well as by the German Federal Ministry

of Education and Research (fund number 01IS23025B). Hiroshi Kera was supported by JSPS KAKENHI Grant Number JP23KK0208.

References

- Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM journal on imaging sciences, 2(1):183–202, 2009.
- Li Deng. The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE signal processing magazine*, 29(6):141–142, 2012.
- Timo Freiesleben. The intriguing relation between counterfactual explanations and adversarial examples. Minds and Machines, 32(1):77–109, 2022.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. Advances in neural information processing systems, 32, 2019.
- Siddharth Joshi, Yu Yang, Yihao Xue, Wenhan Yang, and Baharan Mirzasoleiman. Towards mitigating more challenging spurious correlations: A benchmark & new datasets. arXiv preprint arXiv:2306.11957, 2023.
- Ehsan Kazemi, Thomas Kerdreux, and Liqiang Wang. Minimally distorted structured adversarial attacks. International Journal of Computer Vision, 131(1):160–176, 2023.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. University of Toronto, 2009.
- Soichiro Kumano, Hiroshi Kera, and Toshihiko Yamasaki. Theoretical understanding of learning from adversarial perturbations. International Conference on Learning Representations, 2024a.
- Soichiro Kumano, Hiroshi Kera, and Toshihiko Yamasaki. Wide two-layer networks can learn from adversarial perturbations. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, Advances in Neural Information Processing Systems, volume 37, pages 59755–59807. Curran Associates, Inc., 2024b.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In Proceedings of the IEEE/CVF international conference on computer vision, pages 10012–10022, 2021.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *International Conference on Learning Representations*, 2018a.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018b.
- Martin Pawelczyk, Chirag Agarwal, Shalmali Joshi, Sohini Upadhyay, and Himabindu Lakkaraju. Exploring counterfactual explanations through the lens of adversarial examples: A theoretical and empirical analysis. In *International Conference on Artificial Intelligence and Statistics*, pages 4574–4594. PMLR, 2022.
- Shpresim Sadiku, Moritz Wagner, Sai Ganesh Nagarajan, and Sebastian Pokutta. S-cfe: Simple counterfactual explanations. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2025a.
- Shpresim Sadiku, Moritz Wagner, and Sebastian Pokutta. Gse: Group-wise sparse and explainable adversarial attacks. In *Proceedings of the International Conference on Learning Representations*, 2025b.

- Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *International Conference on Learning Representations*, 2020a.
- Shiori Sagawa, Aditi Raghunathan, Pang Wei Koh, and Percy Liang. An investigation of why overparameterization exacerbates spurious correlations. In *International Conference on Machine Learning*, pages 8346–8356. PMLR, 2020b.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In Proceedings of the IEEE international conference on computer vision, pages 618–626, 2017.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *International Conference on Learning Representations*, 2014.
- Arnaud Van Looveren and Janis Klaise. Interpretable counterfactual explanations guided by prototypes. In Joint European Conference on Machine Learning and Knowledge Discovery in Databases, pages 650–665. Springer, 2021.
- Sandra Wachter, Brent Mittelstadt, and Chris Russell. Counterfactual explanations without opening the black box: Automated decisions and the gdpr. *Harv. JL & Tech.*, 31:841, 2017.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. arXiv preprint arXiv:1708.07747, 2017.
- Wang Xiaosen, Kangheng Tong, and Kun He. Rethinking the backward propagation for adversarial transferability. Advances in Neural Information Processing Systems, 36:1905–1922, 2023.
- Kaidi Xu, Sijia Liu, Pu Zhao, Pin-Yu Chen, Huan Zhang, Quanfu Fan, Deniz Erdogmus, Yanzhi Wang, and Xue Lin. Structured adversarial attack: Towards general implementation and better interpretability. In International Conference on Learning Representations, 2019.
- Yuzhe Yang, Haoran Zhang, Dina Katabi, and Marzyeh Ghassemi. Change is hard: A closer look at subpopulation shift. arXiv preprint arXiv:2302.12254, 2023.
- Wenqian Ye, Guangtao Zheng, Xu Cao, Yunsheng Ma, and Aidong Zhang. Spurious correlations in machine learning: A survey. arXiv preprint arXiv:2402.12715, 2024.
- Songming Zhang, Xiaofeng Chen, Shiping Wen, and Zhongshan Li. Density-based reliable and robust explainer for counterfactual explanation. *Expert Systems with Applications*, 226:120214, 2023.



Figure 4: **Top row**: Original and Grad-CAM visualizations for a misclassified *landbird* (with a water background) from the WaterBirds dataset—incorrectly predicted as a *waterbird*. **Bottom row**: Original and Grad-CAM visualizations for a misclassified *small dog* (with an outdoor background) from the SpuCoAnimals dataset—incorrectly predicted as a *big dog*.

Appendix

A Spurious Correlations - Additional Examples

Fig. 4 illustrates additional examples of spurious correlations. In the top row, the model relies on the presence of water to classify waterbirds. In the bottom row, it associates the outdoor background with the presence of big dogs.

B Learning from p-CFEs - Additional Experiments

In this section, we extend the work of Kumano et al. (2024a) on learning from adversarial perturbations to traditional ℓ_2 CFEs from Wachter et al. (2017), as well as the more recent ℓ_0 p-CFEs proposed by Sadiku et al. (2025a). We generate adversarial examples using Projected Gradient Descent (PGD) (Madry et al., 2018b) with cross-entropy loss under varying norms (ℓ_2, ℓ_∞). CFEs are constructed by minimizing the cross-entropy loss regularized by the unweighted squared Euclidean distance, controlled by the tradeoff parameter λ . We denote λ_{CF} as the learning rate used by the Adam optimizer during CFE optimization. For p-CFE ℓ_0 , we define L as the Lipschitz constant, and λ_{steps} as the number of search steps.

A Results on Artificial Data

We experiment on 2D dim data generated from uniform or Gaussian distribution. The model used for these experiments is a one-layer neural network (Kumano et al., 2024a). Figs. 5 to 12 compare the accuracies of models trained with adversarial examples versus noise-augmented data, across varying input dimensions and numbers of natural and adversarial samples. For CFE ℓ_2 and p-CFE ℓ_0 , we additionally vary the ratio of modified pixels, denoted by d_{δ}/d where d_{δ} is the number of modified pixels and d is the total number of pixels. These results extend the findings of Kumano et al. (2024a) to settings involving both traditional CFEs and p-CFEs.



Figure 5: Comparison of the accuracy of the model trained on CFE ℓ_2 perturbations and noise trained model on the clean dataset. Data was acquired from a **uniform** distribution. The hyper-parameters $\lambda = 0.001$ and $\lambda_{CF} = 0.01$ were used. The algorithm was run for 0.05d iterates, where d is the input dimension.



Figure 6: Comparison of the accuracy of the model trained on CFE ℓ_2 perturbations and noise trained model on the clean dataset. Data was acquired from a **Gaussian** distribution. The hyper-parameters $\lambda = 0.001$ and $\lambda_{CF} = 0.01$ were used. The algorithm was run for 0.05d iterates, where d is the input dimension.



Figure 7: Comparison of the accuracy of the model trained on p-CFE ℓ_0 perturbations and noise trained model on the clean dataset. Data was acquired from a **uniform** distribution. The hyper-parameters $\lambda_{steps} = 5$, L = 1.0 and 0.05d iterations were considered. Here, λ_{steps} and L denote the number of search steps and step size, respectively.



Figure 8: Comparison of the accuracy of the model trained on p-CFE ℓ_0 perturbations and noise trained model on the clean dataset. Data was acquired from a **Gaussian** distribution. The hyper-parameters $\lambda_{steps} = 5, L = 1.0$ and 0.05d iterations were considered. Here, λ_{steps} and L denote the number of search steps and step size, respectively.



Figure 9: Comparison of the accuracy of the model trained on PGD ℓ_2 perturbations and noise trained model on the clean dataset. Data was acquired from a **uniform** distribution. The hyper-parameter $\epsilon = 0.78$ was used for most of the graphs (top-left, top-right, bottom-left). For the graph at bottom-right, multiple perturbation constraints, along with multiple input dimensions, were used. All the attacks were executed for 100 iterations.



Figure 10: Comparison of the accuracy of the model trained on PGD ℓ_2 perturbations and noise trained model on the clean dataset. Data was acquired from a **Gaussian** distribution. The hyper-parameter $\epsilon = 0.78$ was used for most of the graphs (top-left, top-right, bottom-left). For the graph at bottom-right, multiple perturbation constraints, along with multiple input dimensions, were used. All the attacks were executed for 100 iterations.



Figure 11: Comparison of the accuracy of the model trained on PGD ℓ_{∞} perturbations and noise trained model on the clean dataset. Data was acquired from a **uniform** distribution. The hyper-parameter $\epsilon = 0.03$ was used, along with 100 iterations. The hyper-parameter $\epsilon = 0.03$ was used for most of the graphs (top-left, top-right, bottom-left). For the graph at bottom-right, multiple perturbation constraints were used.



Figure 12: Comparison of the accuracy of the model trained on PGD ℓ_{∞} perturbations and noise trained model on the clean dataset. Data was acquired from a **Gaussian** distribution. The hyper-parameter $\epsilon = 0.03$ was used, along with 100 iterations. The hyper-parameter $\epsilon = 0.03$ was used for most of the graphs (top-left, top-right, bottom-left). For the graph at bottom-right, multiple perturbation constraints were used.

B Results on Original (Natural) Datasets

A convolutional neural network was used for the MNIST (Deng, 2012) and Fashion-MNIST (FMNIST) (Xiao et al., 2017) datasets, while a WideResNet was used for CIFAR-10 (Krizhevsky et al., 2009), following the setup of Kumano et al. (2024a). We denote deterministic target labels by D and random target labels by R. The learning rate for the Adam optimizer used in CFE generation is set to $\lambda_{CF} = 0.01$, unless stated otherwise. All algorithms were run for 100 iterations by default.

For each dataset, we first train a model on the clean training set, then use it to generate adversarial samples or CFEs from that same set. A second model is then trained on the perturbed data. Figs. 13 to 21 compare the training and validation accuracies of models trained on clean versus perturbed data. Validation accuracy refers to performance on the clean validation set, while training accuracy reflects performance on the perturbed training data. These results extend the findings of Kumano et al. (2024a) to settings involving both traditional CFEs and p-CFEs for standard benchmark datasets.

B.1 MNIST



Figure 13: Standard training and validation accuracy of a simple convolutional neural network model trained on MNIST dataset for 100 epochs.



Figure 14: Training and Validation accuracies for models trained on MNIST counterfactuals generated using p-CFE ℓ_0 and CFE ℓ_2 . For CFE ℓ_2 , we used $\lambda = 0.1$ (ℓ_2 (regularization strength)) and $\lambda_{CF} = 0.01$ (step-size for Adam Optimizer). For p-CFE ℓ_0 , we used a Lipschitz constant L = 1e-4 and 5 search steps. Both methods were run for 100 epochs.



Figure 15: Training and Validation accuracies for models trained on MNIST adversarial samples generated by PGD ℓ_{∞} and ℓ_2 . The hyper-parameters $\epsilon = 2$ and $\epsilon = 0.3$ were used for the PGD ℓ_2 and ℓ_{∞} attacks, respectively. Both methods were run for 100 epochs.

B.2 FMNIST



Figure 16: Standard training and validation accuracy of a simple convolutional neural network model trained on FMNIST dataset for 200 epochs.



Figure 17: Training and Validation accuracies for models trained on FMNIST adversarial samples generated using CFE ℓ_2 and p-CFE ℓ_0 . For CFE we used $\ell_2 \lambda = 0.01$ and $\lambda_{CF} = 0.01$; for p-CFE ℓ_0 , a Lipschitz constant L = 1e-5 and 5 search steps were considered. Both methods were run for 100 epochs.



Figure 18: Training and Validation accuracies for models trained on FMNIST adversarial samples generated by PGD ℓ_{∞} and ℓ_2 . The hyper-parameters $\epsilon = 2$ and $\epsilon = 0.3$ were considered for PGD ℓ_2 and ℓ_{∞} attack, respectively. Both methods were run for 100 epochs.

B.3 CIFAR-10



Figure 19: Standard training and validation accuracy of a wide ResNet model trained on the CIFAR10 dataset for 200 epochs.



Figure 20: Training and Validation accuracies for models trained on CIFAR10 counterfactuals generated by CFE ℓ_2 and p-CFE ℓ_0 . The hyper-parameters $\lambda = 0.01$ and $\lambda_{CF} = 0.001$ were considered for CFE ℓ_2 , and for p-CFE ℓ_0 , Lipschitz constant L = 1e-4 with 4 search steps were considered. Both methods were run for 100 epochs.



Figure 21: Training and Validation accuracies for models trained on CIFAR10 adversarial samples generated by PGD ℓ_{∞} and ℓ_2 . The hyper-parameters $\epsilon = 0.5$ and $\epsilon = 0.1$ were considered for PGD ℓ_2 and ℓ_{∞} attack, respectively. Both methods were run for 100 epochs.