

Reactive Model Correction: Mitigating Harm to Task-Relevant Features via Conditional Bias Suppression

Dilyara Bareeva¹, Maximilian Dreyer¹, Frederik Pahde¹, Wojciech Samek^{1,2,3,†}, Sebastian Lapuschkin^{1,†}

¹ Fraunhofer Heinrich Hertz Institute, ² Technical University of Berlin,

³ BIFOLD – Berlin Institute for the Foundations of Learning and Data

[†]corresponding authors: {wojciech.samek | sebastian.lapuschkin}@hhi.fraunhofer.de

Abstract

Deep Neural Networks are prone to learning and relying on spurious correlations in the training data, which, for high-risk applications, can have fatal consequences. Various approaches to suppress model reliance on harmful features have been proposed that can be applied post-hoc without additional training. Whereas those methods can be applied with efficiency, they also tend to harm model performance by globally shifting the distribution of latent features. To mitigate unintended overcorrection of model behavior, we propose a reactive approach conditioned on model-derived knowledge and eXplainable Artificial Intelligence (XAI) insights. While the reactive approach can be applied to many post-hoc methods, we demonstrate the incorporation of reactivity in particular for P-ClArC (Projective Class Artifact Compensation), introducing a new method called R-ClArC (Reactive Class Artifact Compensation). Through rigorous experiments in controlled settings (FunnyBirds) and with a real-world dataset (ISIC2019), we show that introducing reactivity can minimize the detrimental effect of the applied correction while simultaneously ensuring low reliance on spurious features.

1. Introduction

Modern Deep Neural Network (DNN) architectures yield remarkable results for a plethora of complex tasks, including high-stake applications, such as medicine [7], finance [35], or criminal justice [42]. However, it has been shown that DNNs are at risk of learning shortcuts based on spurious correlations due to imperfections in the available training data [3, 16, 24, 29], compromising the reliability of these models in high-risk scenarios. Some notable examples include melanoma detection models using visible band-aids as evidence *against* cancerous melanoma, as they only occurred next to benign lesions in the training data [33], or bone age prediction models exploiting the spu-

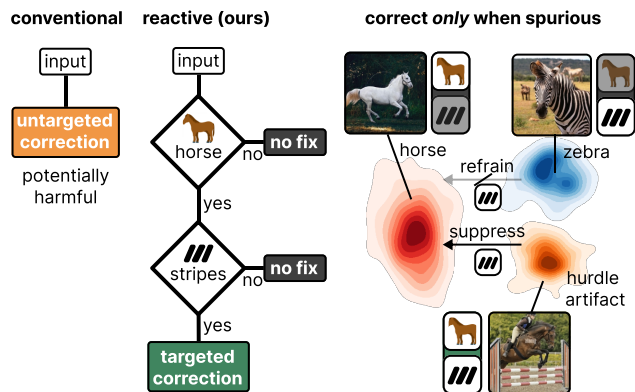


Figure 1. Reactive Model Correction: Whereas traditional post-hoc model correction approaches are applied to all samples *uniformly*, we propose *conditional* suppression of artifacts. One possible condition for triggering correction is the combination of a specific class prediction and the presence of a spurious feature (*left*). This prevents the suppression of concepts when unnecessary or even harmful: When correcting, *e.g.*, for a “hurdle”-artifact (related to “stripe” features), we refrain from suppression for zebra samples, as stripe textures are now valid discriminative features and crucial to discerning zebras from horses (*right*).

rious correlation between the size of lead markers and bone age caused by the specifics of data processing [30]. Another example is the usage of hurdles as features for horse classification [23], as illustrated, *e.g.*, in Fig. 1 (*bottom right*).

To address these model weaknesses arising from compromised training data, a multitude of model correction approaches have been suggested in recent years to unlearn undesired model behavior, either by model re-training on modified data [43, 44], model fine-tuning [14, 33, 34, 38], or by post-hoc model editing [3, 11, 27, 32, 37]. Whereas the former two groups of approaches require access to the training data and are expensive in terms of computational resources, therefore often being infeasible to apply for large models, the latter group provides a cheap solution to modify existing models before deployment, or even during in-

ference time. Commonly, post-hoc model correction approaches model the shortcut concept as a linear direction or subspace in latent space, and modify latent representations or model parameters such that predictions become invariant towards the modeled artifact direction.

However, while such approaches successfully reduce the impact of data artifacts on model predictions, they tend to negatively impact the model performance for “clean” data samples without artifacts, as spurious concepts are suppressed globally (*i.e.*, *always*) due to a constant modification of the model parameters or the elimination of latent subspaces. When artifact representations are modeled inaccurately or entangled with other features, artifact suppression can lead to unintended and harmful suppression of important task-relevant features [22]. For instance, when suppressing the direction encoding the “hurdle” concept in horse detectors, it also affects related concepts like the direction for “stripes”, potentially impacting the model’s ability to differentiate between horses and zebras.

In order to minimize “collateral damage”, we introduce a *reactive* model correction paradigm, where corrective measures are applied during the inference only if certain conditions are met. Such conditions can include the prediction of a certain class and/or presence and relevance of the concept to the model decision-making. To that end, we leverage insights from local and global eXplainable Artificial Intelligence (XAI) methods to understand and recognize the role of certain concepts (*e.g.*, data artifacts) for a single prediction and reactively update the prediction if required, while leaving other predictions unaffected, as illustrated in Fig. 1.

We empirically demonstrate the benefits of our approach compared to non-reactive post-hoc model corrections using the ISIC2019 [8, 9, 28] and FunnyBirds [19] datasets, with both controlled and real-world data artifacts. Our contributions include the following:

1. We introduce *reactive* model correction approach that *only* corrects model behavior when required, *e.g.*, for a specific class and/or when an artifact is predicted to be present.
2. We demonstrate the entanglement of concept representations of artifacts and non-artifacts.
3. We show the superior performance of reactive model correction in quantitative evaluation using controlled and real-world data artifacts.

2. Related Work

Methods for model correction either require full retraining, fine-tuning or can be applied post-hoc. Whereas full retraining is usually necessary for methods that change the data distribution [41, 43, 44], fine-tuning often involves the regularization of a model by applying dedicated loss terms [14, 33, 34]. However, the creation of a representative training data set often comes with its own set of challenges,

as the data cleaning process is not always straightforward and involves significant manual labor [26], and can introduce new biases [12]. (Partial) retraining may be infeasible due to high computational costs and can exacerbate unlabeled shortcuts [25]. In light of the challenges, post-hoc correction methods [5, 6, 32, 37], including the Projective ClArC (P-ClArC) [3] method, represent a preferable alternative. Theoretical and empirical analyses presented in [22] demonstrate that post-hoc methods, specifically INLP [32], have the potential to eliminate not only spurious but also task-relevant features. In this work, we introduce a general framework to address this issue by applying post-hoc model correction only under specific conditions, minimizing collateral damage caused by the suppression of unintended, correlated concepts.

3. Reactive Model Correction Framework

Let $f : \mathcal{X} \rightarrow \mathbb{R}^D$ be a neural network, mapping input samples $x \in \mathcal{X}$ to an output for D labels. Given a hidden layer l of network f with m neurons, we define the *feature extractor* $a : \mathcal{X} \rightarrow \mathcal{A} \subseteq \mathbb{R}^m$ providing the latent activations of layer l , and the *model head* $\hat{f} : \mathcal{A} \rightarrow \mathbb{R}^D$, mapping activations to the output.

We further introduce the notion of *artifacts* as a set \mathcal{C} of concepts with a size $|\mathcal{C}| = n$ that are present in the data but should not be relied upon by the network f (*e.g.*, watermarks or hurdles). For each artifact $c_i \in \mathcal{C}$, $i \in \{1, \dots, n\}$, we are provided with sets of samples with the concepts present $X_1^+, X_2^+, \dots, X_n^+$ (*positive* examples) and corresponding sets without the concepts $X_1^-, X_2^-, \dots, X_n^-$ (*negative* examples). For each set, corresponding activation vectors $\mathcal{A}_i^+ = \{a(x) \in \mathbb{R}^m | x \in X_i^+\} \subseteq \mathcal{A}$ and $\mathcal{A}_i^- = \{a(x) \in \mathbb{R}^m | x \in X_i^-\} \subseteq \mathcal{A}$ are obtained. A wide-spread assumption is that concepts learned and used by a network are encoded as Concept Activation Vectors (CAVs), *i.e.* characteristic directions or subspaces in the latent space \mathcal{A} [20]. Let v_i denote a CAV for concept $c_i \in \mathcal{C}$ pointing from \mathcal{A}_i^- to \mathcal{A}_i^+ .

3.1. Reactive Model Correction

The overall objective of model correction is to *suppress* the influence of artifacts on the model decision-making. Post-hoc model correction methods commonly aim to *globally* and *simultaneously* address the set of all known artifacts, leading to a two-fold problem: Firstly, a concept might be considered spurious for one class, but encodes a valid strategy for another, as, *e.g.*, the “stripes” concept in Fig. 1 (*right*) represents a valid feature for zebras but a shortcut artifact for horses. Secondly, the artifact concept might be entangled with other useful features in latent space. Consequently, artifact suppression also leads to a distorted representation of valid features. Both problems potentially harm

Table 1. Examples of reactive conditions suitable for reactive model correction and their respective informing XAI methods.

	Reactive Condition	Description	Informing XAI Methods
less granular ↓	Predicted class	Correction is only applied when a class that is associated with an artifact is predicted.	none needed
	Artifact presence	Correction is only initiated if the presence of the artifact has been detected in the sample.	Concept Activation Vector (CAV) [20]
more granular ↓	Artifact relevance	Correction is only initiated if the artifact has been identified as relevant for the prediction.	attributions for input [4] or latent [1, 15] features
	Prediction strategy	A sample is categorized into one of the known prediction sub-strategies (clusters), <i>e.g.</i> based on activations and attributions of hidden units. The reactive condition is triggered, if the cluster has been previously identified as undesirable.	PCX[13], strategic clusters[15]
	Domain-expert-in-the-loop	Triggered manually by a domain expert upon inspection of the information provided by XAI methods, including heatmap visualizations, active and identified prediction strategies, <i>etc.</i>	All of the above

model performance (as measured in Sec. 4.3), and motivate a *conditional*, more targeted application of bias suppression.

To address these challenges, we propose *reactive* (post-hoc) model correction. Essentially, this paradigm aims to initiate model correction only under specific conditions. In Tab. 1 we present an overview of potential reactive conditions, including, *e.g.*, the prediction of a specific class, the actual presence of an artifact, or a decision by an expert, as well as XAI methods that can be used to measure whether the pre-defined conditions are met.

Let us formalize a condition (as, *e.g.*, given in Tab. 1) as a condition-generating function:

Definition 3.1 (Condition-generating function) *Given a network f and the set \mathcal{C} of known artifact concepts, we define a condition-generating function $r: \mathcal{X} \rightarrow \mathcal{C}^* \subseteq \mathcal{C}$ as a function, that, for any given input, produces a (reduced) set \mathcal{C}^* of artifacts to be removed.*

The core principle of the reactive model correction framework involves the identification of the artifacts to be corrected for a given sample, based on the condition-generating function. This typically involves a forward pass followed by the computation of the condition-generating function, potentially complemented by expert evaluation. Once the artifacts are determined, if any, a subsequent (partial) forward pass is performed with model correction applied specifically to address these identified artifacts.

While many post-hoc methods, including classifier editing [37], LEACE [5], INLP [32], and P-ClArC [3], can be employed for model correction in the final step, in this work

we demonstrate and evaluate our reactive approach within the P-ClArC framework.

3.2. Reactive ClArC

In the subsequent section, we introduce a reactive variant of P-ClArC, an inference-time model correction method suppressing artifact signals modeled as linear direction in latent space. Two core assumptions constitute the foundation of the P-ClArC method. Firstly, it is assumed that introducing an artifact to a sample in the input domain leads to an increase in the activations along the corresponding CAV direction in the latent space [20]. Originally, linear classifiers, such as Support Vector Machines (SVMs) [10], are trained on latent features to estimate CAVs. Recently, Pahde *et al.* [31] have demonstrated that pattern-based CAVs yield more precise concept representations and therefore, when incorporated into the ClArC framework, superior performance compared to linear classifiers. The second assumption posits that all other concepts are encoded in directions orthogonal to the artifact direction. Consequently, this further implies that there is no variance in the artifact direction for all non-artifact samples.

These two assumptions lead to the P-ClArC *backward* artifact model which suppresses the artifact features in the latent activations.

Definition 3.2 (P-ClArC) *Let $c_i \in \mathcal{C}$ be an artifact, v_i its respective CAV, and $z_i^- = \frac{1}{|X_i^-|} \sum_{z \in X_i^-} a(z)$ the mean activation over non-artifact samples. Then, the P-ClArC artifact backward model $h: \mathcal{A} \rightarrow \mathcal{A}$ for an activation vector*

$a_x = a(x)$ is defined as follows:

$$h(a_x, c_i) = a_x - v_i v_i^T (a_x - z_i^-). \quad (1)$$

The backward artifact model h effectively corrects the output of the feature extractor a during the inference process, before the model head \tilde{f} is applied. We can extend this formulation to accommodate the correction of multiple artifacts simultaneously by defining an *artifact subspace* as the space spanned by all artifact CAVs. We then project onto the subspace orthogonal to this artifact subspace:

Definition 3.3 (Multi-Artifact P-ClArC) Let $C' \subseteq C$ represent a subset of artifacts with a size $|C'| = k$, $1 \leq k \leq n$. Let $V_{C'} = [v_i]_{c_i \in C'}$ be the matrix comprised of the respective CAVs as column vectors. Let $Z_{C'}^- = \bigcap_{c_i \in C'} X_i^-$ be the intersection of negative examples, and $\bar{z}_{C'} = \frac{1}{|Z_{C'}^-|} \sum_{z \in Z_{C'}^-} a(z)$. Then, the multi-artifact P-ClArC artifact backward model $\tilde{h} : \mathcal{A} \times C \rightarrow \mathcal{A}$ for an activation vector $a_x = a(x)$ is defined as follows:

$$\tilde{h}(a_x, C') = a_x - V_{C'} (V_{C'}^T V_{C'})^{-1} V_{C'}^T (a_x - \bar{z}_{C'}^-). \quad (2)$$

with the derivation given in Appendix A.1.

While the P-ClArC method performs effectively under the given assumptions, these conditions are rarely encountered in practice, as motivated in Fig. 1. This is confirmed by experiments in Sec. 4.3, illustrating that artifact CAVs exhibit strong correlations with other concept directions. Consequently, suppressing an artifact can inadvertently alter the representations of other potentially valid and important concepts, whether they are present or absent from a sample, thus impacting decision-making processes. Therefore, we suggest integrating reactivity into the P-ClArC framework. In particular, we apply a condition-generating function to identify the set of artifacts to be suppressed prior to applying the P-ClArC correction:

Definition 3.4 (R-ClArC) For a condition-generating function r and an activation vector $a_x = a(x)$, we define the Reactive ClArC (R-ClArC) model $h_r : \mathcal{A} \rightarrow \mathcal{A}$ as follows:

$$h_r(a_x) = \begin{cases} \tilde{h}(a_x, r(x)), & r(x) \neq \emptyset \\ a_x, & \text{otherwise} \end{cases}. \quad (3)$$

In the following, we discuss and evaluate the detailed implementation of R-ClArC for two reactive conditions outlined in Tab. 1: *class-conditional* R-ClArC based on *label prediction* condition-generating functions, and *artifact-conditional* R-ClArC given by *artifact presence* condition-generating functions.

Definition 3.5 (Class-condition-generating function)

Assume that for every artifact $c \in C$, we are provided with a set of output labels $\mathcal{R}_c \subseteq [D] = \{1, 2, \dots, D\}$. A class-condition-generating function is then as follows:

$$r(x) = \{c \in C \mid \arg \max_{d \in [D]} f^{(d)}(x) \in \mathcal{R}_c\}, \quad (4)$$

where $\arg \max_{d \in [D]} f^{(d)}(x)$ corresponds to the predicted output class of sample x .

Definition 3.6 (Artifact-condition-generating function)

Assume that for every artifact $c \in C$, we are provided with a binary classifier $t_c : \mathcal{X} \rightarrow \{0, 1\}$ that outputs 1 if the artifact c is present in a sample $x \in \mathcal{X}$, and 0 otherwise. An artifact-condition-generating function is then as follows:

$$r(x) = \{c \in C \mid t_c(x) = 1\}. \quad (5)$$

In Appendix A.2, we present pseudocode outlining the algorithms for P-ClArC and R-ClArC, and in Appendix A.3 we provide a 3D toy example to illustrate these methods.

4. Experiments

We address the following research questions:

1. **(Q1)** What is the degree of *dependence* between representations of artifact and non-artifact concepts within models? (Sec. 4.3)
2. **(Q2)** How do the *effectiveness* and the degree of *collateral damage* caused by the reactive approach compare to the traditional model correction approach? (Sec. 4.4)

4.1. Experimental Details

We conduct experiments in two controlled settings using toy datasets and in one setting utilizing real-world artifacts in a benchmark dataset. In the first controlled setup, we generate a synthetic FunnyBirds dataset [19], consisting of two bird classes. We insert a *backdoor* [17] artifact (“green box”, see Fig. 2) into 33% of the training samples originating from class 0 and flip their labels to class 1. This training configuration forces the model always to predict class 1 in the presence of the inserted “green box” object.

In the second setup, we create a FunnyBirds dataset comprising ten classes. We randomly select ten different background object artifacts and insert a random number of these artifacts into 50% of training samples belonging to class 0. Since the artifacts are exclusively present in class 0, we expect a model trained on this dataset to utilize them as *shortcuts* [16] for class 0.

For the real-world dataset, we utilize ISIC2019 [8, 9, 28], a dermatologic dataset designed for skin cancer detection, featuring images of both benign and malignant lesions. Using the Reveal2Revise framework [30], we identify three

artifacts naturally occurring in the dataset, strongly correlating with class labels: “band-aid” (correlating with “NV”), “skin marker” (“NV”, “BKL”), and “reflection” (“BKL”) artifacts. To evaluate our approach, we additionally create a *poisoned* test set for ISIC2019, where all samples have a “reflection” artifact superimposed on them. Visually the “reflection” artifact simulates the spot-wise reflection of a bright illumination source. Additional details for all the datasets are available in Appendix B.1.

For all settings, we train VGG16 [39], ResNet18 [18], and EfficientNet-B0 [40] models, with detailed training information provided in Appendix B.2.

4.2. CAV Calculation

For each experimental setting, we generate datasets of the same images with and without artifacts. For both FunnyBirds datasets and each artifact, we generate pairs of images with and without the artifact. For the ISIC dataset and the “reflection” artifact, we superimpose examples of the artifact onto clean images from the associated class (refer to Appendix B.1 for details). Fig. 2 displays some examples from the generated sets.

This process allows us to compute precise CAVs using sets of positive and negative samples. We refer to CAVs obtained through this method as “generated”. As compared to the conventional method of calculating CAVs using the subsets from the dataset, we expect that the generated CAVs will provide more precise descriptions of the concept direction. Furthermore, the generated sets allow us to calculate pairwise concept directions. By comparing concept directions with CAVs, we assess the faithfulness of these concept representations in Sec. 4.3.

While each pair of images only differs in the presence or absence of an artifact, the difference between the activation vectors of the two samples may still not accurately represent the true concept direction. For instance, as depicted in the first example in Fig. 2, adding a “green box” artifact significantly obscures the beak concept. Therefore, we can expect that the activation vectors of the two samples also differ in their expression of the beak concept. However, for CAV calculation, we use large sets of pairs, which helps mitigate this effect. Further details about the CAV computation procedure can be found in Appendix B.3.

4.3. Orthogonality of Concept Representations

In Sec. 3.1, we introduced the reactive framework for model correction. Regardless of how precise and disentangled a concept representation used by a model correction method may be, the reactive approach can prove beneficial, *e.g.* in cases when an artifact irrelevant for a certain class represents a useful feature for another (*e.g.*, “stripes”). However, many model correction methods may also face challenges related to the entanglement of concept representations. In




Figure 2. Examples of adding artifact concepts in a controlled manner: (Left): For FunnyBirds, we insert a “green box” into images. (Right): For ISIC, we insert “reflections” on the side.

CAV Dataset	FunnyBirds (“green box”)		ISIC2019 (“reflection”)	
	Pattern	Filter	Pattern	Filter
Generated	0.824	0.101	0.617	0.343
Data Subset	0.563	0.042	0.469	0.328

Table 2. Evaluating the alignment of computed artifact CAVs in terms of cosine similarity with the actual change in activations when the concept is added in controlled fashion in ResNet18 models. We compare CAVs computed on original data subsets, and pairs of clean and (generated) poisoned samples.

this context, the reactive approach also enables us to “minimize the damage” on non-artifact samples.

In the following experiments, we evaluate the assumption of concept orthogonality in two controlled settings. Firstly, we investigate the concept representation of the “green box” artifact in the ResNet18 model trained on the backdoored FunnyBirds dataset. Secondly, we analyze the concept representation of the “reflection” artifact in the ResNet18 model trained on ISIC2019. Specifically, we use CAVs as concept representations since they are utilized in both P-CIArC and our proposed approach R-CIArC.

Initially, we assess the faithfulness of the CAV concept representation by comparing them to the pairwise concept directions. For both settings, we compute filter- and pattern-based CAVs using both the generated sets and sampled dataset examples. Then, we determine the alignment score by averaging the cosine similarity of each concept direction and CAV, following the approach in [31]. The results are presented in Tab. 2. Confirming the findings of [31], pattern-based CAV performed better in both settings. Additionally, as anticipated, the generated CAVs offer a more accurate concept representation. We utilize the pattern-based CAV calculated on the generated sets for our subsequent experiments. Additional details regarding the evaluation of CAVs can be found in Appendix B.4.

In our experiment investigating the orthogonality hypothesis, we analyze the distribution of the CAV activations of 500 randomly selected clean samples and all samples containing the artifact. Based on an assumption of the

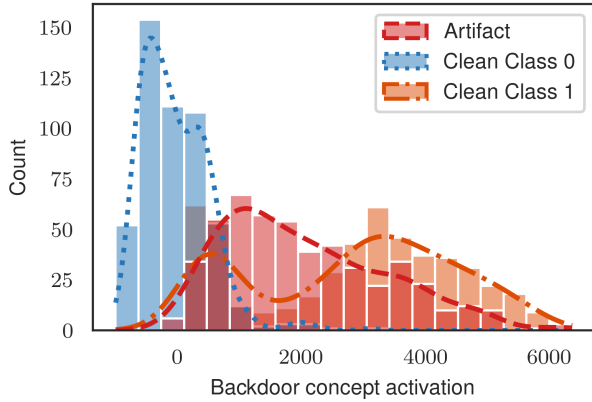


Figure 3. Histogram of CAV activations for the FunnyBirds backdoor (“green box”) artifact in ResNet18: the backdoor CAV aligns with features specific for class 1.

P-CIArC framework, there should be no variance along the artifact direction for all clean samples.

From Fig. 3, we can observe the distribution of pattern-based CAV activations for the backdoor FunnyBirds model. Clean samples exhibit considerable variance along the CAV direction. Specifically, samples belonging to class 1 notably activate the backdoor CAV, indicating entanglement between the backdoor concept and the concept of the class label which is predicted in the backdoor presence. Consequently, suppressing the artifact direction in this scenario coincidentally shifts the clean samples of class 1 towards or across the decision boundary towards class 0.

Fig. 4 illustrates the distribution of clean and “reflection”-artifact ISIC samples along the CAV direction. Firstly, clean samples exhibit variance along the artifact CAV, indicating that the direction encodes information unrelated to the artifact. Secondly, we observe a high level of activation of the concept in samples displaying features related to the “reflection” artifact, such as pale skin or white skin patches. These features could be important for prediction, suggesting that suppressing this information may not be desirable. Fig. 5 further illustrates the cosine similarities between the CAV directions of the three ISIC artifacts and the class directions of clean samples. The non-orthogonality between the CAV and the class directions suggests that suppressing the artifact directions could impact the predictions of clean samples.

4.4. Model Correction with R-CIArC

In the following analysis, we examine the impact of incorporating reactivity into the P-CIArC model correction method, based on three metrics: accuracy, F1 score, and artifact relevance, as measured by the share of Layer-wise Relevance Propagation (LRP) [4] attribution in the

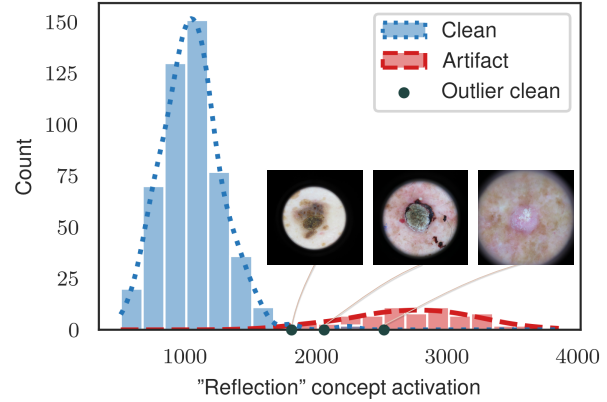


Figure 4. Histogram of CAV activations for the ISIC “reflection” artifact in ResNet18: outlier clean samples with white spots lead to high concept activation.




Figure 5. Cosine similarity between artifact CAV and mean feature direction of each class for the ISIC dataset and ResNet-18: the artifact concept representations are entangled with clean features.

artifact region. Our evaluation includes R-CIArC with class- and artifact-condition-generating functions, as well as their combination. In all settings except for the poisoned ISIC2019, our primary focus is on evaluating the performance on clean samples. We expect the reactive approach to improve the preservation of performance on clean samples while still reducing artifact relevance. In the poisoned ISIC2019 setting, where clean samples are not present, we are mainly interested in recovering performance compared to the original Vanilla model. Additionally, we provide accuracy and F1 scores for artifact samples. Additional details regarding the CIArC parameters and the evaluation procedure are outlined in Appendix B.5.

The outcomes of model correction for VGG16, ResNet18, and EfficientNet-B0 on the FunnyBirds datasets are presented in Tab. 3. In the *backdoor* setting, the R-CIArC class condition is met for the prediction of class 1 and in the *shortcut* setting for the prediction of class 0. As expected, the *backdoor* Vanilla models exhibit low performance on the backdoor samples. As mentioned in Sec. 4.3, in ResNet18, the artifact and class 1 concepts are correlated. Suppressing the artifact often leads to predictions of class

Model	Method	Condition	Accuracy (clean)	Accuracy (artifact)	F1 (clean)	F1 (artifact)	Artifact relevance
Resnet18	Vanilla	-	98.8 93.3	12.1 100.0	98.7 92.4	10.8 100.0	28.6 5.2
	P-CIArC	-	40.1 3.6	100.0 100.0	28.6 0.7	100.0 100.0	17.9 3.8
	R-CIArC	Class	40.1 93.3	100.0 100.0	28.6 92.4	100.0 100.0	18.4 3.8
		Artifact	97.6 92.6	97.0 100.0	97.5 91.6	49.2 100.0	18.0 3.8
	Both	97.6 93.3	97.0 100.0	97.5 92.4	49.2 100.0	18.4 3.8	
VGG16	Vanilla	-	98.8 91.8	27.3 97.0	98.7 91.4	21.4 32.8	22.7 5.2
	P-CIArC	-	59.9 4.6	0.0 100.0	37.5 2.4	0.0 100.0	11.8 4.5
	R-CIArC	Class	98.8 91.8	27.3 97.0	98.7 91.4	21.4 32.8	11.3 4.4
		Artifact	95.2 87.6	3.0 100.0	94.9 85.9	2.9 100.0	9.9 4.4
	Both	98.8 91.8	27.3 97.0	98.7 91.4	21.4 32.8	11.3 4.4	
Efficient Net-B0	Vanilla	-	99.4 86.6	6.1 97.0	99.4 83.1	5.7 32.8	37.0 4.0
	P-CIArC	-	91.6 10.7	3.0 0.0	90.9 1.9	2.9 0.0	24.8 3.5
	R-CIArC	Class	99.4 83.4	6.1 0.0	99.4 74.3	5.7 0.0	24.9 3.5
		Artifact	99.4 82.9	3.0 1.5	99.4 73.9	2.9 1.5	24.8 3.5
	Both	99.4 83.4	6.1 1.5	99.4 74.3	5.7 1.0	24.9 3.5	

Table 3. Model correction results for the background object artifacts inserted in FunnyBirds datasets. We report scores on the (*backdoor* | *shortcuts*) version. The best scores are highlighted in bold.

Model	Method	Condition	Accuracy (clean)	Accuracy (artifact)	F1 (clean)	F1 (artifact)	Artifact relevance
Resnet18	Vanilla	-	- 83.3	50.1 80.4	- 79.4	50.2 33.4	23.9 9.1
	P-CIArC	-	- 61.0	45.1 80.4	- 21.5	14.4 40.2	19.0 8.2
	R-CIArC	Class	- 83.3	59.5 80.4	- 79.1	52.4 32.9	19.1 8.3
		Artifact	- 83.0	46.2 82.1	- 79.0	43.8 41.1	22.1 8.6
	Both	- 83.0	50.3 80.4	- 79.0	50.2 32.9	21.9 8.6	
VGG16	Vanilla	-	- 79.5	20.1 78.6	- 73.2	17.7 32.4	46.3 18.3
	P-CIArC	-	- 63.8	56.8 76.8	- 30.9	24.3 39.3	21.1 15.1
	R-CIArC	Class	- 78.2	47.7 78.6	- 72.4	32.1 32.4	27.2 14.5
		Artifact	- 79.2	27.6 82.1	- 72.4	15.4 41.6	38.6 17.0
	Both	- 79.2	27.4 83.9	- 73.0	20.8 34.4	38.7 17.2	
Efficient Net-B0	Vanilla	-	- 85.3	57.9 83.9	- 82.0	57.1 35.0	22.2 11.2
	P-CIArC	-	- 51.5	51.7 60.7	- 8.5	8.5 18.9	19.4 9.9
	R-CIArC	Class	- 78.6	65.7 58.9	- 71.6	53.0 15.0	18.5 9.5
		Artifact	- 83.5	59.8 60.7	- 79.8	45.1 18.9	18.4 9.7
	Both	- 84.4	63.0 58.9	- 81.2	56.3 15.0	17.7 9.5	

Table 4. Model correction results for the “reflection”, “band aid” and “skin marker” artifacts found in ISIC. We report scores on the (*poisoned* | *original*) version. The best scores are highlighted in bold.

0 for class 1 samples, resulting in less than 50% accuracy for P-CIArC and class-conditional R-CIArC. However, using artifact-conditional and combined approaches helps to reduce the backdoor effect while maintaining good performance on clean samples. Although P-CIArC fails to reduce the backdoor effect for VGG16 and EfficientNet-B0, all variants of R-CIArC allow us to recover accuracy and F1 scores for clean samples. In the *shortcut* setting, sup-

pressing 10 artifacts with P-CIArC notably decreases performance. Fig. 6 further illustrates the dynamics of clean sample accuracy as the number of corrected artifacts increases for the VGG16 model. All reactive CIArCs improve clean sample performance, with the combined approach yielding the best results across all architectures.

Tab. 4 showcases the results for the same architectures, the ISIC2019 dataset and the artifacts “band-aid,” “skin




Figure 6. Clean samples test accuracy by the number of artifacts suppressed with FunnyBirds *shortcut* dataset and VGG16: with *reactive* model correction (R-CIArC), we can retain accuracy significantly better than when with the traditional approach (P-CIArC).

marker,” and “reflection” for both the *poisoned* (left) and *real* (right) evaluation settings. In both configurations, the R-CIArC class conditions link the class “NV” with artifacts “band-aid” and “skin marker”, and the class “BKL” with “skin marker” and “reflection”. In the *poisoned* setting, we observe that R-CIArC consistently outperforms Vanilla models across all architectures, which is not the case with P-CIArC for ResNet18 and EfficientNet-B0. As all test samples are poisoned in this setting, global suppression of the artifact with P-CIArC may prove beneficial, as indicated by the superior performance of P-CIArC compared to R-CIArC in the VGG16 case. In the *real* setting, while P-CIArC significantly harms the performance on clean samples, R-CIArC recovers it to a level close to that of the Vanilla models for all architectures. Fig. 7 additionally depicts LRP heatmaps for the Vanilla models, as well as models corrected with P-CIArC and R-CIArC.

The condition-generating functions may not always be accurate (as demonstrated in Appendix B.5), causing the correction not to get activated when needed. This explains why in some cases P-CIArC demonstrates slightly lower artifact relevance compared to R-CIArC, and in the case of *poisoned* ISIC2019 and VGG16 better model performance.

5. Limitations

While our results show that R-CIArC outperforms P-CIArC in various scenarios, the choice of reactive heuristics is not always straightforward. Furthermore, condition-generating functions may not always be accurate, resulting in the correction not being activated when required.

Additionally, while reactive bias suppression reduces potentially negative (harmful) impact for clean samples (without present artifacts), it does not fully address all the flaws




Figure 7. LRP heatmaps for clean images strongly activating artifact concepts “reflection”, “band-aid,” and “skin marker”: P-CIArC strongly suppresses *all* artifacts, while R-CIArC (combined approach) correctly refrains from suppression.

of P-CIArC. Applying P-CIArC transformations only on artifact samples may still inadvertently alter valid concepts crucial for prediction. Furthermore, merely suppressing a single direction in the latent space may not always suffice to eliminate artifact relevance, as evidenced in Sec. 4.4.

6. Conclusion

In this work, we introduced a general framework for reactive model correction, aimed at mitigating the issue of “collateral damage” to task-relevant features. Specifically, we applied the reactive principle to the P-CIArC post-hoc model correction framework, resulting in its reactive variant, R-CIArC. Our empirical findings demonstrate the detrimental impact of P-CIArC on clean samples while showcasing the superior performance of R-CIArC in this aspect.

Future directions include application of the reactive principle to other post-hoc correction methods, such as classifier editing [37], LEACE[5], and INLP [32]. We also aim to explore more granular reactive conditions, such as those based on prediction strategy or artifact relevance.

Acknowledgements

This work was supported by the Federal Ministry of Education and Research (BMBF) as grant BIFOLD (01IS18025A, 01IS180371I); the European Union’s Horizon 2020 research and innovation programme (EU Horizon 2020) as grant iToBoS (965221); the European Union’s Horizon Europe research and innovation programme (EU Horizon Europe) as grant TEMA (101093003); the German Research Foundation (DFG) as research unit DeSbi (KI-FOR 5363); and the state of Berlin within the innovation support programme ProFIT (IBB) as grant BerDiBa (10174498).

References

- [1] Reduan Achtibat, Maximilian Dreyer, Ilona Eisenbraun, Sebastian Bosse, Thomas Wiegand, Wojciech Samek, and Sebastian Lapuschkin. From attribution maps to human-understandable explanations through Concept Relevance Propagation. *Nature Machine Intelligence*, 5(9):1006–1019, 2023. Number: 9 Publisher: Nature Publishing Group. 3
- [2] Christopher J. Anders, David Neumann, Wojciech Samek, Klaus-Robert Müller, and Sebastian Lapuschkin. Software for dataset-wide XAI: from local explanations to global insights with zennit, corelay, and virelay. *CoRR*, abs/2106.13200, 2021. 3, 4
- [3] Christopher J. Anders, Leander Weber, David Neumann, Wojciech Samek, Klaus-Robert Müller, and Sebastian Lapuschkin. Finding and removing Clever Hans: Using explanation methods to debug and improve deep models. *Information Fusion*, 77:261–295, 2022. 1, 2, 3
- [4] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLOS ONE*, 10(7):1–46, 2015. 3, 6
- [5] Nora Belrose, David Schneider-Joseph, Shauli Ravfogel, Ryan Cotterell, Edward Raff, and Stella Biderman. Leace: Perfect linear concept erasure in closed form, 2023. 2, 3, 8
- [6] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2016. 2
- [7] Giovanni Briganti and Olivier Le Moine. Artificial intelligence in medicine: today and tomorrow. *Frontiers in medicine*, 7:509744, 2020. 1
- [8] Noel C. F. Codella, David Gutman, M. Emre Celebi, Brian Helba, Michael A. Marchetti, Stephen W. Dusza, Aadi Kalloo, Konstantinos Liopyris, Nabin Mishra, Harald Kittler, and Allan Halpern. Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 168–172, 2018. 2, 4, 3
- [9] Marc Combalia, Noel C. F. Codella, Veronica Rotemberg, Brian Helba, Veronica Vilaplana, Ofer Reiter, Cristina Carrera, Alicia Barreiro, Allan C. Halpern, Susana Puig, and Josep Malvehy. Bcn20000: Dermoscopic lesions in the wild, 2019. 2, 4, 3
- [10] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995. 3
- [11] Nicola De Cao, Wilker Aziz, and Ivan Titov. Editing factual knowledge in language models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6491–6506, Online and Punta Cana, Dominican Republic, 2021. Association for Computational Linguistics. 1
- [12] Jesse Dodge, Maarten Sap, Ana Marasovic, William Agnew, Gabriel Ilharco, Dirk Groeneveld, and Matt Gardner. Documenting the english colossal clean crawled corpus. *CoRR*, abs/2104.08758, 2021. 2
- [13] Maximilian Dreyer, Reduan Achtibat, Wojciech Samek, and Sebastian Lapuschkin. Understanding the (Extra-)Ordinary: Validating Deep Model Decisions with Prototypical Concept-based Explanations, 2023. arXiv:2311.16681 [cs]. 3
- [14] Maximilian Dreyer, Frederik Pahde, Christopher J. Anders, Wojciech Samek, and Sebastian Lapuschkin. From hope to safety: Unlearning biases of deep models via gradient penalization in latent space. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(19):21046–21054, 2024. 1, 2
- [15] Thomas Fel, Victor Boutin, Louis Béthune, Remi Cadene, Mazda Moayeri, Léo Andéol, Mathieu Chalvidal, and Thomas Serre. A Holistic Approach to Unifying Automatic Concept Extraction and Concept Importance Estimation. *Advances in Neural Information Processing Systems*, 36, 2023. 3
- [16] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020. Number: 11 Publisher: Nature Publishing Group. 1, 4
- [17] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain, 2019. 4
- [18] 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 (CVPR)*, 2016. 5
- [19] Robin Hesse, Simone Schaub-Meyer, and Stefan Roth. Funnybirds: A synthetic vision dataset for a part-based analysis of explainable ai methods. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3981–3991, 2023. 2, 4
- [20] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, and Rory Sayres. Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV). In *Proceedings of the 35th International Conference on Machine Learning*, pages 2668–2677. PMLR, 2018. ISSN: 2640-3498. 2, 3
- [21] Maximilian Kohlbrenner, Alexander Bauer, Shinichi Nakajima, Alexander Binder, Wojciech Samek, and Sebastian Lapuschkin. Towards best practice in explaining neural network decisions with lrp. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7, 2020. 3, 4
- [22] Abhinav Kumar, Chenhao Tan, and Amit Sharma. Probing classifiers are unreliable for concept removal and detection. In *Advances in Neural Information Processing Systems*, pages 17994–18008. Curran Associates, Inc., 2022. 2
- [23] Sebastian Lapuschkin, Stephan Wäldchen, Alexander Binder, Grégoire Montavon, Wojciech Samek, and Klaus-Robert Müller. Unmasking clever hans predictors and assessing what machines really learn. *Nature Communications*, 10(1):1096, 2019. 1
- [24] Sebastian Lapuschkin, Stephan Wäldchen, Alexander Binder, Grégoire Montavon, Wojciech Samek, and Klaus-

- Robert Müller. Unmasking Clever Hans predictors and assessing what machines really learn. *Nature Communications*, 10(1):1096, 2019. Number: 1 Publisher: Nature Publishing Group. 1
- [25] Zhiheng Li, Ivan Evtimov, Albert Gordo, Caner Hazirbas, Tal Hassner, Cristian Canton Ferrer, Chenliang Xu, and Mark Ibrahim. A whac-a-mole dilemma: Shortcuts come in multiples where mitigating one amplifies others. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20071–20082, 2023. 2
- [26] Weixin Liang, Girmaw Abebe Tadesse, Daniel Ho, Li Fei-Fei, Matei Zaharia, Ce Zhang, and James Zou. Advances, challenges and opportunities in creating data for trustworthy ai. *Nature Machine Intelligence*, 4(8):669–677, 2022. 2
- [27] Lorenz Linhardt, Klaus-Robert Müller, and Grégoire Montavon. Preemptively pruning clever-hans strategies in deep neural networks. *Information Fusion*, 103:102094, 2024. 1
- [28] Masahiro Mitsuhashi, Hiroshi Fukui, Yusuke Sakashita, Takanori Ogata, Tsubasa Hirakawa, Takayoshi Yamashita, and Hironobu Fujiyoshi. Embedding human knowledge into deep neural network via attention map, 2019. 2, 4, 3
- [29] Yannic Neuhaus, Maximilian Augustin, Valentyn Boreiko, and Matthias Hein. Spurious features everywhere—large-scale detection of harmful spurious features in imagenet. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20235–20246, 2023. 1
- [30] Frederik Pahde, Maximilian Dreyer, Wojciech Samek, and Sebastian Lapuschkin. Reveal to revise: An explainable ai life cycle for iterative bias correction of deep models. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*, pages 596–606, Cham, 2023. Springer Nature Switzerland. 1, 4, 3
- [31] Frederik Pahde, Maximilian Dreyer, Leander Weber, Moritz Weckbecker, Christopher J. Anders, Thomas Wiegand, Wojciech Samek, and Sebastian Lapuschkin. Navigating Neural Space: Revisiting Concept Activation Vectors to Overcome Directional Divergence, 2024. arXiv:2202.03482 [cs]. 3, 5, 4
- [32] Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. Null it out: Guarding protected attributes by iterative nullspace projection. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7237–7256, Online, 2020. Association for Computational Linguistics. 1, 2, 3, 8
- [33] Laura Rieger, Chandan Singh, William Murdoch, and Bin Yu. Interpretations are useful: penalizing explanations to align neural networks with prior knowledge. In *International Conference on Machine Learning*, pages 8116–8126. PMLR, 2020. 1, 2
- [34] Andrew Slavin Ross, Michael C Hughes, and Finale Doshi-Velez. Right for the right reasons: training differentiable models by constraining their explanations. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 2662–2670, 2017. 1, 2
- [35] Nusrat Rouf, Majid Bashir Malik, Tasleem Arif, Sparsh Sharma, Saurabh Singh, Satyabrata Aich, and Hee-Cheol Kim. Stock market prediction using machine learning techniques: a decade survey on methodologies, recent developments, and future directions. *Electronics*, 10(21):2717, 2021. 1
- [36] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3): 211–252, 2015. 3
- [37] Shibani Santurkar, Dimitris Tsipras, Mahalaxmi Elango, David Bau, Antonio Torralba, and Aleksander Madry. Editing a classifier by rewriting its prediction rules. In *Advances in Neural Information Processing Systems*, pages 23359–23373. Curran Associates, Inc., 2021. 1, 2, 3, 8
- [38] Patrick Schramowski, Wolfgang Stammer, Stefano Teso, Anna Brugger, Franziska Herbert, Xiaoting Shao, Hans-Georg Luigs, Anne-Katrin Mahlein, and Kristian Kersting. Making deep neural networks right for the right scientific reasons by interacting with their explanations. *Nature Machine Intelligence*, 2(8):476–486, 2020. 1
- [39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations*, 2015. 5
- [40] Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019. 5
- [41] Stefano Teso and Kristian Kersting. Explanatory interactive machine learning. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 239–245, 2019. 2
- [42] Guido Vittorio Travaini, Federico Pacchioni, Silvia Bellumore, Marta Bosia, and Francesco De Micco. Machine learning and criminal justice: A systematic review of advanced methodology for recidivism risk prediction. *International journal of environmental research and public health*, 19(17): 10594, 2022. 1
- [43] Nina Weng, Paraskevas Pegios, Aasa Feragen, Eike Petersen, and Siavash Bigdeli. Fast diffusion-based counterfactuals for shortcut removal and generation. *arXiv preprint arXiv:2312.14223*, 2023. 1, 2
- [44] Shirley Wu, Mert Yuksekogul, Linjun Zhang, and James Zou. Discover and cure: Concept-aware mitigation of spurious correlation. In *International Conference on Machine Learning*, pages 37765–37786. PMLR, 2023. 1, 2