

Towards Robust Classification Model by Counterfactual and Invariant Data Generation

Chun-Hao Chang, George Alexandru Adam, Anna Goldenberg

University of Toronto, Vector Institute, The Hospital for Sick Children

{kingsley.chang, alex.adam}@mail.utoronto.ca, anna.goldenberg@utoronto.ca

Abstract

Despite the success of machine learning applications in science, industry, and society in general, many approaches are known to be non-robust, often relying on spurious correlations to make predictions. Spuriousness occurs when some features correlate with labels but are not causal; relying on such features prevents models from generalizing to unseen environments where such correlations break. In this work, we focus on image classification and propose two data generation processes to reduce spuriousness. Given human annotations of the subset of the features responsible (causal) for the labels (e.g. bounding boxes), we modify this causal set to generate a surrogate image that no longer has the same label (i.e. a counterfactual image). We also alter non-causal features to generate images still recognized as the original labels, which helps to learn a model invariant to these features. In several challenging datasets, our data generations outperform state-of-the-art methods in accuracy when spurious correlations break, and increase the saliency focus on causal features providing better explanations.

1. Introduction

What makes an image be labeled as a cat? What makes a doctor think there is a tumor in a CT scan? What makes a human label a movie review as positive or negative? These questions are inherently causal, but typical machine learning models rely on associations between features and labels rather than causation. Especially in high-dimensional feature spaces with strong correlations, learning which sets of features are right (causal) associations to predict targets becomes difficult, as different sets can result in the same best training accuracy. Because of this, we see issues such as *spurious correlations* [12], *artifacts* [15], *lack of robustness* [2, 5], and *discrimination* [18] happening across many machine learning fields.

Spurious associations happen when factors correlate with labels but are not causal. We might consider factors as spu-

rious associations if *intervening* on such factors would not change the resulting labels. In the context of images, the backgrounds of images can be a source of spurious correlations with labels (e.g. a forest background correlates with bird label) because changing (intervening on) backgrounds should not affect the labels of the foreground classification.

In this paper, we aim to address such spurious associations in the typical ML classification framework by incorporating human causal knowledge. Given a human rationale behind a labeling process (e.g. this part of the image is cat-like), we augment our datasets to break the correlations between backgrounds and labels in two ways. First, we generate counterfactuals that ask "how can we modify the image such that a human would no longer label it as a cat?" That is, by removing the causal features (foreground region containing the cat), and imputing it in a way that is consistent with the background, we generate the counterfactual image that would not be labeled by humans as a cat. Second, we intervene on the non-causal factors (i.e. image backgrounds) to generate new images still containing the cat but with a modified background. This helps the model be invariant to such factors. We experiment on several large-scale datasets and show our methods consistently improve the accuracy and saliency focus on causal features. Our contributions can be summarized as follows:

- We use various counterfactual and invariant data generations to augment training datasets which makes models more robust to spurious correlations.
- We show that our augmentations lead to similar or better accuracy than state-of-the-art saliency regularization and other robustness baselines on challenging datasets in the presence of background shifts. We also find combining our augmentations with saliency regularization can further improve performance.
- Our methods have stronger saliency focus on causal features that provide better explanations, although we find strong saliency on causal features only correlates weakly with good generalization.

2. Related Work

Various works have found that standard machine learning models rely on spurious patterns to make predictions and do not generalize to unseen environments [12]. For instance, Geirhos et al. [11] found standard ImageNet-trained models classify images using object’s texture rather than object’s shape. Several medical imaging classifiers have also been shown to use spurious background to make predictions for COVID-19 [23] and other lung symptoms [38]. Similarly, Young et al. [36] showed deep learning models for CT-scans, although having high accuracy, seem to produce explanations outside of the relevant regions when visualized by Grad-CAM [30] and Shap [20]. Bissoto et al. [6] also found that models trained using public skin lesion datasets tend to have explanations outside of the human-labeled important region, questioning their abilities to generalize across other datasets.

Several methods have been proposed to remove known spurious correlations in concepts (e.g. gender or texture bias). Lu et al. [19], Zmigrod et al. [43] removed gender bias in text by swapping pronouns ("he" becomes "she") to augment the data. Geirhos et al. [11] trained on augmented ImageNet datasets generated in different styles via style transfer to remove texture bias. Zemel et al. [39], Madras et al. [21] directly penalized the model to prevent it from classifying sensitive concepts including race or sex to achieve fairness. In this work we study a different case where the spurious and causal features are separated feature-wise which allows us to remove biases without knowing them in the first place.

Several previous works have attempted to solve the same problem with a different approach: they directly regularize the explanations (saliency) of the model to match the human-labeled important features. Ross et al. [28] were the first to propose regularizing the input gradients toward the causal features and showed improved robustness when the model was evaluated on a different test distribution. Erion et al. [10] used the expected gradient (a stronger saliency method) and regularized it toward other forms of human priors (e.g. sparsity or smoothness). Rieger et al. [26] proposed using Contextual Decomposition [32] which can regularize not just per-pixel saliency but the interactions of the pixels. Several works also found regularizing saliency helps in text classification [9, 13] or medical imaging [42]. In addition, Ross and Doshi-Velez [27] found that input gradient regularization improves adversarial robustness. Bao et al. [3], Mitsuhashi et al. [24] explored regularizing attention (instead of gradient) and showed improvement in text classification. Despite all the aforementioned success, Viviano et al. [33] reported an underwhelming relationship between controlling saliency maps and improving generalization performance in two large-scale medical imaging datasets.

Augmenting counterfactual data to remove spurious correlation has been investigated in the NLP domain [16], but they relied on human efforts to generate counterfactual data.

Several works were also investigated in Visual Question Answering (VQA) fields by augmenting with counterfactual data that changes the answer [1, 8]. Here we instead investigate the effect on the task of image classification, and explore various generation approaches including heuristics and generative models.

3. Proposed Methods

Generating counterfactual data To break the correlation between non-causal features (backgrounds) and labels, we generate counterfactuals that keep the backgrounds in data but remove foregrounds.

Specifically, consider an image \mathbf{x} with U pixels, a label y , and a causal region $\mathbf{r} \subseteq \{0, 1\}^U$ (1 means causal). We define an infilling function ϕ_{cf} that mixes the original image and the infilling value (specific choices are presented later):

$$\phi_{cf}(\mathbf{x}, \mathbf{r}) = (1-\mathbf{r}) \odot \mathbf{x} + \mathbf{r} \odot \hat{\mathbf{x}}, \text{ where } \hat{\mathbf{x}} \sim p_{\text{infill}}(\hat{\mathbf{x}} | \mathbf{x}_{\mathbf{r}=\mathbf{0}}).$$

Then we label such images as "non- y " $\neg y$, which introduces the need for a counterfactual loss function, and we explore 3 different options: (1) Negative log likelihood of $P(\hat{y} \neq y)$ i.e. $-\log(1 - P(\hat{y} = y | x))$; (2) KL divergence between the uniform distribution and the predicted probability i.e. $\text{KL}(\text{Uniform}(y) || P(\hat{y} | x))$. The intuition is that given we remove foregrounds, the model should just predict a uniform distribution among all classes, since it does not have a class for backgrounds; (3) KL divergence between the uniform distribution except original class y , and the predicted probability. We found (2) worked poorly for removing the spurious correlations, and that (1) worked better than (3). We choose (1) as our final objective. Then we augment our training objective with additional counterfactual loss function:

$$L_{cf} = (-\log(1 - P_f(\hat{y} = c | \phi_{cf}(x))))$$

Generating factual data To make a classifier immune to background shifts, we augment our data by perturbing the backgrounds which generates new images with unchanged labels. We define another factual infilling function Φ_f that mixes the foreground with background value $\hat{\mathbf{x}}$:

$$\Phi_f(\mathbf{x}, \mathbf{r}) = \mathbf{r} \odot \mathbf{x} + (1-\mathbf{r}) \odot \hat{\mathbf{x}} \text{ where } \hat{\mathbf{x}} \sim p_{\text{infill}}(\hat{\mathbf{x}} | \mathbf{x}_{\mathbf{r}=\mathbf{1}}).$$

And the final objective (with cross entropy loss) is:

$$L = L_{CE}(y, \hat{y}(x)) + L_{cf} + \overbrace{L_{CE}(y, \hat{y}(\Phi_f(\mathbf{x}, \mathbf{r}))}^{\text{Factual Loss}}$$

Choice of infilling value We describe some methods for producing *counterfactual* infilling values $\hat{\mathbf{x}}$ (some choices

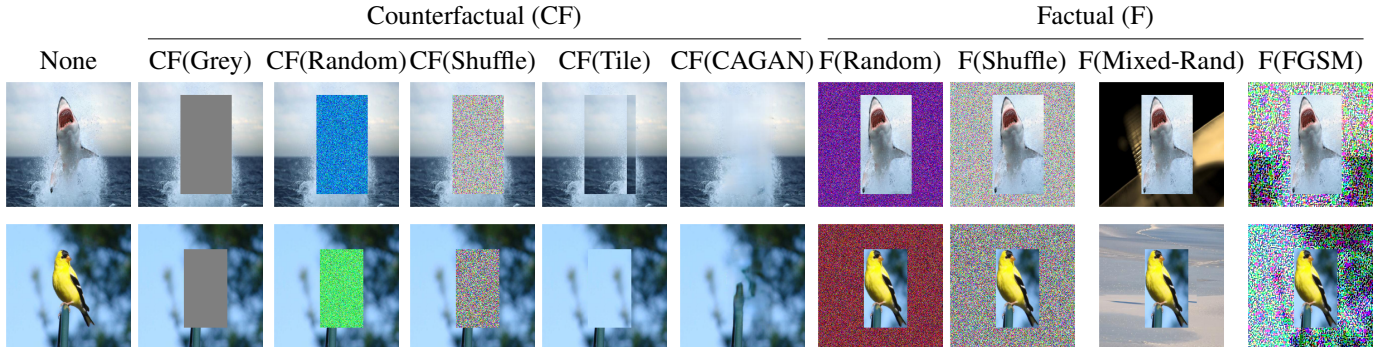


Figure 1: Examples of Counterfactual (CF) or Factual (F) data generation.

are inspired by [7]). **Grey** sets each pixel of \hat{x} to 0.5, which is 0 after being normalized between -1 and 1 . **Random** first samples \hat{x} from a uniform distribution that resembles low-frequency noise, and adds Gaussians of $\sigma = 0.2$ per-channel per-pixel as high-frequency noise, and truncates it between 0 and 1. **Shuffle** randomly shuffles all the pixel values in the specified region; it keeps the marginal distribution but breaks the joint distribution. **Tile** first extracts the largest rectangle from the background that does not intersect with the foreground, and uses that to tile the foreground region. Finally, we use a generative model; **CAGAN** which is the Contextual Attention GAN [37]; we use the authors’ pre-trained ImageNet model to inpaint the removed foreground.

For *factual* infilling, aside from using **Random** and **Shuffle** infilling, we propose **Mixed-Rand** that swaps the background with another randomly-chosen tiled background from images of other classes within the same training batch. We also propose using adversarial attacks to manipulate the non-causal features, i.e. we perform an adversarial attack only on the background region. We adopt the ℓ_{inf} norm and FGSM attack [14] for its fast computation. We tried the PGD attack [22] but found that it performs similarly to FGSM yet is more computationally demanding, so we use FGSM for all experiments. See Figure 1 for examples. We use CF abbreviated for Counterfactual and F for Factual.

Baseline We compare our work with approaches that penalize the model’s saliency (input gradients) outside of bounding boxes. We find the original form of RRR [28] that takes the gradient of sum of the log probabilities across classes do not perform well (also found in Viviano et al. [33]), and thus instead we use GradMask [31] of uncontrast form for multi-class settings that uses the target logit \hat{y} :

$$L_{Sal} = \lambda_{Sal} \sum_{j=1}^U \left(\frac{\partial \hat{y}_{c=y}}{\partial x_j} \right)^2 * (1 - r_j) / \sum_{j=1}^U (1 - r_j)$$

Saliency regularization (**Sal**) explicitly attempt to break

the correlation between non-causal features and labels via model, whereas our augmentations achieve this goal via data.

We also compare with two other methods: **Mix-up** [40] and Label Smoothing (**LS**) [25]. These techniques were not designed to address spurious correlations, but have nonetheless been shown to improve test set accuracy in image classification tasks by non-trivial margins.

Hyperparameters For models, we use a variant of ResNet-50 [17] as our architecture for all our datasets. We call this Original in our experimental section. For data pre-processing, We scales and center-crops images to 224x224 with horizontal flipping and normalizes to $[-1, 1]$. For Sal, we try λ from $\{1e-4, 1e-3, \dots$ to 1000}. For FGSM, we try ϵ of ℓ_{inf} norm for $\{0.15, 0.25, 0.35, 0.5, 0.6, 0.75, 0.9\}$. For Mix-up, we try $\alpha = \{0.1, 0.2, 0.3\}$. For LS, we try $\epsilon = \{0.05, 0.1, 0.2\}$. We run 3 different random seeds. More details are in Appendix A.

4. Experimental Results

We aim to answer the following questions for our augmentations: (i) Do they improve the accuracy under shifted distributions? (ii) Do they make model focus more on foregrounds instead of backgrounds measured by saliency map? (iii) Does focusing on foregrounds indicate better accuracy? (iv) Do our augmentations make models’ predictions less affected by changed backgrounds? We experiment on two controlled datasets that explicitly swaps the foreground and background, and a real-world dataset that has very different backgrounds between train and test images.

4.1. IN-9 dataset: a synthetic perturbed background challenge dataset

ImageNet-9 (IN-9) is a dataset proposed by Xiao et al. [35] to disentangle the relationship between foreground and background. It groups the ImageNet classes into 9 broad classes and filters the images with bounding box annotations, resulting in each class having 5, 045 training images

(a)Original (b)Mixed-Same(c)Mixed-Rand(d)Mixed-Next



Figure 2: Examples of IN9 test sets. Original and Mixed-Same have backgrounds of the same class; Mixed-Rand and Mixed-Next have backgrounds of different classes.

and 450 test images. To disentangle the background and foreground, for each test image they mix the background and foreground in various ways: (1) **Mixed-Same**: the background is swapped with the background of another image belonging to the same class. (2) **Mixed-Rand**: the background is swapped with the background of another image belonging to a different random class. (3) **Mixed-Next**: the background swapped with the background of another image belonging to the next class, i.e. if the class index for the image is 5, then we swap backgrounds with an image from class 6. See Figure 2 for an example. For Sal and F methods, we use the provided foreground segmentation masks as important regions. For CF methods, since the shape of the mask still leaks the information of the foreground objects, we instead use rectangular bounding box. We do not compare to F(Mixed-Rand) since our test sets Mixed-Rand and Mixed-Next are constructed the same way and will give F(Mixed-Rand) an unfair advantage.

We compare our methods in Table 1. For models that rely less on backgrounds, we expect worse performance in Original and Mixed-Same where leveraging backgrounds is beneficial at test time, but expect improvement in Mixed-Rand or Mixed-Next where backgrounds contradict labels. Indeed, Sal and CF methods (except CF(Tile)) perform as expected, with CF(CAGAN) as the best method in Mixed-Rand and Mixed-Next while doing slightly worse in Original and Mixed-Same. To our surprise, F methods like F(Shuffle) perform better in all 4 test sets. We think this is because F methods increase the sample size of images with foregrounds which helps learn more generalizable features. We further combine the best methods in CF, F and Sal and show their combinations improve accuracy even more, suggesting their gains have different causes.

To determine if our augmentations cause models to focus more on foregrounds thereby resulting in higher accuracy, we measure the saliency map of each model and quantify its overlap with provided foreground regions. We try both DeepLiftShap [20] and input gradient as choices of saliency maps and find their ranking is similar, and thus only show results of DeepLiftShap. Specifically, given an image, we take the ℓ_2 norm of the saliency map across chan-

Table 1: The performance results (% of accuracy) of IN9 datasets. We show mean and standard deviation for 3 runs.

	Original	Mixed Same	Mixed Rand	Mixed Next
Orig.	82.0 (0.7)	71.0 (0.4)	51.5 (0.4)	47.5 (0.6)
+ Sal($\lambda = 10^{-4}$)	82.5 (0.6)	71.4 (1.6)	52.4 (2.4)	48.5 (1.4)
+ CF(Grey)	82.4 (0.3)	68.8 (0.9)	51.4 (0.3)	48.8 (0.3)
+ CF(Random)	82.5 (0.2)	69.8 (0.7)	52.1 (0.3)	49.3 (1.1)
+ CF(Shuffle)	82.6 (0.6)	71.5 (0.3)	51.8 (0.3)	48.4 (0.5)
+ CF(Tile)	81.9 (0.8)	58.8 (1.3)	46.5 (0.5)	45.9 (0.1)
+ CF(CAGAN)	80.5 (0.6)	68.7 (0.2)	53.3 (0.2)	50.0 (0.4)
+ F(Random)	83.2 (0.9)	73.9 (1.7)	56.4 (2.8)	53.0 (2.5)
+ F(Shuffle)	83.5 (0.0)	73.8 (0.3)	57.1 (0.8)	53.3 (0.7)
+ F(FGSM $\epsilon = 0.5$)	82.6 (1.3)	70.6 (1.5)	51.6 (2.2)	47.9 (2.5)
+ CF(CAGAN) + F(Shuffle)	81.2 (0.7)	71.5 (1.3)	58.2 (1.7)	54.2 (1.0)
+ CF(CAGAN)+Sal	80.5 (0.2)	68.0 (0.9)	53.2 (0.8)	49.9 (0.9)
+ F(Shuffle) + Sal	83.0 (0.7)	74.5 (0.1)	57.1 (0.9)	52.8 (1.6)
+ CF(CAGAN) + F(Shuffle) + Sal	81.9 (1.2)	72.5 (0.5)	58.0 (1.2)	55.5 (2.1)
+ Mixup($\alpha = 0.2$)	82.8 (1.7)	72.1 (1.0)	52.2 (1.7)	48.2 (1.8)
+ LS($\epsilon = 0.1$)	82.9 (0.4)	71.8 (0.7)	52.5 (0.5)	48.5 (0.7)

Table 2: The model’s AUPR of using saliency map to predict binary foregrounds, and its corresponding accuracy in IN-9 Original and Mixed-Next. High Saliency AUPR does not necessarily indicate high accuracy in both datasets.

	Saliency AUPR		Accuracy	
	Original	Mixed-next	Original	Mixed-next
Orig.	56.7 (0.8)	58.2 (0.5)	82.0 (0.7)	47.5 (0.6)
+ Sal($\lambda = 10^{-4}$)	56.0 (1.6)	58.8 (0.5)	82.5 (0.6)	48.5 (1.4)
+ CF(CAGAN)	57.5 (0.9)	58.7 (0.4)	80.5 (0.6)	50.0 (0.4)
+ F(shuffle)	59.8 (0.9)	61.6 (0.5)	83.5 (0.0)	53.3 (0.7)
+ Combined	60.8 (1.0)	61.2 (0.9)	81.9 (1.2)	55.5 (2.1)
+ Sal($\lambda = 100$)	60.5 (0.7)	61.2 (0.4)	80.8 (0.6)	45.9 (1.2)

nels per pixel as prediction score, and set targets as 1 for pixels in foregrounds and 0 in backgrounds, and compute the binary Area Under the Precision-Recall curve (AUPR). We show saliency AUPR and accuracy in Table 2. Overall F(Shuffle) increases the foreground focus the most compared to Sal and CF(CAGAN), suggesting its large accuracy improvement comes from better focus on the foreground. But, CF(CAGAN) in Mixed-Next with almost the same foreground AUPR still achieves higher accuracy. On the other hand, Sal with high penalty ($\lambda = 100$) which explicitly encourages higher foreground focus has similar AUPR as F(Shuffle) while having much lower accuracy. To further

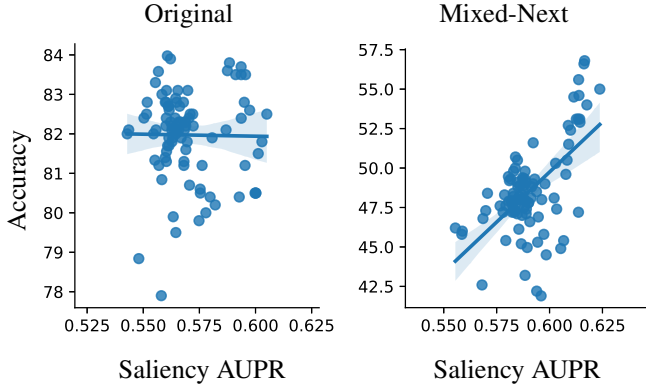


Figure 3: The scatter plot between Saliency map (DeepLift-Shap) AUPR and accuracy for all models we trained in IN-9. Only in Mixed-Next there is a positive correlation ($R^2 = 0.4$) between saliency AUPR and accuracy.

Table 3: The difference of the probability (%) of the next class between Mixed-Next and Original IN9 test sets. The lower the number the model is less affected by backgrounds.

Orig	Sal	CF(CAGAN)	F(shuffle)	Combined
10.1 (0.3)	9.9 (0.4)	7.4 (0.2)	9.2 (0.1)	6.8 (0.3)

understand the relationship, in Figure 3 we plot the Saliency AUPR v.s. Accuracy for all models we have trained. We find no positive correlation in the Original test set which makes sense because getting high accuracy on this partition does not require focusing on foregrounds. In Mixed-Next we do find a stronger positive correlation ($R^2 = 0.4$), although we still have a few outliers in the lower right corner; they are Sal with high λ or F(FGSM) with high ϵ . These results show that accuracy only correlates with foreground AUPR when backgrounds disagree with labels (e.g. Mixed-Next), but does not necessarily correlate well. In fact there may be a tradeoff when strong regularization is used.

To investigate if our models indeed learn to ignore spurious correlations (backgrounds), we measure the difference of probability for the next class between Mixed-Next and Original test sets. Given that these two test sets have identical foregrounds, a model relying more on backgrounds to predict will increase the probability more for the next class when backgrounds are swapped with next-class backgrounds like Mixed-Next does. We show results in Table 3. F(Shuffle), although being the most accurate model, is not the least reliant on backgrounds. Instead CF(CAGAN) is the best from this perspective. This confirms our intuition that F(Shuffle) improves accuracy not only by ignoring backgrounds but also generalizing better in foregrounds (and thus have stronger foreground focus). Both CF and F augmentations outperform Sal, and the combination of the three further decreases the reliance on backgrounds.

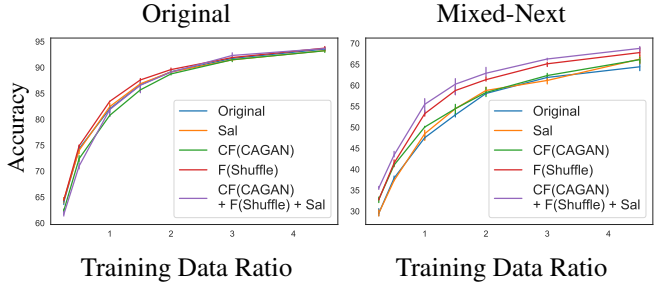


Figure 4: Our methods improve the accuracy under various training data size. Note that we only augment maximum data ratio 1 data since only these images have bounding box annotations, and thus gaps decrease after 1.

To understand how training set size affects our improvements, in Figure 4 we train models using various training data sizes. Here data ratio as 1 means using the original IN-9 dataset around 45k images, and we subset images to make ratio less than 1, or include remaining ImageNet images without bounding boxes to make it bigger than 1. And in Figure 4 we measure the performance gain across different data ratio. When the data ratio is ≤ 1 , our methods continue improving as more data becomes available. When the data ratio is > 1 , the performance gap narrows since the size for our additional data augmentation stays fixed as 1. In summary, our methods improve over baselines across different training sizes, and the more data the better.

In Figure 5 we show examples in the Mixed-Next dataset which the best model (CF(CAGAN)+F(Shuffle)+Sal) predicts correctly and the Original model fails on. We find that their top-5 predictions can be very different.

4.2. Waterbirds dataset

Waterbirds is proposed by Sagawa et al. [29] to study the effect of spurious correlations. It combines bird photographs from the Caltech-UCSD Birds-200-2011 (CUB) dataset [34] with image backgrounds from the Places dataset [41]. They label each bird as one of $Y = \{\text{waterbird, landbird}\}$ and place it on one of $A = \{\text{water background, land background}\}$, with waterbirds (landbirds) more frequently appearing against a water (land) background. In the training set, they place 95% of waterbirds against a water background and the remaining 5% against a land background. Similarly, 95% of all landbirds are placed against a land background with the remaining 5% against water. The number is balanced for the validation set. For test set, we divide test images into Original and Flip, where Original contains images with waterbirds on water background and landbirds on land background, while Flip has opposite bird/backgrounds mix. There are 4,795, 1199, 2897 and 2897 examples in the training, validation, Original Test, and Flip Test set. We follow the original paper to finetune a pretrained ResNet-50 on this dataset.

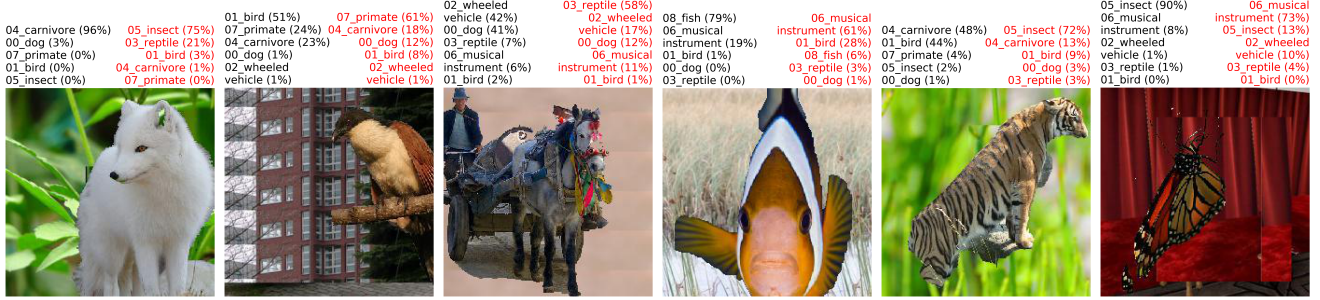


Figure 5: Examples in the IN9 Mixed-Next dataset that the best model (CF(CAGAN)+F(Shuffle)+Sal) predicts correctly while Original model fails. Grey texts are the predictions from the best model, and Red text are from Original model.

Table 4: The accuracy on the Waterbird. The Improv(%) is the relative improvement of Flip Test w.r.t. the Orig. method. We show mean and standard deviation for 3 runs.

	Orig. Test	Flip Test	Improv(%)
Orig.	98.4 (0.1)	76.1 (0.4)	0.0
+ Sal($\lambda = 10$)	98.4 (0.1)	77.6 (0.9)	1.9
+ CF(Grey)	96.0 (1.8)	81.9 (4.9)	7.6
+ CF(Random)	98.1 (0.1)	73.2 (1.2)	-3.8
+ CF(Shuffle)	98.1 (0.1)	71.7 (0.6)	-5.8
+ CF(Tile)	96.1 (1.1)	81.6 (3.6)	7.2
+ CF(CAGAN)	95.2 (0.1)	85.6 (0.3)	12.5
+ F(Random)	98.4 (0.2)	80.7 (0.2)	6.0
+ F(Shuffle)	98.4 (0.0)	79.6 (0.5)	4.5
+ F(FGSM)	98.2 (0.1)	69.3 (1.5)	-8.9
+ CF(CAGAN)	94.8 (0.4)	89.6 (0.7)	17.7
+ F(Random)	94.8 (0.1)	90.6 (0.0)	19.0
+ Mixup($\alpha = 0.1$)	98.4 (0.0)	76.7 (1.5)	0.8
+ LS($\epsilon = 0.1$)	98.4 (0.0)	78.7 (0.5)	3.4

We show the accuracy of each method in Table 4. Similar to what we found in IN-9, Sal maintains the accuracy in Original Test set and improves slightly in Flip Test. CF methods decrease accuracy slightly in the Original Test set while improving accuracy in the Flip Test set, with our most natural augmentations CF(CAGAN) improving the most. F methods remain similarly accurate in Original Test while improving Flip accuracy quite a bit with F(Random) as the best. Further combining the best methods - CF(CAGAN), F(Random) and Sal - improves our Flip accuracy even more up to 19% relative improvement compared to Original method.

To investigate the relationship between foregrounds focus and accuracy (details in Section 4.1), we again show saliency

Table 5: The AUPR of the saliency map and the accuracy on Waterbirds Original and Flip test sets.

	Saliency AUPR		Accuracy	
	Original	Flip	Original	Flip
Orig.	56.2 (0.4)	48.8 (1.0)	98.4 (0.1)	76.1 (0.4)
+ Sal($\lambda = 10$)	60.3 (0.7)	55.0 (0.7)	98.4 (0.1)	77.6 (0.9)
+ CF(CAGAN)	56.7 (0.8)	48.4 (1.0)	95.2 (0.1)	85.6 (0.3)
+ F(Random)	58.3 (0.3)	50.6 (0.1)	98.4 (0.2)	80.7 (0.2)
+ Combined	58.2 (0.6)	50.2 (1.3)	94.8 (0.1)	90.6 (0.0)
+ Sal($\lambda = 10^3$)	62.0 (0.4)	57.7 (0.9)	97.6 (0.1)	72.4 (1.1)




Figure 6: The scatter plot between Saliency map (DeepLift-Shap) AUPR and accuracy for all models we trained in Waterbirds. In Flip test set there is a stronger positive correlation ($R^2 = 0.08$) between saliency AUPR and accuracy.

AUPR and accuracy in Table 5, and scatter plot all models in both Original and Flip test sets in Figure 6. In Table 5, all methods except CF(CAGAN) improve saliency AUPR (left column) in both Original and Flip test set with Sal($\lambda = 10^3$) as the best method. However, the improvement in saliency AUPR does not come with better accuracy in both Original and Flip. For example, Sal($\lambda = 10^3$) gets the best saliency




Figure 7: Examples in the Waterbird Flip Test set that the best model (CF(CAGAN)+F(Random)+Sal) predicts correctly while Original model fails. Grey texts are the predictions from the best model, and Red text are from Original model.

Table 6: The performance results (% of AUC) of CCT datasets. The Improv(%). is the relative improvement of Trans-Test w.r.t. the Orig. method. We show mean and standard deviation for 3 runs.

	Cis-Test	Trans-Test	Improv(%)
Orig.	82.1 (1.9)	70.2 (1.2)	0.0
+ Sal($\lambda = 100$)	86.6 (1.2)	72.2 (1.1)	2.8
+ CF(Grey)	83.8 (2.1)	69.7 (2.5)	-0.7
+ CF(Random)	85.7 (1.8)	70.8 (2.7)	0.9
+ CF(Shuffle)	84.2 (0.2)	68.9 (1.5)	-1.9
+ CF(Tile)	86.2 (1.6)	70.4 (2.8)	0.3
+ CF(CAGAN)	84.4 (2.3)	70.1 (2.1)	-0.1
+ F(Random)	82.9 (1.1)	70.9 (1.6)	1.0
+ F(Shuffle)	84.2 (2.4)	70.8 (1.9)	0.9
+ F(Mixed-Rand)	82.7 (1.7)	70.8 (0.9)	0.9
+ F(FGSM $\epsilon = 0.5$)	84.9 (2.1)	71.9 (1.8)	2.4
+ CF(Tile) + Sal	86.9 (0.9)	74.0 (1.0)	5.4
+ CF(Tile) + F(Shuffle)	85.0 (1.3)	72.3 (2.1)	3.0
+ F(Shuffle) + Sal	88.5 (0.5)	73.2 (2.3)	4.3
+ CF(Tile)	87.0 (1.4)	73.2 (1.1)	4.3
+ F(Shuffle) + Sal	87.0 (1.4)	73.2 (1.1)	4.3
+ Mixup($\alpha = 0.2$)	85.0 (0.5)	69.4 (1.7)	-1.1
+ LS ($\epsilon = 0.2$)	81.2 (0.7)	67.8 (3.0)	-3.4

AUPR while having the worst accuracy in Flip. In Figure 6, we also find that in Flip there is a slightly better correlation between saliency focus and accuracy than in Original ($R^2 = 0.08 > 0.004$), but it is also not very strong.

In Figure 7, we show some example images that our best model predicts correctly while the Original model fails.

4.3. Caltech Camera Trap (CCT) dataset

We test on a real-world dataset faced with background shifts across training and test sets. Camera traps are motion- or heat-triggered cameras placed in locations of interest by biologists to monitor and study animal populations and be-

Table 7: The AUPR of the saliency map (DeepLiftShap) of the ground truth class on CCT Cis and Trans datasets. The higher the model has better foreground focus.

	Saliency AUPR		AUC	
	Cis-Test	Trans-Test	Cis-Test	Trans-Test
Orig	5.7 (0.1)	7.2 (0.0)	82.1 (1.9)	70.2 (1.2)
Sal($\lambda = 100$)	6.4 (0.5)	7.5 (0.3)	86.6 (1.1)	72.2 (1.1)
CF(Tile)	6.2 (0.5)	7.8 (0.5)	86.2 (1.6)	70.4 (2.8)
F(Shuffle)	7.3 (1.1)	9.3 (1.8)	84.2 (2.4)	70.8 (1.9)
CF(Tile) + Sal	7.9 (0.8)	8.8 (0.7)	86.9 (0.9)	74.0 (1.0)
F(Shuffle) + Sal	8.1 (0.2)	9.7 (0.5)	88.5 (0.5)	73.2 (2.3)
Sal($\lambda = 10^4$)	6.5 (0.8)	7.6 (0.5)	78.8 (15.8)	69.4 (7.8)

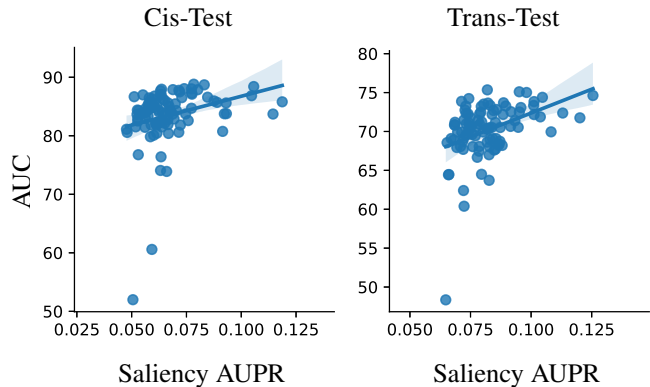


Figure 8: The scatter plot between Saliency map (DeepLiftShap) AUPR and accuracy for all models we trained in CCT. In Trans-Test there is a stronger positive correlation ($R^2 = 0.16$) than Cis-Test ($R^2 = 0.07$) between saliency AUPR and accuracy.

havior. The goal is to train a classifier that recognizes the same species of animals but with different camera backgrounds. Caltech Camera Traps-20 (CCT) dataset [4], consists of 57, 868 images across 20 locations in the American Southwest, each labeled with one of 15 classes of animals. We follow the setup of the original paper that divides test

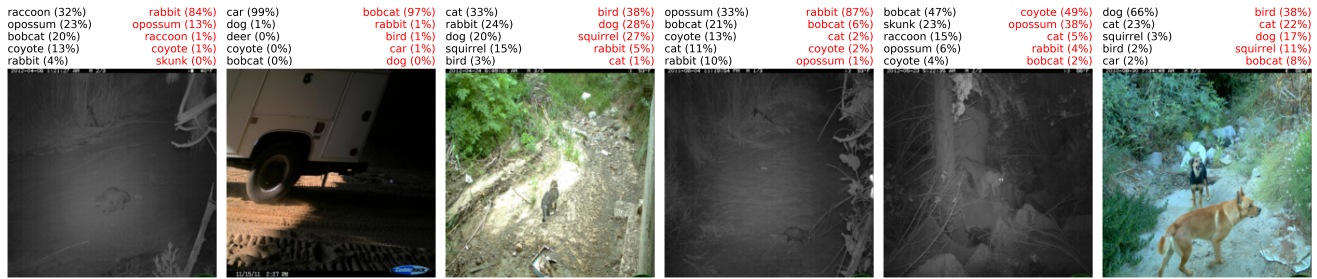


Figure 9: Examples of the CCT dataset that the best model (+CF(Tile)+Sal) predicts correctly but Original model fails. Grey texts are the predictions from the best model, and Red text are from Original model.

images into "cis-locations" and "trans-locations", where "cis-locations" are images with locations seen during training, and "trans" are new locations not seen before. This gives us 13, 553 training images, 3, 484 validation and 15, 827 test images from cis-locations, and 23, 275 from trans-locations. Since cis and trans locations have imbalanced classes, we use multi-class AUC to measure the performance.

In Table 6, we find $\text{Sal}(\lambda = 100)$ performs better than CF and F augmentations alone in both Cis and Trans split, although our combined method CF(Tile)+F(Shuffle) performs similarly with Sal in Trans-Test (72.3 v.s. 72.2). For CF augmentations, they mainly improve on Cis rather than Trans; this shows that even within the same camera location like Cis dataset, there still exists spurious correlations probably due to wide variety of backgrounds (lightning, occlusions), relatively small bounding boxes, and small amount of training data. Note that the pretrained generative model CAGAN is not fine-tuned on the CCT dataset, and further fine-tuning on this dataset can potentially give better inpaintings that may further improve the performance. For F augmentations, F(FGSM) helps the most in the Trans split. Surprisingly, although F(FGSM) works best alone, when combining with CF(Tile) or Sal it gets worse performance; we think the adversarial nature of F(FGSM) might make the optimization harder. Instead we try combinations with F(Shuffle). We find both CF(Tile) + Sal and F(Shuffle) + Sal can further improve the performance in both Cis and Trans splits, again suggesting that CF, F and Sal improve in different ways.

We again analyze the relationship between saliency AUPR and accuracy (details in Section 4.1). In Table 7, Sal, CF(Tile) and F(Shuffle) focus better on foregrounds with F(Shuffle) as the best. The combination F(Shuffle)+Sal further improves focus. $\text{Sal}(\lambda = 1e4)$ with too strong regularization again increases saliency focus at the expense of lower accuracy. In Figure 8, we scatter plot the relationship between test set AUC and foreground saliency AUPR. They do not necessarily correlate well, and Trans-Test has better correlations than Cis-Test.

In Figure 9, we show examples where the best model (CF(Tile)+Sal) succeeds but Original model fails in the Trans

split. Although it is unclear if the changed camera locations make the Original model wrongly classify, it can have very different top predictions from our best model.

5. Discussions

In this paper we focus on a particular type of spuriousness where the spurious and causal features are separated feature-wise which enables us to remove spuriousness using foreground annotations alone without knowing what the spurious factors are. In the case where spuriousness happens in both foreground and background such as color or texture bias, our data augmentations would still work if we know what the spurious factors are, and generate corresponding Factual and Counterfactual data. For example, if color blue is correlated with label, we can generate blue images without shape as Counterfactual images and generate random color images as Factual images.

We find our augmentations may vary performance from dataset to dataset, e.g. CF(Tile) does worse than no augmentation in IN-9, yet is one of the better performing methods for CCT. Overall we find CF methods perform best when imputation is natural as evidenced by better performance of CF(CAGAN) in IN-9 and Waterbirds. And we recommend F(Shuffle) or F(Random) for their better performance in our experiments. The best approach is to try all different imputations to see what works well.

Lastly, we recognize that requiring additional annotations such as bounding boxes or segmentation maps can be costly for some datasets. This limitation can be overcome by using pre-trained segmentation models or heatmaps of pretrained models to obtain reasonable annotations. If some classes are novel, then few-shot semantic (FSS) segmentation can be used instead such that only a few images of novel class require manual segmentations, and the rest can be handled by the FSS model.

We believe that developing methods which make models more robust to spurious correlation is essential to overcoming the inherent obstacles to generalization posed by ambiguities in real-world datasets.

References

- [1] Vedika Agarwal, Rakshith Shetty, and Mario Fritz. Towards causal vqa: Revealing and reducing spurious correlations by invariant and covariant semantic editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9690–9698, 2020.
- [2] Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. Analyzing the behavior of visual question answering models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1955–1960, 2016.
- [3] Yujia Bao, Shiyu Chang, Mo Yu, and Regina Barzilay. Deriving machine attention from human rationales. In *EMNLP*, 2018.
- [4] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 456–473, 2018.
- [5] Yonatan Belinkov and Yonatan Bisk. Synthetic and natural noise both break neural machine translation. In *International Conference on Learning Representations*, 2018.
- [6] Alceu Bissoto, Michel Fornaciali, Eduardo Valle, and Sandra Avila. (de) constructing bias on skin lesion datasets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [7] Chun-Hao Chang, Elliot Creager, Anna Goldenberg, and David Duvenaud. Explaining image classifiers by counterfactual generation. In *International Conference on Learning Representations*, 2018.
- [8] Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. Counterfactual samples synthesizing for robust visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10800–10809, 2020.
- [9] Mengnan Du, Ninghao Liu, Fan Yang, and Xia Hu. Learning credible deep neural networks with rationale regularization. In *2019 IEEE International Conference on Data Mining (ICDM)*, pages 150–159. IEEE, 2019.
- [10] Gabriel Erion, Joseph D Janizek, Pascal Sturmfels, Scott Lundberg, and Su-In Lee. Learning explainable models using attribution priors. *arXiv preprint arXiv:1906.10670*, 2019.
- [11] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, 2018.
- [12] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *arXiv preprint arXiv:2004.07780*, 2020.
- [13] Reza Ghaeini, Xiaoli Fern, Hamed Shahbazi, and Prasad Tadepalli. Saliency learning: Teaching the model where to pay attention. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4016–4025, 2019.
- [14] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [15] Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R Bowman, and Noah A Smith. Annotation artifacts in natural language inference data. *arXiv preprint arXiv:1803.02324*, 2018.
- [16] Divyansh Kaushik, Eduard Hovy, and Zachary Lipton. Learning the difference that makes a difference with counterfactually-augmented data. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SkIgs0NFvvr>.
- [17] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Large scale learning of general visual representations for transfer. *arXiv preprint arXiv:1912.11370*, 2019.
- [18] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. In *Advances in neural information processing systems*, pages 4066–4076, 2017.
- [19] Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. Gender bias in neural natural language processing. In *Logic, Language, and Security*, pages 189–202. Springer, 2020.
- [20] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in neural information processing systems*, pages 4765–4774, 2017.
- [21] David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. Learning adversarially fair and transferable representations. In *International Conference on Machine Learning*, pages 3384–3393, 2018.
- [22] Aleksander Madry, Aleksandar Makelev, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.
- [23] Gianluca Maguolo and Loris Nanni. A critic evaluation of methods for covid-19 automatic detection from x-ray images. *arXiv preprint arXiv:2004.12823*, 2020.
- [24] Masahiro Mitsuhara, H. Fukui, Yusuke Sakashita, T. Ogata, Tsubasa Hirakawa, T. Yamashita, and H. Fujiyoshi. Embedding human knowledge in deep neural network via attention map. *ArXiv*, abs/1905.03540, 2019.
- [25] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? In *Advances in Neural Information Processing Systems*, pages 4694–4703, 2019.

- [26] 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.
- [27] Andrew Ross and Finale Doshi-Velez. Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [28] 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.
- [29] Shiori Sagawa, Pang Wei Koh, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=ryxGuJrFvS>.
- [30] 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.
- [31] Becks Simpson, Francis Dutil, Yoshua Bengio, and Joseph Paul Cohen. Gradmask: Reduce overfitting by regularizing saliency. In *International Conference on Medical Imaging with Deep Learning—Extended Abstract Track*, 2019.
- [32] Chandan Singh, W James Murdoch, and Bin Yu. Hierarchical interpretations for neural network predictions. In *International Conference on Learning Representations*, 2018.
- [33] Joseph D Viviano, Becks Simpson, Francis Dutil, Yoshua Bengio, and Joseph Paul Cohen. Underwhelming generalization improvements from controlling feature attribution. *arXiv preprint arXiv:1910.00199*, 2019.
- [34] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical report, 2011.
- [35] Kai Xiao, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. Noise or signal: The role of image backgrounds in object recognition. *arXiv preprint arXiv:2006.09994*, 2020.
- [36] Kyle Young, Gareth Booth, Becks Simpson, Reuben Dutton, and Sally Shrapnel. Deep neural network or dermatologist? In *Interpretability of Machine Intelligence in Medical Image Computing and Multimodal Learning for Clinical Decision Support*, pages 48–55. Springer, 2019.
- [37] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5505–5514, 2018.
- [38] John R Zech, Marcus A Badgeley, Manway Liu, Anthony B Costa, Joseph J Titano, and Eric K Oermann. Confounding variables can degrade generalization performance of radiological deep learning models. *arXiv preprint arXiv:1807.00431*, 2018.
- [39] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *International Conference on Machine Learning*, pages 325–333, 2013.
- [40] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018.
- [41] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. 2014.
- [42] Jiaxin Zhuang, Jiabin Cai, Ruixuan Wang, Jianguo Zhang, and Weishi Zheng. Care: Class attention to regions of lesion for classification on imbalanced data. In *International Conference on Medical Imaging with Deep Learning*, pages 588–597, 2019.
- [43] Ran Zmigrod, Sebastian J Mielke, Hanna Wallach, and Ryan Cotterell. Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1651–1661, 2019.