

Explainable Deep Classification Models for Domain Generalization

Andrea Zunino^{*,†,1}, Sarah Adel Bargal^{*,2}, Riccardo Volpi^{†,3}, Mehrnoosh Sameki⁴,
Jianming Zhang⁵, Stan Sclaroff², Vittorio Murino^{1,6,7}, Kate Saenko²

¹Huawei Ireland Research Center ²Department of Computer Science, Boston University

³Naver Labs Europe ⁴Microsoft ⁵Adobe Research

⁶Pattern Analysis & Computer Vision, Istituto Italiano di Tecnologia ⁷Università di Verona

{andrea.zunino, riccardo.volpi, vittorio.murino}@iit.it,

{sbargal, sclaroff, saenko}@bu.edu, mehrnoosh.sameki@microsoft.com, jianmzha@adobe.com

Abstract

Conventionally, AI models are thought to trade off explainability for lower accuracy. We develop a training strategy that not only leads to a more explainable AI system for object classification, but as a consequence, suffers no perceptible accuracy degradation. Explanations are defined as regions of visual evidence upon which a deep classification network makes a decision. This is represented in the form of a saliency map conveying how much each pixel contributed to the network's decision. Our training strategy enforces a periodic saliency-based feedback to encourage the model to focus on the image regions that directly correspond to the ground-truth object. We quantify explainability using an automated metric, and using human judgement. We propose explainability as a means for bridging the visual-semantic gap between different domains where model explanations are used as a means of disentangling domain specific information from otherwise relevant features. We demonstrate that this leads to improved generalization to new domains without hindering performance on the original domain.

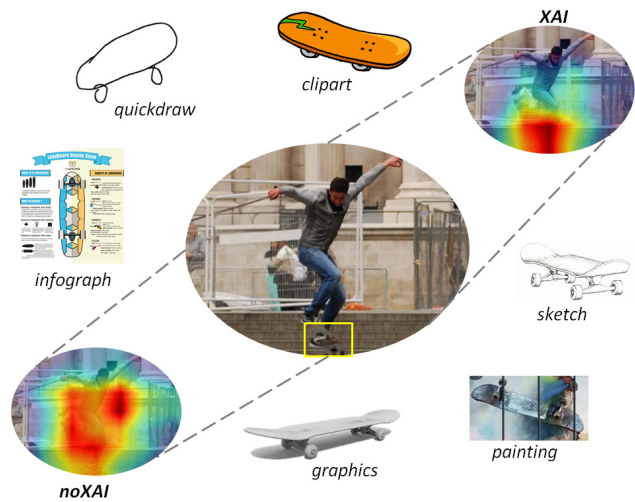


Figure 1. In this figure we demonstrate how explainability (XAI) can be used to achieve domain generalization from a *single source*. Training a deep neural network model to enforce explainability, e.g. focusing on the skateboard region (red is most salient, and blue is least salient) for the ground-truth class skateboard in the central training image, enables improved generalization to other domains where the background is not necessarily class-informative.

1. Introduction

Increased explainability in machine learning is traditionally associated with lower performance, e.g. a decision tree is more explainable, but less accurate than a deep neural network. In this paper we argue that, in fact, increasing the explainability of a deep classifier can improve its generalization, especially to novel domains. End-to-end deep models often exploit biases unique to their training dataset which leads to poor generalization on novel datasets or en-

vironments. We develop a training strategy for deep neural network models that increases explainability, suffers no perceptible accuracy degradation on the training domain, and improves performance on unseen domains.

Domain adaptation and generalization are formulations that mitigate the problem of dataset bias. In domain adaptation one needs to know a priori the target distribution, which limits applicability [12, 4, 25]. In standard domain generalization techniques, one needs several source domains for training, which may not be available in practice. A more generic formulation is *single-source* domain generalization,

^{*}Equal contribution.

[†]This work was done while A. Zunino and R. Volpi were in PAVIS, at Istituto Italiano di Tecnologia.

where one would like to avoid learning dataset bias for better generalization, but only has access to a *single* source distribution. In this work, we address this challenging *single-source* setting. Data augmentation approaches were shown to be successful for improving generalization to unseen domains by randomizing and perturbing the way training images are portrayed, and therefore learning invariance against some inherent biases of the source dataset [33, 37, 36].

A limitation of such data augmentation techniques is that some biases do not depend on color features, but are more structured or context dependent. Consider the case where every sample associated with a given class always has the same background, *e.g.* an object recognition dataset where a soccer ball is mostly seen on a soccer field. In this case, the classifier might learn the background features instead of the ball characteristics -if they are sufficient to obtain good accuracy on training. While context and context-correlations aid classification when the source and target domains come from similar distributions [21, 34, 7], this becomes a limitation when the correlation is corrupted in unseen target domains. The drawback is evident: a model that learns to recognize soccer balls by evaluating whether grass is present or not in a scene will poorly generalize to scenarios where grass is not present or visible. This vulnerability, commonly referred to as “domain bias”, significantly limits the applicability of machine learning systems into the wild.

We posit that the design of algorithms that better mimic the way humans reason, or “explain”, can help mitigating the domain bias issue. Our approach utilizes explainability as a means for bridging the visual-semantic gap between different domains as presented in Figure 1. Specifically, our training strategy is guided by model explanations and available human-labeled explanations. Explanations are defined as regions of visual evidence upon which a network makes a decision. This is represented in the form of a saliency map conveying how much each pixel contributed to the network’s decision.

Our training strategy periodically guides the forward activations of spatial layer(s) of a Convolutional Neural Network (CNN) trained for object classification. The activations are guided to focus on regions in the image that directly correspond to the ground-truth (GT) class label, as opposed to context that may more likely be domain dependent. The proposed strategy aims to reinforce explanations that are non-domain specific, and alleviate explanations that are domain specific. Classification models are compact and fast in comparison to more complex semantic segmentation models. Our approach allows the compact classification model to possess some properties of a segmentation model without increasing model complexity or test-time overhead.

We show how the identification of evidence within a visual input using top-down neural attention formulations [26] can be a powerful tool for domain analysis. Inspired

by these findings, we demonstrate that more explainable deep classification models could be trained without hindering their performance. We then conduct a human study to confirm our intuitive quantification of an “explainable” model. Finally, we demonstrate how the explainable model better generalizes on six unseen target domains, although it was trained only using a single-source domain.

In summary, our contributions are:

- We propose a training strategy that leads to more explainable deep classification models. We quantify explainability computationally and using human judgement.
- We demonstrate benefits of having a more explainable model for single-source domain generalization.

2. Related Work

Robustness to Domain Shift. Several problem formulations have been proposed with the aim of learning models which are more robust in out-of-distribution settings. One formulation that received a large amount of interest from the community is domain adaptation [12, 4, 25]. The assumption here is to have access to a set of samples from a “source” domain for which annotations are available, and a set of samples from a “target” distribution on which we desire to perform well, for which annotation is unavailable or only partially available. There is a significant body of works that propose effective solutions to this problem (*e.g.*, [12, 4, 25, 9, 35, 32]); the limitation is that one needs to know the target distribution a priori.

In domain generalization [20, 13, 19, 14, 28, 15, 46] this assumption is relaxed; the goal here is to learn models that better generalize to unseen domains, without fixing target distributions a priori. As it was originally conceived, domain generalization requires access to several source domains to learn models that better generalize. Recently, different works have proposed ways to learn more general representations by relying on a single-source distribution. Volpi *et al.* [37] propose to rely on adversarial robustness [10], generating samples that are hard for the model over iterations. A related method [36] proposes to find new data augmentation rules over iterations by evaluating the image transformations that the current model is more vulnerable to. Carlucci *et al.* [6] rely on self-supervised learning to learn representations that are less biased towards the source distribution. Domain randomization [33] allows improving performance on unseen, real data when training on rendered data. Hendricks *et al.* [11] enforce looking at a person as opposed to looking at other background elements in an image to make gender prediction less biased.

In this work, we show that our proposed training procedure improves domain generalization performance without assumptions on the target distributions, and without the

need of multiple source distributions. With respect to models trained with data augmentation strategies [36, 37] for domain generalization, the method proposed here allows overcoming more complex dataset biases. We show that our learning procedure is complementary to data augmentation, and can be efficiently used in tandem.

Saliency for Explainability. The “black-box” nature of end-to-end deep neural networks creates highly non-linear and inexplicable feature representations that make it difficult to understand what causes the models to make certain decisions -evidence of a model prediction. Various methods have been introduced that investigate this major drawback of such powerful models. For visual data, interpretability/explainability has been addressed in the form of saliency maps highlighting image regions that a model uses to make a prediction, *i.e.* evidence. Explanation of visual models has been addressed using white-box (black-box) methods where the model parameters are known (unknown).

White-box methods include [40, 1, 39, 29, 31, 43, 26]. Zeiler *et al.* [39] use a variant of the standard backpropagation error from neuron activations in higher layers down to the image level. Selvaraju *et al.* [26] obtain activation maps of a specific class using a weighted sum of deep convolutional features. In [40] top-down attention of a convolutional neural network (CNN) classifier is modelled for generating task-specific attention maps. This work was then extended in the temporal dimension in [1] for Recurrent Neural Networks (RNNs) to provide visual explanations of spatiotemporal models. Black-box methods include [8, 24] where image regions are perturbed and network output is monitored to determine regions of discriminative evidence.

Several works have employed explainability in developing training time and testing time frameworks to further improve model predictions. Cao *et al.* [5] use explanation maps to feed regions of highest importance into the same model and use the predicted class-conditional probabilities to improve the original ones corresponding to the whole image. Zunino *et al.* [45] propose a guided dropout regularizer for deep networks based on the explanation of a network prediction defined as the firing of neurons in specific paths. The explanation at each neuron is utilized to determine the probability of dropout, rather than dropping out neurons uniformly at random as in standard dropout. Bargal *et al.* [2, 3] employ explainability by making sure the model has “the right reasons” for a prediction, defined as reasons that are coherent with those used to make similar correct decisions at training time. Selvaraju *et al.* [27] optimize the alignment between human attention maps and gradient-based network importance for improving performance on Visual Question Answering and Image Captioning tasks. Explainability has also been used in spatial semantic segmentation tasks [16, 44, 38], and object localization tasks [42]. In contrast to previous works, our approach proposes a

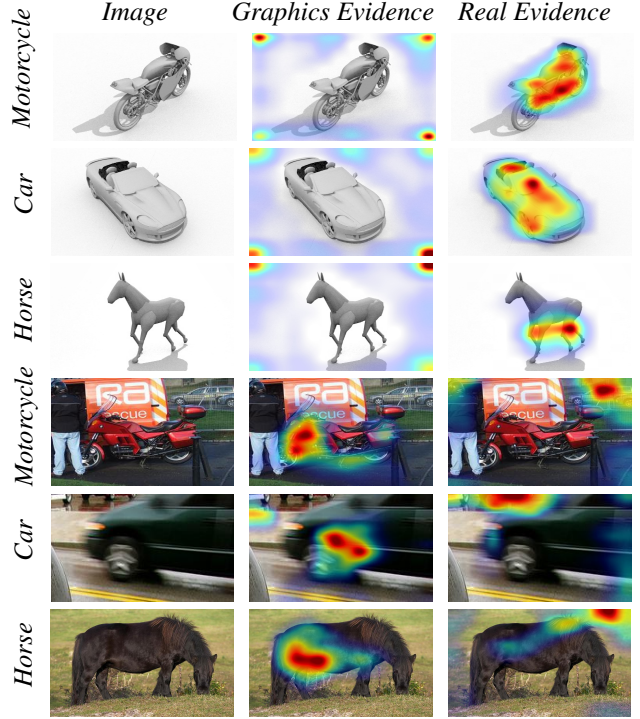


Figure 2. **Domain Evidence.** (Left) Graphics and Real images from the classes *motorcycle*, *car* and *horse*. (Middle) Evidence the domain discrimination network identifies for the Graphics domain. (Right) Evidence the domain discrimination network identifies for the Real domain. For Graphics images, the model selects the white background as evidence for the Graphics domain and selects the object as evidence for the Real domain. For the Real images, the model selects the objects to be evidence for the Graphics domain and selects the background as evidence for the Real domain.

training strategy to ensure that models periodically learn to rely on object-related visual concepts for the task of object classification in order to achieve more explainable models that generalize better on unseen target domains.

3. Explainability for Domain Analysis

We start by motivating how saliency can be used to highlight discriminative evidence found in each domain in a domain transfer setting. We then demonstrate how saliency can be used to highlight how different training strategies shift the model focus within the image. In this section, we use the Syn2Real dataset [23], which is constructed from a graphics source domain rendered from 3D CAD models and a real images target domain of the following classes: Airplane, Bicycle, Bus, Car, Horse, Knife, Motorcycle, Person, Plant, Skateboard, Train, and Truck.

Highlighting Domain Evidence. We set up an experiment to visualize image regions that are domain specific. We train a VGG16 [30] network to differentiate between the

graphics and real images domains of the Syn2Real dataset in a binary classification problem setting. Having a classifier trained to differentiate between domains, we can then visualize why the model processes an unseen image as belonging to a specific domain and not the other. As saliency methods can visualize evidence of classes that are not necessarily GT, we visualize the evidence for each domain in images of the source domain and images of the target domain in Figure 2. For graphics images, the model uses the white background as evidence for the graphics domain, and the object as evidence for the real images domain. For real images, the model uses the object as evidence for the graphics domain, and the busy background as evidence for the real images domain. The evidence associated with the GT domain of an image is observed to be context-dependent, and the evidence associated with an alternative domain is observed to be object-dependent. This capability of interpreting models and visually analyzing differences between domains suggests the possibility of building models that bridge exactly that highlighted domain gap.

Highlighting the Evidence Shift. We now set up an experiment to visualize the shift of focus in input images when different training strategies are used. The first training strategy is vanilla CNN training with no domain adaptation for object classification, *i.e.* training on graphics images only and testing on real images only. The second training strategy employs domain adaptation. We train our model on the Syn2Real classification task without Domain Adaptation. We then repeat training with the domain adaptation approach of Long *et al.* [18]. This approach aligns distributions of the source and target domains based on a joint maximum mean discrepancy. We highlight the shift of the object’s class evidence in test images of the target domain that are misclassified by the model that does not perform domain adaptation, and are correctly classified by the domain adaptation model. Examples are presented in Figure 3. The saliency maps demonstrate the shift toward more discriminative evidence when the domain adaptation model is used. For example, before domain adaptation a metallic surface of the airplane was the evidence the network used to incorrectly classify the airplane as a motorcycle. However, after domain adaptation the image is correctly classified as an airplane and the evidence has shifted to the wings of the airplane - a more discriminative feature of airplane.

Testing on data from domains not used in training poses several challenges that are addressed in the fields of domain adaptation and generalization. We find that different training strategies, *e.g.* with or without domain adaptation, make models reason based on different evidence. This inspired us to encourage models to focus on object-dependent evidence, rendering models that (1) are more explainable, and (2) improve generalization on unseen domains.

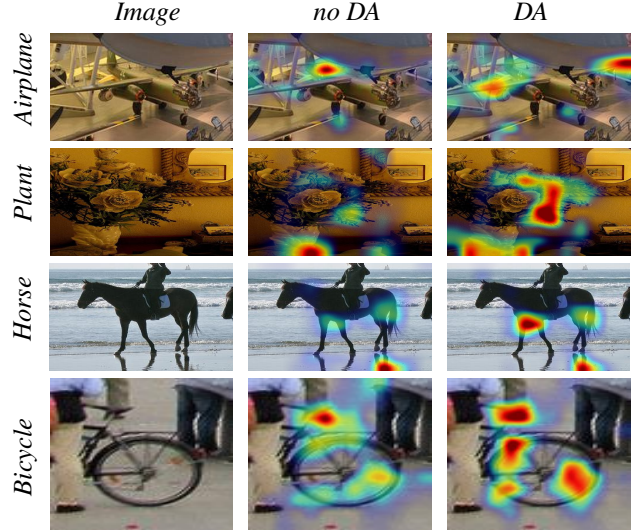


Figure 3. **Domain Shift.** We train two models on the Graphics source domain of Syn2Real, with and without Domain Adaptation (DA), and then test on the Real target domain. (Left) Three sample images from the target domain of the Syn2Real dataset that were misclassified before domain adaptation, then were correctly classified after domain adaptation. (Middle) The saliency of a classification model that does not employ domain adaptation. (Right) The saliency of a classification model that employs domain adaptation. It is clear how the evidence, after domain adaptation, is more focused on discriminative evidence of the GT class.

4. Method

We propose to explicitly disentangle domain specific information from otherwise relevant features using model explanations. We train models to produce saliency maps that are more explainable, in the sense that they better localize GT class objects. As we will show, this results in improved performance on unseen domains. We will refer to such explainable models as Explainable AI (XAI) models, and other vanilla CNN models trained using the conventional approaches without explicitly requiring any notion of explainability as noXAI models. At training time, we periodically (with frequency $freq$) force the model to focus, within an image x^i , on the objects corresponding to the GT label y^i using the GT spatial annotation g^i , rather than focus on the surrounding evidence which may be more domain specific. We assume that g^i is 1) a 2D binary map that has 1 in locations within the GT segmentation, and 0 otherwise; and 2) rescaled to the spatial dimension of the layer XAI is applied to. During an epoch where explainability is enforced, we compute saliency maps for the GT class and examine whether their peak overlaps with the GT spatial annotation. If overlapping occurs, we classify this saliency map to be explainable, *i.e.* classified for the correct reasons. Otherwise, we enforce explainability by utilizing the GT spatial annotation as an improved explanation.

We enforce focusing on objects in an image by scaling the forward activations of a particular spatial layer l in the network at certain epochs. We generate a multiplicative binary *mask* for guiding the focus of the network in the layer in which we are enforcing XAI. For an explainable image x^i , the binary mask is a binarization of the achieved saliency map, i.e. $mask_{j,k}^i = \mathbb{1}(s_{j,k}^i > 0) \forall j \forall k, j = 1, \dots, W$ and $k = 1, \dots, H$, where W and H are the spatial dimension of a layers' output neuron activations; The mask is active at locations of non-zero saliency. This re-inforces the activations corresponding to the active saliency regions that have been classified as being explainable. For images that need an improved explanation, the binary mask is assigned to be the GT spatial annotation $mask_{j,k}^i = g_{j,k}^i \forall j \forall k, j = 1, \dots, W$ and $k = 1, \dots, H$; The mask is active at GT locations. This increases the frequency at which the network reinforces activations at locations that are likely to be non-domain specific and suppresses activations at locations that are likely to be domain specific. We then perform element-wise multiplication of our computed *mask* with the forward activations of layer l ; i.e. $a_{j,k}^{l,i} = mask_{j,k}^i * a_{j,k}^{l,i} \forall j \forall k, j = 1, \dots, W$ and $k = 1, \dots, H$. Our XAI approach is summarized in Algorithm 1.

5. Experiments

In this section, we present experimental setup and results that quantify the explainability of XAI and noXAI classification models. We then present how XAI models lead to better single-source domain generalization.

Datasets. We use Microsoft Common Objects in Context (MSCOCO) [17] for training models using the noXAI and XAI strategies. MSCOCO provides GT spatial annotations consisting of object segmentations of the corresponding GT class. We use such annotations to guide the XAI training strategy. We then test how both strategies generalize from a single-source dataset to six unseen target domains from DomainNet [22] and Syn2Real [23]: graphics, clipart, infograph, painting, quickdraw, and sketch. To do so, we train models for the single-label classification task, once using MSCOCO as the single source, and another using the PASCAL VOC having similar annotations. We select the common subset of classes between source and target domains for each scenario. For MSCOCO, this leads to $\sim 25K$ training images belonging to the following classes: Airplane, Bicycle, Bus, Car, Horse, Knife, Motorcycle, Skateboard, Train, Truck. For PASCAL VOC this leads to $\sim 11K$ training images belonging to the following classes: Airplane, Bicycle, Bird, Bus, Car, Cat, Chair, Cow, Dog, Horse, Motorbike, Sheep, Television, Train.

Experimental Setup. For all experiments considering noXAI, XAI we focus on a vanilla Resnet-50 architecture. We resize input images to be 224×224 , and train for 50 epochs using a learning rate of 0.00001. Saliency maps are

Algorithm 1: XAI Training Strategy

Input: $x^i, i \in 1, \dots, m$ training images; $y^i, i \in 1, \dots, m$ corresponding class label; $g^i, i \in 1, \dots, m$ corresponding annotation; *freq* of training feedback; initial model M ; n epochs; layer l

Output: Trained Model M'

Procedure:

- 1 For every epoch $e \in 1, \dots, n$
- 2 if $e \bmod freq == 0$:
- 3 For every training example $x^i, i \in 1, \dots, m$
- 4 $s^i = \text{saliency}(x^i, y^i)$
- 5 $w, h = \text{argmax}(s^i)$
- 6 if $g_{w,h}^i == 1$
- 7 $\parallel j = 1, \dots, W; k = 1, \dots, H$
- 8 $mask_{j,k}^i = \mathbb{1}(s_{j,k}^i > 0) \forall j \forall k$
- 9 else:
- 10 $mask^i = g^i$
- 11 $a^{l,i} = mask^i * a^{l,i}$
- 12 Compute gradients and update weights
- 13 else:
- 14 Compute gradients and update weights

computed using the GradCAM [26] algorithm after the last block layer l of the ResNet-50. We choose the last spatial layer since it performed best, as it models higher level spatial patterns. In all experiments we set the frequency of XAI training to be five epochs. We compare the performance of our models against the data augmentation strategy by Volpi and Murino [36], implemented following the recipe proposed by the authors (random concatenation of five different transformations among *sharpness*, *brightness*, *color*, *contrast*, *RGB-to-grayscale conversion*, *RGB-channel perturbations* applied to each batch during training).

5.1. More Explainable Classification Models

In this section, we compare the explainability of noXAI vs. XAI classification models using an intuitive automated metric from the computer vision community, and then using human judgment in a crowdsourcing setting.

5.1.1 Automated Metric

We explore how the noXAI and XAI models perform on the unseen validation set of MSCOCO. We investigate how localized the evidence is with respect to the GT spatial annotations of MSCOCO. Figure 4 presents examples of saliency maps from the noXAI and XAI models after training is complete with comparable classification accuracy. Saliency is better localized over the object corresponding to the GT class when the XAI strategy is adopted for training. We use the pointing game of [41] to compute the number of hits; the number of correctly classified images where the peak

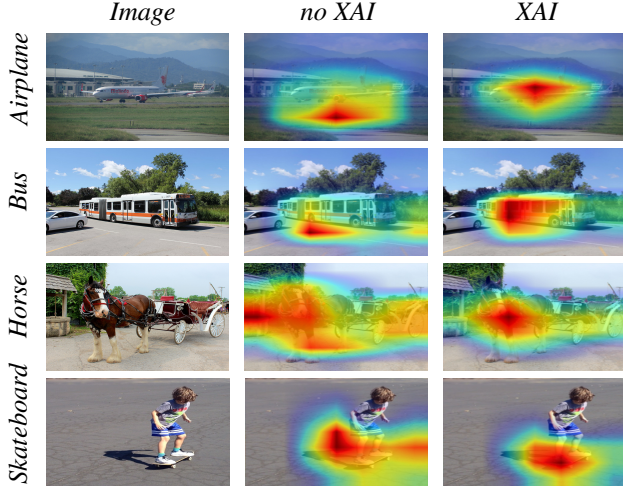


Figure 4. (Left) Four sample images from the unseen MSCOCO validation set that are correctly classified by both the noXAI and XAI models. (Middle) Saliency associated with the correctly predicted class using the noXAI model. (Right) saliency associated with the correctly predicted class using the XAI model. The XAI model, based on human spatial annotations, provides feedback that enables saliency to be better localized over the objects corresponding to the GT class compared to the noXAI vanilla training.

saliency overlaps with the GT spatial annotation. This is depicted in Figure 5 for the XAI and noXAI models over the training epochs to quantify explainability. We also demonstrate how data augmentation techniques of [36] do not significantly affect localization of explanation maps. At epoch 50, the model accuracy for MSCOCO’s validation set is: 61.87% for the noXAI model and 62.04% for the XAI averaged for three runs on the last four epochs. In essence, we are improving the explanation the model is providing for its classification result, assuming that a good explanation focuses more on the object corresponding to the GT class, without hindering classification accuracy. The XAI model has learnt to rely less on context information, without hurting the performance. At epoch 50, the XAI model has 1321 images with better localization/explainability as per our automated metric.

5.1.2 Human Judgment

In Section 5.1.1 we assumed that a saliency map whose peak overlaps with the GT spatial annotation of the object is a better explanation. We now design an unbiased crowdsourcing user study which asks users what they think is a better explanation for the presence of an object.

Annotation Tool Settings. We use the Amazon Mechanical Turk (AMT) crowdsourcing marketplace to recruit crowd workers. We accept AMT workers who had previously completed at least 1000 tasks (a.k.a HITs), and maintained an approval rating of at least 98%. We compensate

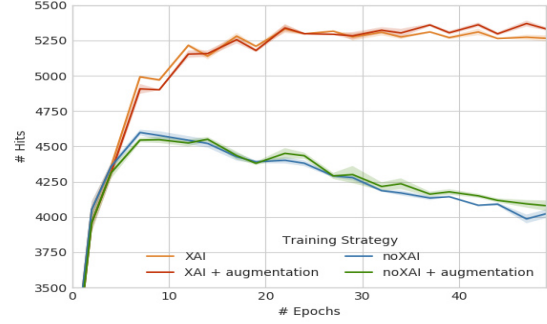


Figure 5. In this plot we present the number of hits, *i.e.* the number of unseen MSCOCO images, among the 16K validation set, where the model is able to provide an accurate explanation for, among the correctly classified ones during training. A model explanation is defined as a saliency map over the image, and a better explanation is defined to be one that has a higher number of pointing game hits, *i.e.* a higher number of image explanations overlapping with the annotation of the GT class label. We can see that the noXAI model fits the dataset bias at training time, while the XAI model improves its explainability over time for validation data. The standard deviation is depicted around the mean using a lighter shade.

the work of all crowd workers who participated in our tasks.

Crowdsourcing Task Details. We collect annotations from crowd workers for a task (HIT) that contains ten subtasks. Ten unique workers provide annotations for each task (annotating the ten subtasks). Each subtask presents the worker with one image and two evidence images, one is a saliency map generated from the model trained using our XAI strategy, and the second is a saliency map generated from the conventional noXAI training strategy. The worker is asked to select which of the evidence images (described to the user as “highlights”) is a better explanation for the presence of the object of interest. Figure 6 presents the sample interface presented to the worker for a particular subtask. We do not bias the users’ definition of explainability by not providing any specific example/instruction images. We post all HITs simultaneously, for all 1321 images, while randomizing the presentation order of XAI and noXAI saliency maps in each subtask. We allot a maximum of ten minutes to complete each HIT and paid \$0.10 per HIT.

Crowdsourcing Task Results. We processed the 13,210 (1,321 images * 10 votes per image) crowdsourced results to assign a winning label, “XAIwinner” or “noXAIwinner”, for each image. 190 unique crowd workers contributed to the 10 votes per HIT. It took each crowd worker an average of 137 seconds to complete a HIT consisting of ten voting subtasks. We used majority voting to combine the ten crowd-collected selections into a single vote for each image. For 225 out of the 1,320 images, there was no clear winner as there was a tie between XAI evidence and noXAI evidence. For the remaining 1,096 images, XAI evidence won

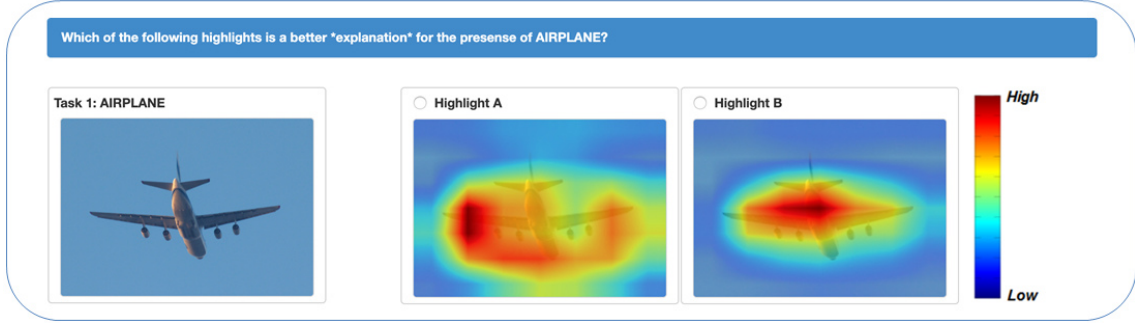


Figure 6. A snapshot of the interface presented to Amazon Mechanical Turk users for a subtask. The interface asks the users to select the evidence (“highlight”) they think is a better explanation for the presence of an object of a certain class, *e.g.* airplane, together with the original image. The order in which the XAI and noXAI evidence maps are presented is randomized for every subtask image.

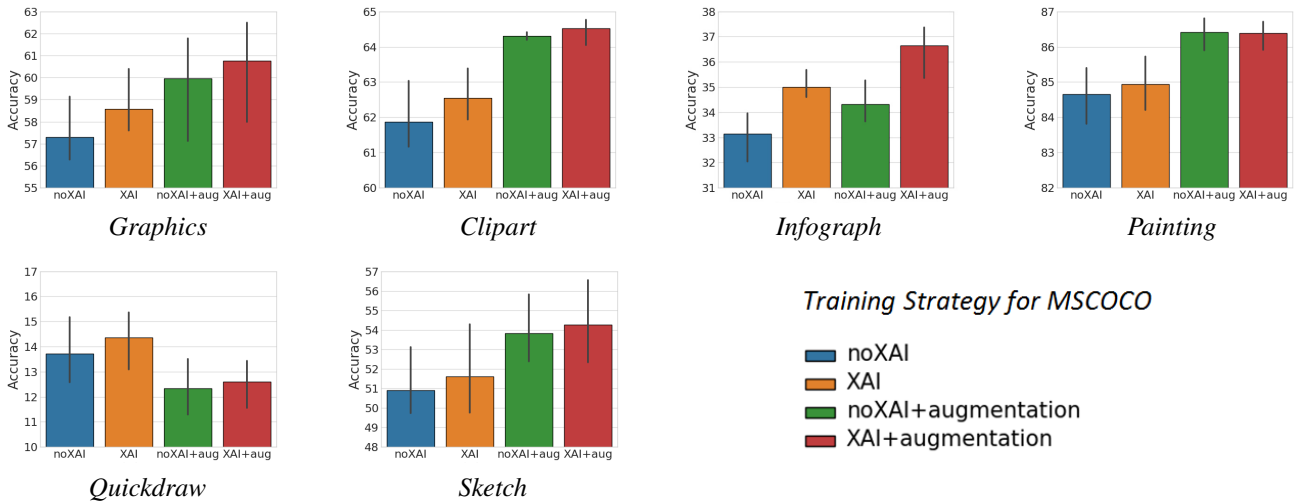


Figure 7. We present results for domain generalization on six *unseen* target domains. The Graphics domain data is obtained from the Syn2Real dataset, and the other five domains are obtained from the DomainNet dataset. We note that training has been conducted on a single source: the MSCOCO dataset, and no data from any of the target domains is presented to the model at training time. Results are demonstrated for the common classes between the three datasets. For each training strategy we report the average test accuracy of the last four epochs for three trained models. The black bar demonstrates the minimum and maximum average over three trained models.

for 887 of the images (67% of the whole image population and 80% of the images with a winner choice).

5.2. Improved Domain Generalization

In this section, we explore whether training a model to be explainable helps generalization to unseen domains. We consider the challenging problem of domain generalization from a *single-source*. This means that the target domains are *unseen* during training. We use MSCOCO as our single source and utilize its GT spatial annotations at training time to enforce a more explainable model. We test noXAI and XAI models on domains that are unseen at training time: graphics, clipart, infograph, painting, quickdraw, sketch.

Figure 7 presents the classification accuracy on unseen domains of XAI and noXAI models, with and without data augmentation [36]. Accuracy is the average of the last

four epochs of three trained models. The XAI model consistently results in improved domain generalization over the six target domains tested. The XAI training strategy achieves a relative 2.56% accuracy improvement averaged over all target domains. Figure 7 also demonstrates the complementarity of our proposed XAI training strategy to that of Volpi and Murino [36]. Figure 8 presents a similar trend when the PASCAL VOC is used as the single-source dataset for enforcing explainability, and unseen target domains are used for testing from DomainNet and Syn2Real.

By guiding the network to periodically focus on regions that contain the object of a ground truth class, we are able to train models that are more explainable in the sense that they classify based on the correct evidence, generalize better to unseen target domains, and suffer no degradation in performance on the original domain.

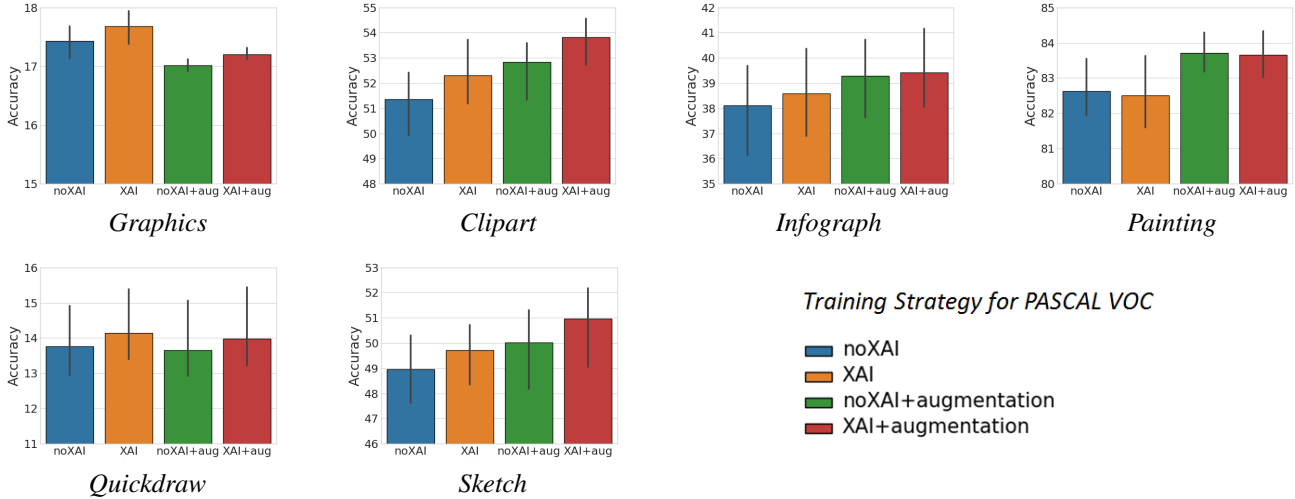


Figure 8. We present results for domain generalization on six *unseen* target domains. The Graphics domain data is obtained from the Syn2Real dataset, and the other five domains are obtained from the DomainNet dataset. We note that training has been conducted on a single source: the PASCAL VOC dataset, and no data from any of the target domains is presented to the model at training time. Results are demonstrated for the common classes between the datasets. For each training strategy we report the average test accuracy of the last four epochs for three trained models. The black bar demonstrates the minimum and maximum average over three trained models.

Domain	<i>block 1</i>	<i>block 2</i>	<i>block 3</i>	<i>block 4</i>
Graphics	58.67	58.02	57.09	58.56
Clipart	63.90	63.21	62.55	62.59
Infograph	32.74	32.68	32.91	35.16
Painting	85.30	85.31	85.10	84.98
Quickdraw	14.29	13.85	13.13	14.49
Sketch	51.74	51.40	51.32	51.76

Table 1. Test accuracies on the six unseen domains from DomainNet and Syn2Real, where XAI training was applied at different locations (After *block1*, *block2*, etc.) of the ResNet-50.

Domain	<i>freq = 5</i>	<i>freq = 10</i>	<i>freq = 15</i>
Graphics	58.56	58.06	58.08
Clipart	62.59	61.69	62.22
Infograph	35.16	33.91	34.07
Painting	84.98	84.86	84.96
Quickdraw	14.49	13.96	14.13
Sketch	51.76	51.04	51.42

Table 2. Test accuracies on the six unseen domains from DomainNet and Syn2Real, where XAI training was applied at *freq = 5, 10, 15*, *i.e.* every 5, 10, or 15 training epochs.

5.3. Where and When to apply XAI training

In this section we perform two ablation studies on MSCOCO to explore where and when XAI training could be applied to maximize performance gains.

Where to apply XAI training? In this section we explore where in a network architecture is best for applying the XAI training. This is where we would be computing the saliency map for the images, computing a mask, and applying the mask on forward activations. We apply XAI training after every block of the ResNet-50 architecture and present the results in Table 1. The majority of the domains had highest accuracy when XAI training was applied after the last block, however, we found that accuracies are comparable regardless to where XAI training was applied.

When to apply XAI training? We apply XAI training using *freq = 5, 10, 15* epochs in Table 2. Frequencies greater than 1 obtain a comparable performance, with best performance achieved at *freq = 5*. Applying XAI every it-

eration would involve a computational overhead and would completely disregard context hindering performance on the source and target domains.

Conclusions

In this work, we make object classification models possess the valuable property of explainability exposing their internal decision process in a human-interpretable way. We do so by periodically forcing the model at training time to focus on object/class evidence. We demonstrate that this training strategy encourages models to classify objects by looking at the objects themselves rather than their surrounding context, leading to better generalization across domains. While our approach leads to more explainable classification models and better generalization to unseen domains, it has no associated accuracy degradation on the original domain and no added test-time complexity.

References

- [1] Sarah Adel Bargal, Andrea Zunino, Donghyun Kim, Jianming Zhang, Vittorio Murino, and Stan Sclaroff. Excitation backprop for RNNs. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [2] Sarah Adel Bargal, Andrea Zunino, Vitali Petsiuk, Jianming Zhang, Kate Saenko, Vittorio Murino, and Stan Sclaroff. Guided zoom: Questioning network evidence for fine-grained classification. In *Proc. British Machine Vision Conference (BMVC)*, 2019. 3
- [3] Sarah Adel Bargal, Andrea Zunino, Vitali Petsiuk, Jianming Zhang, Kate Saenko, Vittorio Murino, and Stan Sclaroff. Guided zoom: Zooming into network evidence to refine fine-grained model decisions. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2021. 3
- [4] John Blitzer, Ryan McDonald, and Fernando Pereira. Domain adaptation with structural correspondence learning. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, 2006. 1, 2
- [5] Chunshui Cao, Xianming Liu, Yi Yang, Yinan Yu, Jiang Wang, Zilei Wang, Yongzhen Huang, Liang Wang, Chang Huang, Wei Xu, et al. Look and think twice: Capturing top-down visual attention with feedback convolutional neural networks. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, pages 2956–2964, 2015. 3
- [6] Fabio M. Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [7] Myung Jin Choi, Joseph J Lim, Antonio Torralba, and Alan S Willsky. Exploiting hierarchical context on a large database of object categories. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010. 2
- [8] R. C. Fong and A. Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2018. 3
- [9] Yaroslav Ganin and Victor S. Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proc. International Conference on Machine Learning (ICML)*, 2015. 2
- [10] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *Proc. International Conference on Learning Representations (ICLR)*, 2015. 2
- [11] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 2
- [12] Hal Daumé III and Daniel Marcu. Domain adaptation for statistical classifiers. *CoRR*, abs/1109.6341, 2011. 1, 2
- [13] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Deeper, broader and artier domain generalization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [14] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Proc. AAAI Conference on Artificial Intelligence (AAAI)*, 2018. 2
- [15] Da Li, Jianshu Zhang, Yongxin Yang, Cong Liu, Yi-Zhe Song, and Timothy M. Hospedales. Episodic training for domain generalization. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [16] Kunpeng Li, Ziyang Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. Tell me where to look: Guided attention inference network. *arXiv preprint arXiv:1802.10171*, 2018. 3
- [17] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *Proc. European Conference on Computer Vision (ECCV)*, 2014. 5
- [18] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. *arXiv preprint arXiv:1605.06636*, 2016. 4
- [19] Saeid Motiian, Marco Piccirilli, Donald A. Adjeroh, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [20] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *Proc. International Conference on Machine Learning (ICML)*, 2013. 2
- [21] Aude Oliva and Antonio Torralba. The role of context in object recognition. *Trends in cognitive sciences*, 11(12):520–527, 2007. 2
- [22] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 5
- [23] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshop (CVPRw)*, 2018. 3, 5
- [24] Vitali Petsiuk, Abir Das, and Kate Saenko. RISE: Randomized input sampling for explanation of black-box models. In *Proc. British Machine Vision Conference (BMVC)*, 2018. 3
- [25] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *Proc. European Conference on Computer Vision (ECCV)*, 2010. 1, 2
- [26] 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 *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017. 2, 3, 5
- [27] Ramprasaath R Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Dhruv Batra, and Devi Parikh. Taking a hint: Leveraging explanations to make vision and language models more grounded. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 3
- [28] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. In

- Proc. International Conference on Learning Representations (ICLR)*, 2018. 2
- [29] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *Proc. International Conference on Learning Representations (ICLRw)*, 2014. 3
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3
- [31] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin A. Riedmiller. Striving for simplicity: The all convolutional net. *CoRR*, abs/1412.6806, 2014. 3
- [32] Baochen Sun and Kate Saenko. Deep CORAL: correlation alignment for deep domain adaptation. In *Proc. European Conference on Computer Vision Workshop (ECCVw)*, 2016. 2
- [33] Joshua Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. *CoRR*, 2017. 2
- [34] A Torralba, KP Murphy, WT Freeman, and MA Rubin. Context-based vision system for place and object recognition. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2003. 2
- [35] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [36] Riccardo Volpi and Vittorio Murino. Addressing model vulnerability to distributional shifts over image transformation sets. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 2, 3, 5, 6, 7
- [37] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 2, 3
- [38] Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [39] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *Proc. European Conference on Computer Vision (ECCV)*, 2014. 3
- [40] Jianming Zhang, Sarah Adel Bargal, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. Top-down neural attention by excitation backprop. *International Journal of Computer Vision (IJCV)*, 126(10):1084–1102, 2018. 3
- [41] Jianming Zhang, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. Top-down neural attention by excitation backprop. In *Proc. European Conference on Computer Vision (ECCV)*, pages 543–559. Springer, 2016. 5
- [42] Xiaolin Zhang, Yunchao Wei, Jiashi Feng, Yi Yang, and Thomas Huang. Adversarial complementary learning for weakly supervised object localization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [43] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [44] Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Weakly supervised instance segmentation using class peak response. *arXiv preprint arXiv:1804.00880*, 2018. 3
- [45] Andrea Zunino, Sarah Adel Bargal, Pietro Morerio, Jianming Zhang, Stan Sclaroff, and Vittorio Murino. Excitation dropout: Encouraging plasticity in deep neural networks. *International Journal of Computer Vision (IJCV)*, 2021. 3
- [46] Andrea Zunino, Jacopo Cavazza, Riccardo Volpi, Pietro Morerio, Andrea Cavallo, Cristina Becchio, and Vittorio Murino. Predicting intentions from motion: The subject-adversarial adaptation approach. *International Journal of Computer Vision (IJCV)*, 2020. 2