

Invisible Shortcuts: Why Vision Encoders Know Your Camera

Vladan Stojnić¹, Ryan Ramos², Giorgos Kordopatis-Zilos¹,
Noa Garcia², and Giorgos Tolas¹

¹ VRG, FEE, Czech Technical University in Prague, Prague, Czechia

² The University of Osaka, Osaka, Japan




Fig. 1: Metadata distribution per semantic category. Different classes (ImageNet21k) and topics (Re-LAION-2B) have distinct distributions, indicating correlations between metadata and semantics. These correlations create shortcut signals that models exploit during large-scale pretraining. Global: overall dataset distribution.

Abstract. Deep vision models exploit shortcuts, relying on cues that correlate with supervision signals. Prior work has focused on visible biases, such as object-background or texture correlations. We identify a different source of shortcut learning: invisible metadata traces embedded at the pixel level, for metadata such as image processing and photo acquisition. We hypothesize that large-scale semantic supervision, whether through categorical labels (ImageNet) or billion-scale captions (LAION), naturally induces metadata-semantics correlations during pretraining, leading models to convert low-level signals into predictive features. By introducing controlled metadata-semantics correlations, we show that stronger ones produce systematically higher sensitivity to metadata traces and larger performance degradation under metadata distribution shifts. We further explore mitigation strategies applied during and after pretraining that reduce sensitivity not only to targeted metadata but also to unseen ones, without sacrificing performance on downstream tasks. Metadata sensitivity also has a positive side: it partly explains the strong generated-image detection ability of some encoders, while its mitigation can improve out-of-distribution generalization. Code: <https://github.com/ryan-caesar-ramos/visual-encoder-traces>

1 Introduction

Deep neural networks are known to rely on shortcuts [4, 14, 60, 64, 70, 78]. When a predictive cue is present in training data, models exploit it regardless of whether it aligns with the intended task, since the objective is simply to minimize error.

Whether this is harmful depends on the stability of these correlations; when they shift, performance degrades. This phenomenon, studied under shortcut learning [14] and spurious correlations [33, 54], has been extensively analyzed in vision, including object–background dependencies [21, 43, 62], color–label correlations [39], texture versus shape bias [15], watermark artifacts [32], and synthetic benchmarks [2, 54]. These cues are visible and interpretable to humans.

We study shortcuts largely invisible to the human eye, namely metadata traces, including processing parameters such as JPEG compression and sharpening, and acquisition parameters such as focal length and camera model. Although imperceptible at a semantic level, they leave low-level signals.

Recent work shows that distribution shifts in image metadata affect the predictions of pretrained vision models [50], revealing that representations encode metadata traces, but the origin remains unclear. We argue that metadata sensitivity arises from metadata–semantics correlations in pretraining datasets: images from certain semantic categories are often captured or processed under similar conditions, creating systematic associations between labels or captions and metadata attributes such as camera parameters or JPEG quality (Fig. 1). Models trained with semantic supervision exploit these correlations as shortcuts, relying on metadata-derived pixel signals rather than semantic content, as such signals are simpler than high-level concepts and optimization favors them [60].

We test this hypothesis across both regimes of large-scale semantic supervision: categorical label supervision on ImageNet and caption supervision on LAION, measuring the metadata–semantics correlations in each and introducing controlled ones during training. As correlation strength increases, models become progressively more sensitive to metadata. Motivated by this, we investigate mitigation both during and after pretraining. Interestingly, using a single attribute, *e.g.* JPEG compression, as the correlation source or mitigation target also changes sensitivity to other metadata, *e.g.* resizing or camera model, suggesting that models exploit a broad class of low-level traces rather than attribute-specific artifacts. Notably, this contradicts prior work [32], where mitigating one type of bias amplified others.

Our main contribution is to show that metadata sensitivity is a result of shortcut learning arising from metadata–semantics correlations in training data, manifesting at scale across both supervised ImageNet and caption-guided LAION pretraining, rather than being specific to a single supervision paradigm. We further show it can be mitigated, both during and after pretraining, with mitigation targeting one attribute generalizing to unseen ones. Finally, such sensitivity is not only a disadvantage: since it reflects a model’s tendency to capture pixel-level signal, which is key to separating generated from real images, it partly explains the strong generated-image detection ability of some encoders [44]. Meanwhile its mitigation can improve out-of-distribution generalization.

2 Diagnostic metrics

We use metrics that quantify (i) the extent to which frozen model representations encode metadata and (ii) the degree to which such encoding interferes with semantic prediction ability. We follow Ramos *et al.* [50], see original work for details. We consider two metadata categories: (i) *image processing attributes*: JPEG compression parameters (quality, chroma subsampling), sharpening (strength), resizing (factor), and interpolation (algorithm type), and (ii) *acquisition attributes*: camera make/model, exposure time, aperture, ISO, and focal length.

Metadata Prediction (MP). We measure the strength of metadata encoding in frozen model representations by training a linear classifier to predict metadata labels. For processing metadata, we use ImageNet1k (IN1k) [53], randomly assign processing labels to training and test images, and process them accordingly. For acquisition metadata, we use FlickrExif [50] with predefined train/test splits per acquisition attribute. Accuracy is denoted by MP_p and MP_a for processing and acquisition metadata, respectively. Accuracy substantially above chance indicates that the representation contains linearly separable metadata information, while accuracy close to chance suggests invariance to metadata labels.

Semantic Prediction (SP) ability is evaluated by the accuracy of a k -NN classifier predicting semantic labels with original IN1k train/val sets.

Semantic Prediction Distraction – processing (SPD_p). To quantify the impact of processing-related metadata on semantic predictions, we use IN1k with a k -NN classifier for semantic labels, where metadata labels of training/exemplar images are modified per test image guided by semantic labels. Images are processed according to the assigned metadata label. Two evaluation setups: (i) *Pos-same*: the test image and its semantic positives are assigned to the same processing label, while negatives are assigned to a different one. Accuracy in this setup is denoted by $A_{pos-same}$. (ii) *Neg-same*: the test image and its semantic negatives are assigned to the same processing label, while positives are assigned to a different one. Accuracy in this setup is denoted by $A_{neg-same}$. We evaluate distraction with $\Delta_p = |A_{pos-same} - A_{neg-same}|$ quantifying the extent to which processing metadata impacts semantic recognition. Lower values indicate stronger invariance to processing metadata, while higher values indicate stronger sensitivity.

Semantic Prediction Distraction – acquisition (SPD_a). To measure how acquisition metadata interferes with semantic similarity, we use the PairCams [50] dataset. Given a query image from one camera class *e.g.* smartphones, the task is to retrieve a semantically positive image captured with a different camera class *e.g.* digital cameras. Using recall@1, denoted by R , we evaluate two cases where semantically negative images vary: (i) *same*: negatives originate from the same camera class as the query; performance denoted by R_{same} , and (ii) *different*: negatives originate from a different camera class than the query; performance denoted by R_{diff} . We define the distraction score as $\Delta_a = |R_{same} - R_{diff}|$, quantifying the extent to which acquisition metadata impacts semantic recognition.

Overall. MP quantifies metadata encoding, SPD quantifies metadata interference with semantic recognition, and SP captures whether training data correlations or mitigation strategies degrade semantic prediction performance.

3 Why are vision encoders sensitive to image metadata

The work of Ramos *et al.* [50] shows that vision encoders are sensitive to image metadata, and that the presence of metadata traces in features of pretrained models has consequences on downstream performance. However, the causes of these traces remain unclear. We hypothesize that they stem from correlations between metadata and semantic labels in pretraining datasets: such correlations let metadata labels and their pixel-level signal act as a shortcut for predicting semantic labels, leading the model to encode metadata information. We corroborate our hypothesis in several steps. We first show that these correlations, with both processing and acquisition metadata, exist across standard pretraining datasets, spanning categorical (ImageNet) and caption-based (LAION) supervision. We then use controlled experiments to verify that an encoder’s metadata sensitivity grows as a direct function of correlation strength.

3.1 Correlations of semantics and processing metadata in ImageNet

Measuring existing correlations. We focus on JPEG compression parameters, specifically the quality and chroma subsampling, as processing metadata. Each combination of values for these two parameters defines a discrete categorical label for each image. We analyze the association between these metadata and semantic labels in ImageNet1k (IN1k) and ImageNet21k (IN21k) using Cramér’s V [7] measure of association between categorical variables. The values are 0.047 and 0.067 in IN1k and IN21k, respectively. Even though the value itself is low and not easily interpretable, we focus on the relative comparison between the two datasets; correlations are stronger in IN21k than in IN1k.

Impact of correlations. We use models³ pretrained on IN1k and IN21k to evaluate the impact of correlation strength. In Fig. 2, we present all the diagnostic metrics, SP, MP, and SPD, for different backbones. We observe that both types of models have a good ability to predict metadata labels (MP), both processing and acquisition ones, at rates substantially above random chance. While there exist minor differences in training recipes, the models trained on IN21k consistently demonstrate stronger predictive ability. This indicates that features trained on IN21k encode more metadata information than those trained on IN1k. The SPD measurements further show that metadata affects the semantic prediction abilities of both models, but more so for IN21k. Together, these results show that IN21k features not only encode more metadata information but also let it interfere more strongly with semantic prediction. Overall, these results suggest that the stronger correlations between metadata and semantic labels in IN21k, relative to IN1k, have a greater impact on both the models’ ability to encode metadata information and the sensitivity of their semantic predictions to metadata.

³ We use pretrained models available in the timm library.




Fig. 2: IN1k vs. IN21k pretraining: which training set causes higher meta-data sensitivity? Reporting Metadata Prediction (MP), Semantic Prediction (SP), and Semantic Prediction Distraction (SPD) for ConvNeXt [37], FlexiViT [5], Swin [35], and ViT [10] pretrained on IN1k and IN21k. The amount of correlation between JPEG compression parameters and semantic labels is measured to be 0.047 and 0.067 in IN1k and IN21k, respectively, according to Cramér’s V measure. The stronger correlations in IN21k result in consistently more sensitive models.

3.2 Controlling correlations of semantics and processing metadata

To further verify that the presence and extent of metadata-semantics correlations impacts model’s sensitivity to metadata, we synthetically introduce correlations into the IN1k training set, train models on these correlated datasets, and evaluate all diagnostic metrics.

Creating stronger correlations. We define 22 distinct metadata labels by combining 11 JPEG quality levels with 2 chroma subsampling modes. All images are then recompressed according to one of these labels, creating a correlation between semantic labels and metadata labels. We randomly assign one specific JPEG compression, *i.e.* metadata label, to each semantic class. We introduce a probability p_i that determines whether an image is processed using the metadata label assigned to its semantic class or using a randomly sampled metadata label. When $p_i = 100$, correlation is maximal, while when $p_i = 0$, metadata is independent of semantics. Note that Cramér’s V is roughly equal to the value of p_i , *e.g.* 0.1 and 0.8 for $p_i = 10$ and $p_i = 80$, respectively.

Impact of correlations. We train ResNet50 [18]⁴ on these versions of the IN1k training set created using the mechanism described above, and evaluate the models using the diagnostic metrics. Fig. 3 present the results, where the larger values of p_i mean stronger correlation between metadata and semantics.

⁴ ViT-S/16 [10], and ViT-B/16 [10] results are presented in the supplementary.




Fig. 3: Impact of correlations between semantics and processing metadata on the model’s sensitivity to metadata. ResNet50 trained on different versions of IN1k with manipulated JPEG labels. The larger the p_i , the stronger the correlation between JPEG labels and semantic labels. *None*: model trained on original IN1k. Difference between $p_i = 0$ and *none*: the former has no JPEG-semantics correlations, while the latter has the original IN1k correlations. SP: two types of test sets are used, the original one (*none*, green color) which is fixed across values of p_i , and 7 different versions (purple color) created by manipulating the JPEG labels of the original test set in the same way as the corresponding training set for each value of p_i .

We observe that accuracy of predicting JPEG metadata labels (MP) increases noticeably for stronger correlations. For zero correlations it is close to random chance, and lower than for the model trained on the original IN1k, verifying that the original IN1k has some correlations, though weaker than the ones we introduce with $p_i \geq 20$. Since the model is trained on a set where JPEG compression drives the correlation, its ability to predict JPEG labels is expected. More notably, the ability to predict other processing and acquisition attributes (resizing, make, aperture) also increases with the JPEG-semantics correlation, suggesting that the model learns to encode pixel-level signal carrying such metadata.

The semantic prediction ability (SP) is affected by the correlation strength (green bars of bottom-left plot). There is no noticeable drop for weak correlations compared to zero correlations, but as the correlation becomes stronger, the drop becomes more significant. This suggests that strong correlations make the model capture pixel-level signal related to metadata labels rather than complex semantic-level features. Performance is additionally reported for testing on 7 modified versions (purple bars of bottom-left plot) of the IN1k validation, each processed using the same metadata-to-semantics associations as the training set for the corresponding value of p_i . Observe that when evaluating on these transformed test sets, the accuracy increases for increased correlations; this is because the test set is transformed identically to the training set. The model uses the same pixel-level signal that it exploited during training, which is not the case when testing on the original set.

Finally, we observe that the distraction of semantic prediction abilities by metadata (SPD) also increases with the increase in correlation, which suggests that the model relies more on the metadata information for semantic prediction when the correlation is stronger, and thus the interference of metadata information with semantic prediction becomes stronger as well.

3.3 Correlations of semantics and acquisition metadata in LAION

We additionally investigate whether correlations exist between semantics and acquisition metadata. The subset of Re-LAION-2B [28] carrying Exif tags, about 40M images, provides an image set annotated with acquisition metadata labels. Since Re-LAION-2B contains captions rather than categorical class labels, we apply topic modeling [17] to cluster caption embeddings and assign each caption to one of the resulting topics, which serves as a discrete semantic label per image. Fig. 1 illustrates examples of such correlations, where for a given acquisition attribute, we plot the frequency distribution of its metadata labels within Re-LAION-2B for two topics with clearly distinct distributions. For instance, images whose captions fall under the topic of football tend to have longer focal lengths than those under plots of land. Similar distinctions arise between weddings and eyewear for aperture, and between birds and kitchens for exposure time. We presume that such correlations can act as a source of shortcut learning, analogous to the processing metadata case, which we verify in the following.

3.4 Controlling correlations of semantics and acquisition metadata

Creating stronger and weaker correlations. We create three Re-LAION-2B subsets of equal size, *i.e.* 6.4M images, with varying amounts of correlation between semantic topics and acquisition metadata labels. Metadata labels are defined as tuples comprising the labels for the following attributes: camera manufacturer, aperture, exposure, and ISO speed. (i) *baseline*: we sample uniformly, reflecting the natural correlations present in Re-LAION-2B. (ii) *stronger*: for each metadata label, we select images belonging to its 40 most frequent topics, concentrating each label on a few topics and thus strengthening the topic-metadata association. (iii) *weaker*: for each metadata label, we identify its least frequent topic and downsample all other topics to that size, equalizing the topic distribution within each metadata label and thus weakening the association.

Impact of correlations. We train ResNet50 models⁵ using CLIP loss [49], with results in Fig. 4. Training on the subset with stronger correlations consistently leads to higher SPD, while training on the weaker one tends to yield lower SPD, indicating that acquisition metadata correlations are also exploited as shortcuts. Furthermore, consistently with Sec. 3.2, models trained on the weaker subset exhibit the strongest semantic prediction ability, and those trained on the stronger subset the weakest. The MP results are less conclusive; nevertheless, interference with semantic predictions reflected by SPD is the key signal of metadata sensitivity.

⁵ ViT-S/16 [10] and ConvNeXt-T [37] results are presented in the supplementary.




Fig. 4: Impact of correlations between semantics and acquisition metadata on the model’s sensitivity to metadata. ResNet50 trained from scratch using CLIP loss [49] on three Re-LAION-2B subsets with Cramér’s V correlations of 0.396, 0.255, and 0.166, referred to as stronger, baseline, and weaker, respectively.

4 Analysis of mitigation approaches

We investigate mitigation in two settings: during training on a standard benchmark (IN1k), and as post-processing of released, widely used foundational models that already exhibit metadata sensitivity. Unlike Sec. 3, these experiments do not rely on artificially introduced correlations, but instead target sensitivity reduction under naturally occurring ones.

4.1 Mitigation during training with augmentations

Augmentations during training. Ramos *et al.* [50] argue that strong data augmentations during training reduce metadata sensitivity. We explore this by training and testing ResNet50 on original IN1k. The default training recipe includes RandomAugment (RA) [8], while we remove some or add additional augmentations to test their impact.

Impact of augmentations. Fig. 5 presents the results of this experiment. Removing any of the RA augmentations commonly increases the model’s ability to capture metadata and its sensitivity. Adding Gaussian blur noticeably reduces sensitivity to processing metadata, but has little impact on acquisition metadata; blurring is thus relatively effective at removing the pixel-level signal that metadata leaves, while harming semantic prediction (SP) only slightly. Most notably, the combination of color jittering, blurring, and grayscale conversion consistently reduces sensitivity by a large factor for both metadata types. This is the set of augmentations used by DINOv2 [45], which is reported to be among the least sensitive foundational models [50].




Fig. 5: Mitigation during training with augmentations – impact on model sensitivity to metadata. ResNet50 trained on IN1k with a varying set of augmentations. The default case is Random Augment (RA) [8], while we remove or add augmentations from/to the default. *all*: color jitter, grayscale, and blur.

4.2 Post-hoc adversarial mitigation

Adversarial mitigation. Given a pretrained model that maps an image to a D -dimensional representation space, we train a linear layer $f : \mathbb{R}^D \rightarrow \mathbb{R}^D$, as a mitigation component, to transform the representation space and reduce metadata sensitivity. A linear classifier $h : \mathbb{R}^D \rightarrow \mathbb{R}^M$, applied to the transformed space, predicts among M metadata labels. Training is adversarial [1]: h tries to predict metadata labels, while f tries to make the predictions of h as bad as possible. It proceeds in two steps. (i) *Train the metadata prediction model h while keeping the mitigation component frozen*, via cross-entropy loss and supervision from metadata labels. (ii) *Train the mitigation component f while keeping the metadata prediction model frozen*, using three loss terms: one that keeps the transformed feature $f(x)$ close to the pretrained feature by maximizing their dot product $f(x)^\top x$; one that pushes apart features $f(x)$ and $f(y)$ of images sharing the same metadata label by minimizing their dot product $f(x)^\top f(y)$; and one that maximizes the entropy of the metadata prediction $h(f(x))$ for each feature x . These two steps are repeated L times. The process requires access to metadata labels, which we obtain by reprocessing IN1k with different JPEG compression parameters; can be used with other attributes similarly.

Impact of adversarial mitigation. Fig. 6 presents the results for this mitigation strategy for different contrastive vision language (CLIP [49], ConvNeXt [37], SigLIP [80]), supervised (ViT [10], ConvNeXt [37]), and self-supervised (DINO [6], DINOv2 [45]) models. We observe that such process consistently reduces the sensitivity for a variety of models, both in terms of metadata prediction and also in terms of semantic prediction distraction. The most interesting fact is that the mitigation training is guided by the JPEG compression parameters, but it also reduces the sensitivity to other types of metadata. Regarding the overall semantic prediction ability, it is either maintained or even slightly improved.




Fig. 6: Post-hoc adversarial mitigation with a linear layer – impact on model sensitivity to metadata. Performance is reported for the original frozen foundational models and for the models adjusted by an additional linear layer trained to mitigate the metadata sensitivity guided by manipulated JPEG labels on IN1k.

5 Further analysis

Do Stable Diffusion-generated images carry metadata traces? We hypothesize that images generated by models such as Stable Diffusion [52] may not contain pixel-level metadata traces, since they do not arise from a physical image formation process and the post-processing pipeline can be kept uniform after generation. To test this, we train models on synthetically generated images for IN1k classes [12, 56] compressed with fixed high-quality JPEG parameters. The resulting models still encode and rely on metadata, with MP_p - JPEG: 29.91, MP_p - Resizing: 61.58, SPD_p - JPEG: 2.00, and SPD_p - Resizing: 6.06, comparable to training on the original ImageNet (*none*) in Fig. 3. This suggests generative models may implicitly reproduce pixel-level traces correlated with the semantics in their training data.

Metadata sensitive models are good detectors of generated images.

Recent work [44, 59, 63] shows that frozen CLIP features are effective at detecting generated images, which we hypothesize is due to their strong metadata encoding ability. To test this, we apply a k -NN classifier ($k = 1$) to detect generated images, as in Ojha *et al.* [44], using features from the models in Sec. 3.2 with varying metadata sensitivity. Tab. 1 shows that more sensitive models, trained on strong JPEG-semantics correlations, *i.e.* $p_i = 50$ or 100, detect generated images noticeably better than the one trained on the original IN1k. This validates our hypothesis: metadata-semantics correlations encourage the model to focus on low-level details, a beneficial trait for synthetic image detection.

Model	Generative Adversarial Networks						Deep fakes	Low level vision		Perceptual loss		Guided	LDM		Glide			DALL-E	Avg	
	Pro-GAN	Cycle-GAN	Big-GAN	Style-GAN	Gau-GAN	Star-GAN		SITD	SAN	CRN	IMLE		200 steps	200 w/ CFG	100 steps	100	50			100
$p_i = \text{none}$	69.1	65.3	49.5	53.5	58.7	55.4	51.1	<u>55.0</u>	49.8	57.3	59.6	<u>51.2</u>	<u>62.2</u>	53.4	62.6	58.3	<u>60.7</u>	<u>58.7</u>	<u>57.4</u>	57.3
$p_i = 50\%$	71.4	<u>66.4</u>	49.0	58.6	58.1	60.1	<u>56.1</u>	<u>54.7</u>	52.3	64.7	68.2	49.9	61.3	53.7	<u>61.3</u>	59.5	60.4	57.9	<u>57.4</u>	<u>59.0</u>
$p_i = 100\%$	77.2	67.0	48.4	<u>58.4</u>	<u>51.4</u>	76.0	58.7	58.3	53.9	61.3	<u>61.0</u>	55.0	62.9	56.2	60.3	60.2	62.5	61.1	57.5	60.4

Table 1: Metadata sensitive models are good detectors of generated images. Used models are ResNet50 trained on different versions of IN1k with manipulated JPEG labels from Sec. 3.2 (original IN1k/*none*, $p_i = 50$ and $p_i = 100$ with low, strong and maximal correlations, respectively). Binary classification accuracy reported.




Fig. 7: Ability of internal model layers to predict metadata. Metadata prediction (MP) accuracy across blocks of ResNet50 trained on: (i) original IN1k, (ii) original IN1k using extra augmentations (color jitter, grayscale, and blur augmentations), and (iii) modified IN1k with strong JPEG-semantics correlations. In all cases Random Augment is used.

Layer-wise sensitivity analysis. In Fig. 7, we examine the metadata prediction ability across the 5 blocks of ResNet50 models trained in different ways. All models reach a strong metadata prediction ability already after two blocks, since metadata traces are low-level signals easily captured by early layers [79]. The ability further improves after the third block for all models, but then diverges: it starts dropping in deeper layers for the model trained on the original IN1k, while continuing to improve for the model trained on the modified IN1k with strong correlations. This shows that the former learns to ignore metadata in deeper layers, delegating semantic feature extraction to them, whereas the latter learns to rely on metadata and cannot ignore it, suggesting that semantic discrimination is compromised in favor of metadata reliance. The use of mitigation effective augmentations, as identified in Sec. 4.1, reduces the MP accuracy. Interestingly, the model trained on strong JPEG-semantics correlations behaves similarly for resizing: it captures the resizing metadata in early layers and cannot ignore it in deeper ones, even though it was not trained with strong resizing-semantics correlations. This is further evidence that the model learns to rely on low-level traces in general, not only on the specific attribute strongly correlated with semantics during training.




Fig. 8: t-SNE visualization of features before (top) and after (bottom) post-hoc mitigation for two classes (●, ▲) from IN1k processed by different JPEG parameters (■, ■). Used models are ResNet50 trained on different versions of IN1k with manipulated JPEG labels from Sec. 3.2 (original IN1k/*none*, $p_i = 80$ and $p_i = 100$ with low, strong and maximal correlations, respectively).

Representation space visualization before/after post-hoc mitigation.

Fig. 8 provides a t-SNE [38] visualization of ResNet50 features for two IN1k classes, where all images are processed twice with two different sets of JPEG parameters. The models are trained on the original IN1k and on modified versions with strong processing metadata-semantics correlations (Sec. 3.2), shown before and after post-hoc mitigation. For the model trained on the original IN1k, the classes are decently separated while points with different JPEG parameters overlap well; the small signs of misalignment are consistent with the quantitative measurements in Fig. 3, indicating mild sensitivity to metadata. The model thus learns semantic features and is only subtly affected by metadata. In contrast, the two classes start to mix for models trained under strong correlations, while points with different JPEG parameters become more separated, showing that the model relies on metadata instead of semantics and that metadata leaves strong traces in the representation space. At $p_i = 100\%$, the classes are almost completely mixed, and the features can instead be separated by JPEG processing. Finally, post-hoc mitigation improves semantic class separability even for the model trained with $p_i = 100\%$ correlation.

Impact of mitigation on OOD tasks. Since metadata traces can be a spurious low-level signal that models exploit as a shortcut, reducing reliance on them should be most beneficial under distribution shift, where such shortcuts no longer hold. We therefore investigate whether reducing the models’ metadata sensitivity improves performance on common out-of-distribution (OOD) ImageNet test




Fig. 9: Post-hoc adversarial mitigation with a linear layer – impact on ImageNet OOD test sets. Performance is reported for the original frozen foundational models and for the models adjusted by an additional linear layer trained to mitigate the metadata sensitivity guided by manipulated JPEG labels on IN1k. Accuracy reported on ImageNet-C instead of the commonly reported error.

sets: ImageNet-C [20], ImageNet-R [19], and ImageNet-Sketch [68]. Fig. 9 reports results without and with the post-hoc mitigation. The mitigation helps in many cases, sometimes considerably, and the gains are largest for the models most sensitive to metadata (according to Fig. 6: CVLs and supervised ConvNeXt), while sometimes hurting less sensitive ones like DINOv2. We observe similar trends for ResNet50 trained on the original IN1k (ImageNet-C: 44.2 $\rightarrow$ 45.0, ImageNet-R: 34.5 $\rightarrow$ 37.7, ImageNet-Sketch: 26.4 $\rightarrow$ 30.0) when mitigating during training with augmentations.

6 Related work

6.1 Shortcut learning

In empirical risk minimization (ERM) settings, neural networks have been shown to rely on simple input features as shortcuts to solve the task being learned, even if the shortcuts are only spuriously correlated with features more relevant to the task [14, 64]. This is because under ERM, the objective is simply to minimize error, meaning that shortcut features will be learned as long as they help with the training objective [40, 55]. Not only are shortcut features learned, but they are even preferred by models due to their simplicity. Even in the presence of features that are more salient and more informative, neural networks will still rely more strongly on the simple spurious features if the salient features are more complex to learn [60].

A frequent origin of shortcuts is dataset bias. Due to data sources and collection methods [36, 61], as well as their properties such as confounding factors [78], data scarcity [46], collection location [13], and collection timeframe [51], resulting datasets may exhibit unique characteristics to the point that it is possible to distinguish them from one another despite having similar design goals [36, 66]. More importantly, among these idiosyncrasies may be correlations that create

exploitable shortcuts. While these commonly manifest as object-background correlations, as is the case with ImageNet [21, 43, 62], *e.g.* the side of freight cars being associated with graffiti, non-object-centric correlations also exist. In LAION, images of cardboard boxes are commonly accompanied by watermarks [32]. Even spatial characteristics, such as object location, and lower-level features, such as color distributions, have been found to be correlated with different semantic concepts [23, 39]. Models trained on these datasets in standard ERM settings typically learn these shortcuts and predict the semantic label even when the image contains only the shortcut feature [21, 32, 43, 62].

In contrast to the aforementioned immediately perceptible cues, we focus on largely imperceptible shortcuts, mainly JPEG compression artifacts. Previous studies that have examined their role in shortcut learning focus on specialized tasks [16] or exist outside the context of their correlation with semantic labels [9], while our work investigates their presence in large general-purpose training datasets and consequent ability to interfere with foundation models.

Additionally, neural networks have been shown to rely on frequency-based shortcuts [71, 72], *e.g.* zebra patterns being easily distinguishable in the spectrum, and such shortcuts arise even in models trained at ImageNet scale. Although some types of metadata could also introduce frequency-related shortcuts, these works do not account for them.

6.2 Shortcut mitigation

A diverse range of strategies has been proposed to mitigate shortcut learning. One prominent direction is *data-centric*, where the training distribution is modified to reduce spurious correlations. These approaches include data augmentation techniques [15, 47, 76], such as image stylization or mixup, which expand the dataset with semantically consistent transformations. Other methods rely on resampling or reweighting strategies [27, 30] to balance class distributions, as well as concept discovery and pseudo-label extraction techniques [42, 58, 73] that incorporate additional supervisory signals. A second line of work mitigates shortcuts directly during training. Some approaches introduce additional classification heads [1, 24] or train parallel models [57, 58, 69] to explicitly capture shortcut-related features, thereby encouraging the main model to ignore them.

Post-training approaches have also been explored, typically in a two-stage framework. These methods often leverage biased models to identify potential shortcuts [67, 74, 77], employ ensembles of bias models [3, 25, 29, 41], or focus on misclassified examples using contrastive learning techniques [81]. In addition, optimization-based methods modify the training objective to improve robustness to spurious correlations [33, 34, 54]. Other approaches fine-tune biased models either entirely [65] or only at the final layer [11, 22, 26, 48], often using feature reweighting strategies. Recent work has also begun exploring bias mitigation in foundation models, particularly in vision-language systems [67, 75]. Our mitigation strategy draws inspiration from these directions. Our post-hoc contrastive separation term is similar to prior work [1, 81], but unlike prior methods, our approach specifically targets *metadata traces* in vision encoders.

7 Discussion and conclusions

We study why vision encoders become sensitive to image metadata and show that this sensitivity is the result of shortcut learning, emerging when metadata correlates with the semantic supervision used during pretraining. This accounts for its presence under categorical supervision on ImageNet and, as our controlled LAION experiments indicate, under caption supervision in vision-language pretraining. Self-supervised methods such as the DINO family are less sensitive, since they do not optimize for semantic targets, yet not fully invariant, as metadata still serves as a shortcut for matching views of the same image. More broadly, we establish the image formation and processing pipeline itself as a source of bias, imperceptible to humans yet readily exploited by models.

Crucially, metadata sensitivity is not purely a disadvantage. The tendency to capture pixel-level signal that harms robustness is what makes an encoder a strong detector of generated images: models trained under stronger metadata-semantics correlations, and thus more sensitive, detect synthetic images noticeably better. This reframes recent observations that frozen CLIP features excel at fake-image detection as a direct consequence of metadata encoding, making sensitivity double-edged, a disadvantage for semantic generalization but an asset for forensics. In our analysis we identify further connections than detection: even images synthesized by Stable Diffusion, which never pass through a physical acquisition pipeline, seem to carry pixel-level metadata traces correlated with semantics, suggesting generators implicitly imitate the acquisition and processing signatures of their training data.

Since metadata distributions drift across cameras, capture conditions, and processing pipelines, encoders relying on such shortcuts are fragile precisely where deployment demands robustness. We show this sensitivity can be mitigated both during and after pretraining, generalizing across attributes while preserving semantic utility. Consistently, reducing it improves out-of-distribution generalization on ImageNet-C, ImageNet-R, and ImageNet-Sketch, with larger gains for the more sensitive models, suggesting metadata invariance may need to be a first-class objective when training or adapting foundation encoders.

We believe this work opens a new horizon, shifting attention from visible, semantic biases to the invisible traces left by the imaging pipeline. It remains open how strongly metadata sensitivity contaminates downstream tasks and the benchmarks used to evaluate encoders.

Acknowledgments: This work was supported by the GACR grant GA26-24228S, the Horizon MSCA-PF grant No. 101154126, the Czech Technical University in Prague grant No. SGS26/074/OHK3/1T/13, the JSPS KAKENHI No. 23H00497 and No. 22K12091. We acknowledge VSB – Technical University of Ostrava, IT4Innovations National Supercomputing Center, Czech Republic, for awarding this project (OPEN-35-7) access to the LUMI supercomputer, owned by the EuroHPC Joint Undertaking, hosted by CSC (Finland) and the LUMI consortium through the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (grant ID: 90254).

References

1. Alvi, M., Zisserman, A., Nellåker, C.: Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings. In: ECCV (2018)
2. Arjovsky, M., Bottou, L., Gulrajani, I., Lopez-Paz, D.: Invariant risk minimization. arXiv (2019)
3. Bahng, H., Chun, S., Yun, S., Choo, J., Oh, S.J.: Learning de-biased representations with biased representations. In: ICML (2020)
4. Beery, S., Van Horn, G., Perona, P.: Recognition in terra incognita. In: ECCV (2018)
5. Beyer, L., Izmailov, P., Kolesnikov, A., Caron, M., Kornblith, S., Zhai, X., Minderer, M., Tschannen, M., Alabdulmohsin, I., Pavetic, F.: FlexiViT: One model for all patch sizes. In: CVPR (2023)
6. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: ICCV (2021)
7. Cramér, H.: Mathematical methods of statistics. Princeton university press (1946)
8. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: RandAugment: Practical automated data augmentation with a reduced search space. In: CVPRW (2020)
9. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: ICCV (2015)
10. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2020)
11. Du, M., Mukherjee, S., Wang, G., Tang, R., Awadallah, A., Hu, X.: Fairness via representation neutralization. In: NeurIPS (2021)
12. Fan, L., Chen, K., Krishnan, D., Katabi, D., Isola, P., Tian, Y.: Scaling laws of synthetic images for model training ... for now. In: CVPR (2024)
13. Gaviria Rojas, W., Diamos, S., Kini, K., Kanter, D., Janapa Reddi, V., Coleman, C.: The dollar street dataset: Images representing the geographic and socioeconomic diversity of the world. NeurIPS (2022)
14. Geirhos, R., Jacobsen, J.H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., Wichmann, F.A.: Shortcut learning in deep neural networks. Nature Machine Intelligence (2020)
15. Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: ImageNet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In: ICLR (2018)
16. Grommelt, P., Weiss, L., Pfreundt, F.J., Keuper, J.: Fake or jpeg? revealing common biases in generated image detection datasets. In: ECCVW (2024)
17. Grootendorst, M.: BERTopic: Neural topic modeling with a class-based TF-IDF procedure. arXiv preprint arXiv:2203.05794 (2022)
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
19. Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al.: The many faces of robustness: A critical analysis of out-of-distribution generalization. In: CVPR (2021)
20. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. In: ICLR (2019)
21. Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., Song, D.: Natural adversarial examples. In: CVPR (2021)

22. Izmailov, P., Kirichenko, P., Gruver, N., Wilson, A.G.: On feature learning in the presence of spurious correlations. In: NeurIPS (2022)
23. Jahanian*, A., Chai*, L., Isola, P.: On the "steerability" of generative adversarial networks. In: ICLR (2020)
24. Kim, B., Kim, H., Kim, K., Kim, S., Kim, J.: Learning not to learn: Training deep neural networks with biased data. In: CVPR (2019)
25. Kim, N., Hwang, S., Ahn, S., Park, J., Kwak, S.: Learning debiased classifier with biased committee. In: NeurIPS (2022)
26. Kirichenko, P., Izmailov, P., Wilson, A.G.: Last layer re-training is sufficient for robustness to spurious correlations. In: ICLR (2022)
27. LaBonte, T., Hill, J., Zhang, X., Muthukumar, V., Kumar, A.: The group robustness is in the details: Revisiting finetuning under spurious correlations. In: NeurIPS (2024)
28. LAION: Releasing re-laion-5b: transparent iteration on laion-5b with additional safety fixes. <https://laion.ai/blog/re-laion-5b/> (2024)
29. Lee, Y., Yao, H., Finn, C.: Diversify and disambiguate: Out-of-distribution robustness via disagreement. In: ICLR (2023)
30. Li, Y., Vasconcelos, N.: Repair: Removing representation bias by dataset resampling. In: CVPR (2019)
31. Li, Z., Evtimov, I., Gordo, A., Hazirbas, C., Hassner, T., Ferrer, C.C., Xu, C., Ibrahim, M.: A whac-a-mole dilemma: Shortcuts come in multiples where mitigating one amplifies others. In: CVPR (2023)
32. Li, Z., Evtimov, I., Gordo, A., Hazirbas, C., Hassner, T., Ferrer, C.C., Xu, C., Ibrahim, M.: A whac-a-mole dilemma: Shortcuts come in multiples where mitigating one amplifies others. In: CVPR (2023)
33. Liu, E.Z., Haghighi, B., Chen, A.S., Raghunathan, A., Koh, P.W., Sagawa, S., Liang, P., Finn, C.: Just train twice: Improving group robustness without training group information. In: ICML (2021)
34. Liu, S., Zhang, X., Sekhar, N., Wu, Y., Singhal, P., Fernandez-Granda, C.: Avoiding spurious correlations via logit correction. In: ICLR (2022)
35. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV (2021)
36. Liu, Z., He, K.: A decade's battle on dataset bias: Are we there yet? In: ICLR (2025)
37. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. In: CVPR (2022)
38. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. JMLR (2008)
39. Meister, N., Zhao, D., Wang, A., Ramaswamy, V.V., Fong, R., Russakovsky, O.: Gender Artifacts in Visual Datasets. In: ICCV (2023)
40. Nagarajan, V., Andreassen, A., Neyshabur, B.: Understanding the failure modes of out-of-distribution generalization. In: ICLR (2021)
41. Nam, J., Cha, H., Ahn, S., Lee, J., Shin, J.: Learning from failure: Training debiased classifier from biased classifier. In: NeurIPS (2020)
42. Nam, J., Kim, J., Lee, J., Shin, J.: Spread spurious attribute: Improving worst-group accuracy with spurious attribute estimation. In: ICLR (2022)
43. Neuhaus, Y., Augustin, M., Boreiko, V., Hein, M.: Spurious features everywhere - large-scale detection of harmful spurious features in ImageNet. In: ICCV (2023)
44. Ojha, U., Li, Y., Lee, Y.J.: Towards universal fake image detectors that generalize across generative models. In: CVPR (2023)

45. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: DINOv2: Learning robust visual features without supervision. TMLR (2024)
46. Parashar, S., Lin, Z., Liu, T., Dong, X., Li, Y., Ramanan, D., Caverlee, J., Kong, S.: The neglected tails in vision-language models. In: CVPR (2024)
47. Puli, A.M., Joshi, N., Wald, Y., He, H., Ranganath, R.: Nuisances via negativa: Adjusting for spurious correlations via data augmentation. TMLR (2024)
48. Qiu, S., Potapczynski, A., Izmailov, P., Wilson, A.G.: Simple and fast group robustness by automatic feature reweighting. In: ICML (2023)
49. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
50. Ramos, R., Stojnić, V., Kordopatis-Zilos, G., Nakashima, Y., Tolias, G., Garcia, N.: Processing and acquisition traces in visual encoders: What does clip know about your camera? In: ICCV (2025)
51. Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do ImageNet classifiers generalize to ImageNet? In: ICML (2019)
52. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)
53. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M.S., Berg, A.C., Fei-Fei, L.: ImageNet large scale visual recognition challenge. IJCV (2015)
54. Sagawa, S., Koh, P.W., Hashimoto, T.B., Liang, P.: Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. In: ICLR (2020)
55. Sagawa, S., Raghunathan, A., Koh, P.W., Liang, P.: An investigation of why over-parameterization exacerbates spurious correlations. In: ICML (2020)
56. Sariyildiz, M.B., Alahari, K., Larus, D., Kalantidis, Y.: Fake it till you make it: Learning transferable representations from synthetic ImageNet clones. In: CVPR (2023)
57. Sarriidis, I., Koutlis, C., Papadopoulos, S., Diou, C.: Flac: Fairness-aware representation learning by suppressing attribute-class associations. PAMI (2024)
58. Sarriidis, I., Koutlis, C., Papadopoulos, S., Diou, C.: MAVias: Mitigate any visual bias. In: ICCV (2025)
59. Sha, Z., Li, Z., Yu, N., Zhang, Y.: De-fake: Detection and attribution of fake images generated by text-to-image generation models. In: SIGSAC (2023)
60. Shah, H., Tamuly, K., Raghunathan, A., Jain, P., Netrapalli, P.: The pitfalls of simplicity bias in neural networks. NeurIPS (2020)
61. Shirali, A., Hardt, M.: What makes ImageNet look unlike LAION. TMLR (2025)
62. Singla, S., Feizi, S.: Salient ImageNet: How to discover spurious features in deep learning? In: ICLR (2022)
63. Smeu, S., Oneata, E., Oneata, D.: Declip: Decoding clip representations for deep-fake localization. In: WACV (2025)
64. Suhail, P., Sethi, A.: Shortcut Learning Susceptibility in Vision Classifiers. In: Workshop on Spurious Correlation and Shortcut Learning: Foundations and Solutions (2025)
65. Taghanaki, S.A., Khani, A., Khani, F., Gholami, A., Tran, L., Mahdavi-Amiri, A., Hamarneh, G.: Masktune: mitigating spurious correlations by forcing to explore. In: NeurIPS (2022)
66. Torralba, A., Efros, A.A.: Unbiased look at dataset bias. In: CVPR (2011)

67. Varma, M., Delbrouck, J.B., Chen, Z., Chaudhari, A., Langlotz, C.: Ravi: Discovering and mitigating spurious correlations in fine-tuned vision-language models. In: NeurIPS (2024)
68. Wang, H., Ge, S., Lipton, Z.C., Xing, E.P.: Learning robust global representations by penalizing local predictive power. In: NeurIPS (2019)
69. Wang, H., He, Z., Lipton, Z.C., Xing, E.P.: Learning robust representations by projecting superficial statistics out. In: ICLR (2019)
70. Wang, Q., Lin, Y., Chen, Y., Schmidt, L., Han, B., Zhang, T.: A sober look at the robustness of clips to spurious features. NeurIPS (2024)
71. Wang, S., Veldhuis, R.N.J., Brune, C., Strisciuglio, N.: What do neural networks learn in image classification? A frequency shortcut perspective. In: ICCV (2023)
72. Wang, S., Veldhuis, R.N.J., Strisciuglio, N.: Do imagenet-trained models learn shortcuts? the impact of frequency shortcuts on generalization. In: CVPR (2025)
73. Wu, S., Yuksekgonul, M., Zhang, L., Zou, J.: Discover and cure: Concept-aware mitigation of spurious correlation. In: ICML (2023)
74. Yang, Y., Gan, E., Dziugaite, G.K., Mirzasoleiman, B.: Identifying spurious biases early in training through the lens of simplicity bias. In: AISTATS (2024)
75. Yang, Y., Nushi, B., Palangi, H., Mirzasoleiman, B.: Mitigating spurious correlations in multi-modal models during fine-tuning. In: ICML (2023)
76. Yao, H., Wang, Y., Li, S., Zhang, L., Liang, W., Zou, J., Finn, C.: Improving out-of-distribution robustness via selective augmentation. In: ICML (2022)
77. Yenamandra, S., Ramesh, P., Prabhu, V., Hoffman, J.: Facts: First amplify correlations and then slice to discover bias. In: ICCV (2023)
78. Zech, J.R., Badgeley, M.A., Liu, M., Costa, A.B., Titano, J.J., Oermann, E.K.: Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. PLoS Medicine **15**(11), e1002683 (2018)
79. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: ECCV (2014)
80. Zhai, X., Mustafa, B., Kolesnikov, A., Beyer, L.: Sigmoid loss for language image pre-training. In: ICCV (2023)
81. Zhang, M., Sohoni, N.S., Zhang, H.R., Finn, C., Ré, C.: Correct-n-contrast: A contrastive approach for improving robustness to spurious correlations. In: NeurIPS (2022)