

Inverting Adversarially Robust Networks for Image Synthesis

Renan A. Rojas-Gomez¹, Raymond A. Yeh²,
 Minh N. Do¹, and Anh Nguyen³

¹ University of Illinois at Urbana-Champaign

² Purdue University

³ Auburn University

{renanar2,minhdo}@illinois.edu rayyeh@purdue.edu anh.ng8@gmail.com

Abstract. Despite unconditional feature inversion being the foundation of many image synthesis applications, training an inverter demands a high computational budget, large decoding capacity and imposing conditions such as autoregressive priors. To address these limitations, we propose the use of adversarially robust representations as a perceptual primitive for feature inversion. We train an adversarially robust encoder to extract disentangled and perceptually-aligned image representations, making them easily invertible. By training a simple generator with the mirror architecture of the encoder, we achieve superior reconstruction quality and generalization over standard models. Based on this, we propose an adversarially robust autoencoder and demonstrate its improved performance on style transfer, image denoising and anomaly detection tasks. Compared to recent ImageNet feature inversion methods, our model attains improved performance with significantly less complexity.⁴

1 Introduction

Deep classifiers trained on large-scale datasets extract meaningful high-level features of natural images, making them an essential tool for manipulation tasks such as style transfer [1,2,3], image inpainting [4,5], image composition [6,7], among others [8,9,10]. State-of-the-art image manipulation techniques use a decoder [5,10], *i.e.*, an *image generator*, to create natural images from high-level features. Extensive work has explored how to train image generators, leading to models with photorealistic results [11]. Moreover, by learning how to invert deep features, image generators enable impressive synthesis use cases such as anomaly detection [12,13] and neural network visualization [7,14,15,16].

Inverting ImageNet features is a challenging task that often requires the generator to be more complex than the encoder [17,18,19,5], incurring in a high computational cost. Donahue et al. [17] explained this shortcoming by the fact that the encoder bottleneck learns entangled representations that are hard to invert. An alternative state-of-the-art technique for inverting ImageNet features

⁴ Code available at https://github.com/renanrojasg/adv_robust_autoencoder

requires, in addition to the encoder and decoder CNNs, *an extra* autoregressive model and vector quantization [20,21] or a separate invertible network [16].

In this paper, we propose a **novel mechanism for training effective ImageNet autoencoders that do not require extra decoding layers or networks besides the encoder and its mirror decoder**. Specifically, we adopt a pre-trained classifier as encoder and train an image generator to invert its features, yielding an autoencoder for real data. Unlike existing works that use feature extractors trained on natural images, we train the encoder on adversarial examples [22]. This fundamental difference equips our *adversarially robust* (AR) autoencoder with representations that are perceptually-aligned with human vision [23,9], resulting in favorable inversion properties.

To show the advantages of learning how to invert AR features, our generator corresponds to the *mirror* architecture of the encoder, without additional decoding layers [17,6] or extra components [24,20,16,21,25]. To the best of our knowledge, we are the first to show the benefits of training an autoencoder on both adversarial and real images. Our main findings are as follows:

- A generator trained to invert AR features has a substantially higher reconstruction quality than those trained on standard features. Our method generalizes to different models (AlexNet [26], VGG-16 [27], and ResNet [28]) and datasets (CIFAR-10 [29] and ImageNet [30])(Sec. 5.1).
- Our proposed AR autoencoder is remarkably robust to resolution changes, as shown on natural and upscaled high-resolution images (Fig. 8). Experiments on DIV2K [31] show it accurately reconstructs high-resolution images without any finetuning, despite being trained on low-resolution images (Sec. 5.3).
- Our generator outperforms state-of-the-art inversion methods based on iterative optimization techniques [23] in terms of PSNR, SSIM, and LPIPS [10]. It also attains comparable accuracy to the well-established DeepSiM model [32] with a much lower model complexity (Sec. 5.4).
- Our AR model outperforms standard baselines on three downstream tasks: style transfer [2], image denoising [5] (Sec. 6) and anomaly detection [12]. The latter is covered in detail in the Appendix (Sec. A1).

2 Related Work

Inverting Neural Networks. Prior work exploring deep feature inversion using optimization approaches are either limited to per-pixel priors or require multiple steps to converge and are sensitive to initialization [36,61,23,9]. Instead, we propose to map contracted features to images via a generator, following the work by Dosovitskiy et al. [19] and similar synthesis techniques [6,5,7]. By combining natural priors and AR features, we get a significant reconstruction improvement with much less trainable parameters.

Our results are consistent to prior findings on AR features being more invertible via optimization [23] and more useful for transfer learning [52]. As part of our contribution, we complement these by showing that (i) learning a map from the AR feature space to the image domain largely outperforms the original

optimization approach, (ii) such an improvement generalizes to models of different complexity, and (iii) inverting AR features shows remarkable robustness to scale changes. We also show AR encoders with higher robustness can be more easily decoded, revealing potential security issues [62].

Regularized Autoencoders. Prior work requiring data augmentation to train generative and autoencoding models often requires learning an invertible transformation that maps augmented samples back to real data [63]. Instead, our approach can be seen as a novel way to regularize bottleneck features, providing an alternative to contractive, variational and sparse autoencoders [11,54,64].

3 Preliminaries

Our model exploits AR representations to reconstruct high-quality images, which is related to the feature inversion framework. Specifically, we explore AR features as a strong prior to obtain photorealism. For a clear understanding of our proposal, we review fundamental concepts of feature inversion and AR training.

Feature Inversion. Consider a target image $x_0 \in \mathbb{R}^{W \times H \times C}$ and its contracted representation $f_0 \triangleq F_\theta(x_0) \in \mathbb{R}^{W' \times H' \times C'}$. Here, F_θ denotes the target model, *e.g.* AlexNet, with parameters $\theta \in \mathbb{R}^T$ and $W'H'C' \ll WHC$. Features extracted by F_θ encapsulate rich input information that can either be used for the task it was trained on, transferred to a related domain [33] or used for applications such as image enhancement and manipulation [34,1].

An effective way to leverage these representations is by training a second model, a generator, to map them to the pixel domain. This way, deep features can be manipulated and transformed into images [32,19]. Also, since deep features preserve partial input information, inverting them elucidates what kind of attributes they encode. Based on these, *feature inversion* [35,36,32] has been extensively studied for visualization and understanding purposes as well as for synthesis and manipulation tasks. Typically, feature inversion is formulated as an optimization problem:

$$\hat{x} = \arg \min_x \mathcal{F}(F_\theta(x), f_0) + \lambda \mathcal{R}(x), \quad (1)$$

where \hat{x} is the estimated image and $\mathcal{F}(F_\theta(x), f_0)$ the fidelity term between estimated and target representations, $F_\theta(x)$ and f_0 respectively. $\mathcal{R}(x)$ denotes the regularization term imposing *a priori* constraints in the pixel domain and $\lambda \in \mathbb{R}_{++}$ balances between fidelity and regularization terms.

Adversarial Robustness. Adversarial training adds perturbations to the input data and lets the network learn how to classify in the presence of such adversarial attacks [37,22,38]. Consider the image classification task with annotated dataset \mathcal{K} . Let an annotated pair correspond to image $x \in \mathbb{R}^{W \times H \times C}$ and its one-hot encoded label $y \in \{0, 1\}^{|\mathcal{C}|}$, where \mathcal{C} is the set of possible classes. From the definition by Madry et al. [22], a perturbed input is denoted by $x' = x + \delta$, where x' is the perturbed sample and δ the perturbation. Let the set of perturbations be bounded by the ℓ_p ball for $p \in \{2, \infty\}$, $\mathcal{S} : \{\delta, \|\delta\|_p \leq \varepsilon\}$. Then, the AR training corresponds to an optimization problem:




Fig. 1: By training it to invert adversarially robust features, our proposed autoencoder obtains better reconstructions than models trained on standard features.

$$\tilde{\theta} = \arg \min_{\theta} \mathbb{E}_{(x,y) \sim \mathcal{K}} \left[\max_{\delta \in \mathcal{S}} \mathcal{L}_{x',y}(\theta) \right], \quad (2)$$

where $\tilde{\theta} \in \mathbb{R}^T$ are the optimal weights and $\mathcal{L}_{x',y}(\theta)$ the negative log-likelihood. The goal is to minimize $\mathcal{L}_{x',y}(\theta)$ in the presence of the worst possible adversary.

4 Proposed Method

4.1 Adversarially Robust Autoencoder

We propose an autoencoder architecture (Fig. 1) to extract bottleneck AR features of arbitrary input images, manipulate them for a given synthesis task, and map the results back to images. We denote the AR feature extractor as $F_{\tilde{\theta}}$, where $\tilde{\theta}$ are the AR model weights, as explained in Sec. 3. Robust features are transformed into images using a CNN-based generator denoted as $G_{\tilde{\phi}}$. Here, $\tilde{\phi}$ are the generator weights learned by inverting AR features.

Following prior works [39,19], we use AlexNet as the encoder and extract AR features from its conv5 layer. We also explore more complex encoders from the VGG and ResNet families and evaluate their improvement over standard encoders (See Sec. A4.1 for architecture details).

4.2 Image Decoder: Optimization Criteria

Given a pre-trained AR encoder $F_{\tilde{\theta}}$, the generator $G_{\tilde{\phi}}$ is trained using ℓ_1 pixel, ℓ_2 feature and GAN losses, where the feature loss matches AR representations, known to be *perceptually aligned* [23].

In more detail, we denote $\hat{x} = G_{\tilde{\phi}}(f)$ to be the reconstruction of image x , where $f = F_{\tilde{\theta}}(x)$ are its AR features. Training the generator with fixed encoder's weights $\tilde{\theta}$ corresponds to the following optimization problem:

$$\tilde{\phi} = \arg \min_{\phi} \lambda_{\text{pix}} \mathcal{L}_{\text{pix}}(\phi) + \lambda_{\text{feat}} \mathcal{L}_{\text{feat}}(\phi, \tilde{\theta}) + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}}(\phi, \psi), \quad (3)$$

$$\mathcal{L}_{\text{pix}}(\phi) \triangleq \mathbb{E}_{x \sim \tilde{\mathcal{K}}} \|x - G_{\phi}(f)\|_1, \quad (4)$$

$$\mathcal{L}_{\text{feat}}(\phi, \tilde{\theta}) \triangleq \mathbb{E}_{x \sim \tilde{\mathcal{K}}} \|f - F_{\tilde{\theta}} \circ G_{\phi}(f)\|_2^2, \quad (5)$$

$$\mathcal{L}_{\text{adv}}(\phi, \psi) \triangleq \mathbb{E}_{x \sim \tilde{\mathcal{K}}} [-\log D_{\psi} \circ G_{\phi}(f)], \quad (6)$$

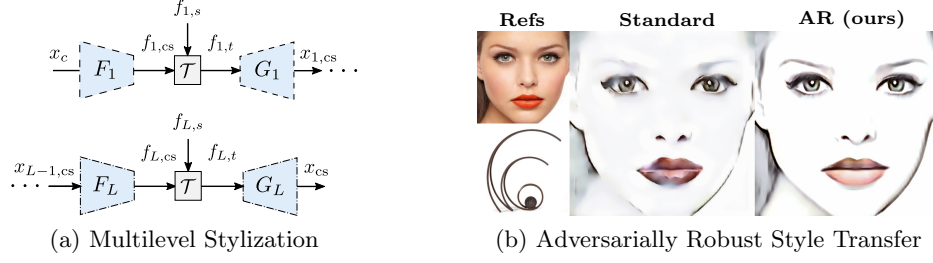


Fig. 2: Example-based Style Transfer using adversarially robust features.

where $\lambda_{\text{pix}}, \lambda_{\text{feat}}, \lambda_{\text{adv}} \in \mathbb{R}_{++}$ are hyperparameters, $D_\psi : \mathbb{R}^{W \times H \times C} \mapsto [0, 1]$ denotes the discriminator with weights ψ and predicts the probability of an image being real. The pixel loss $\mathcal{L}_{\text{pix}}(\phi)$ is the ℓ_1 distance between prediction $G_\phi(f)$ and target x . The feature loss $\mathcal{L}_{\text{feat}}(\phi, \theta)$ is the ℓ_2 distance between the AR features of prediction and target. The adversarial loss $\mathcal{L}_{\text{adv}}(\phi, \psi)$ maximizes the discriminator score of predictions, *i.e.*, it increases the chance the discriminator classifies them as real. On the other hand, the discriminator weights are trained via the cross-entropy loss, *i.e.*,

$$\min_{\psi} \mathcal{L}_{\text{disc}}(\phi, \psi) \triangleq \mathbb{E}_{x \sim \tilde{\mathcal{K}}} [-\log D_\psi(x) - \log(1 - D_\psi \circ G_\phi(f))]. \quad (7)$$

This discriminative loss $\mathcal{L}_{\text{disc}}(\phi, \psi)$ guides D_ψ to maximize the score of real images and minimize the score of reconstructed (fake) images. Similar to traditional GAN algorithms, we alternate between the generator and discriminator training to reach the equilibrium point.

4.3 Applications

The trained AR autoencoder can be used to improve the performance of tasks such as style transfer [2], image denoising [40], and anomaly detection [12]. In what follows, we describe the use of our model on style transfer and image denoising. The task of anomaly detection is covered in the Appendix (Sec. A1).

Example-based Style Transfer. Style transfer [1] aligns deep features to impose perceptual properties of a style image x_s over semantic properties of a content image x_c . This is done by matching the content and style distributions in the latent space of a pre-trained encoder to then transform the resulting features back into images. We adopt the Universal Style Transfer framework [2] to show the benefits of using our AR model for stylization (Fig. 2).

We train three AR AlexNet autoencoders $\{F_{l,\tilde{\theta}}, G_{l,\tilde{\phi}}\}_{l=1}^{L=3}$ and use them to sequentially align features at each scale. $F_{1,\tilde{\theta}}, F_{2,\tilde{\theta}}$ and $F_{3,\tilde{\theta}}$ extract AR conv5, conv2 and conv1 features, respectively. First, style features $f_{l,s} = F_{l,\tilde{\theta}}(x_s)$ are extracted at each stage. We then use the content image as initialization for the stylized output $x_{0,cs} \triangleq x_c$ and extract its conv5 features $f_{1,cs} = F_{1,\tilde{\theta}}(x_{0,cs})$.

At stage $l = 1$, the style distribution is imposed over the content features by using the whitening and coloring transform [2,41] denoted by \mathcal{T} . The resulting representation $f_{1,t} = \mathcal{T}(f_{1,s}, f_{1,cs})$ is characterized by the first and second




Fig. 3: Image denoising using our adversarially robust autoencoder.

Table 1: AlexNet feature inversion on ImageNet. Under distinct training losses, inverting AR features via our proposed generator is consistently more accurate than inverting standard features.

Losses	Model	PSNR (dB) \uparrow	SSIM \uparrow	LPIPS \downarrow
Pixel	Standard	17.562 ± 2.564	0.454 ± 0.167	0.624 ± 0.099
	AR (ours)	19.904 ± 2.892	0.505 ± 0.169	0.596 ± 0.104
Pixel, Feature	Standard	14.462 ± 1.884	0.103 ± 0.044	0.713 ± 0.046
	AR (ours)	17.182 ± 2.661	0.284 ± 0.111	0.601 ± 0.034
Pixel, Feature, GAN	Standard	15.057 ± 2.392	0.307 ± 0.158	0.547 ± 0.055
	AR (ours)	17.227 ± 2.725	0.358 ± 0.163	0.567 ± 0.056

moments of the style distribution. An intermediate stylized image $x_{1,cs}$ incorporating the style at the first scale is then generated as $x_{1,cs} = G_{1,\tilde{\phi}}(f_{1,t})$.

The process is repeated for $l \in \{2, 3\}$ to incorporate the style at finer resolutions, resulting in the final stylized image $x_{cs} = x_{3,cs}$.

Image Denoising. Motivated by denoising autoencoders (DAE) [40] where meaningful features are extracted from distorted instances, we leverage AR features for image enhancement tasks. Similarly to deep denoising models [42], we incorporate skip connections in our pre-trained AR AlexNet autoencoder to extract features at different scales, complementing the information distilled at the encoder bottleneck (Fig. 3). Skip connections correspond to Wavelet Pooling [3], replacing pooling and upsampling layers by analysis and synthesis Haar wavelet operators, respectively. Our skip-connected model is denoted by $\{F_{s,\tilde{\theta}}, G_{s,\tilde{\phi}}\}$.

Similarly to real quantization scenarios [43,44,45], we assume images are corrupted by clipped additive Gaussian noise. A noisy image is denoted by $b = \rho(x + \eta) \in \mathbb{R}^{W \times H \times C}$, where $\eta \sim \mathcal{N}(0, \sigma)$ is the additive white Gaussian noise term and $\rho(x) = \max[0, \min(1, x)]$ a pointwise operator restricting the range between 0 and 1. Denoised images are denoted by $\hat{x}_s = G_{s,\tilde{\phi}} \circ F_{s,\tilde{\theta}}(b)$.

$G_{s,\tilde{\phi}}$ is trained to recover an image x from the features of its corrupted version $F_{s,\tilde{\theta}}(b)$. The training process uses the optimization criteria described in Sec. 4.2.

5 Experiments on Feature Inversion

We begin analyzing the reconstruction accuracy achieved by inverting features from different classifiers and empirically show that learning how to invert AR features via our proposed generator improves over standard feature inversion. Refer to Sec. A2 and Sec. A4 for additional inversion results and training details.

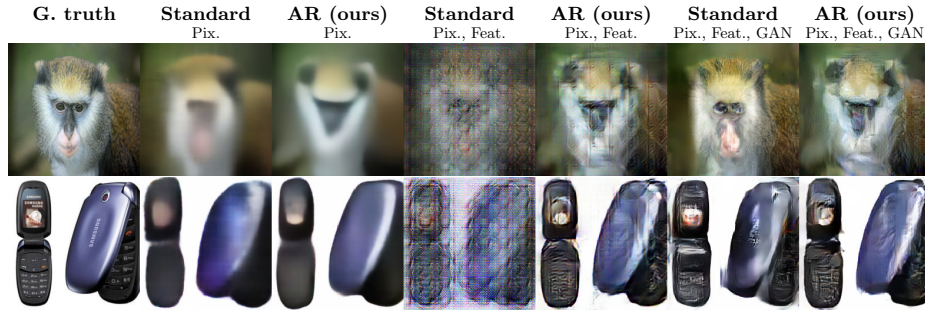


Fig. 4: AlexNet feature inversion on ImageNet. Conv5 features are inverted using our proposed generator under three different training criteria. Reconstructions from AR features are more faithful to the ground-truth image.

5.1 Reconstruction Accuracy of AR Autoencoders

Inverting AlexNet features. Standard and AR AlexNet autoencoders are trained as described in Sec. 4.1 on ImageNet for comparison purposes. The AR AlexNet classifier is trained via ℓ_2 -PGD attacks [22] of radius $\varepsilon = \frac{3}{255}$ and 7 steps of size 0.5. Training is performed using 90 epochs via SGD with a learning rate of 0.1 reduced 10 times every 30 epochs. On the other hand, the standard AlexNet classifier is trained on natural images via cross-entropy (CE) loss with the same SGD setup as in the AR case.

Next, generators are trained using pixel, feature and GAN losses to invert AlexNet conv5 features (size $6 \times 6 \times 256$). Both AR and standard models use the same generator architecture, which corresponds to the mirror network of the encoder. We deliberately use a simple architecture to highlight the reconstruction improvement is due to inverting AR features and not the generator capacity. We also train generators using (i) pixel and (ii) pixel and feature losses to ablate their effect. Reconstruction quality is evaluated using PSNR, SSIM and LPIPS.

Under all three loss combinations, reconstructions from AR AlexNet features obtain better PSNR and SSIM than their standard counterparts (Tab. 1). Specifically, inverting AR AlexNet features gives an average PSNR improvement of over 2 dB in all three cases. LPIPS scores also improve, except when using pixel, feature and GAN losses. Nevertheless, inverting AR features obtain a strong PSNR and SSIM improvement in this case as well. Qualitatively, inverting AR features better preserves the natural appearance in all cases, reducing the checkerboard effect and retaining sharp edges (Fig. 4).

Inverting VGG features. We extend the analysis to VGG-16 trained on ImageNet-143 and evaluate the reconstruction improvement achieved by inverting its AR features. We use the AR pre-trained classifier from the recent work by Liu et al. [46] trained using ℓ_∞ -PGD attacks of radius $\varepsilon = 0.01$ and 10 steps of size $\frac{1}{50}$. Training is performed using 80 epochs via SGD with a learning rate of 0.1 reduced 10 times every 30, 20, 20 and 10 epochs. On the other hand, its standard version is trained on natural images via CE loss with the same SGD setup as in the AR case.

Table 2: AR VGG-16 [46] feature inversion on ImageNet. Training our generator via pixel and feature losses, reconstruction largely improves by inverting AR representations.

	Standard Model	AR Model (ours)
Standard Accuracy	65.0	48.7
ℓ_∞ PGD Accuracy	0	23.0
PSNR (dB) \uparrow	18.35 ± 2.471	21.063 ± 3.132
SSIM \uparrow	0.466 ± 0.2	0.538 ± 0.165
LPIPS \downarrow	0.327 ± 0.101	0.225 ± 0.057



Fig. 5: AR VGG-16 reconstruction on ImageNet.

Generators are trained on pixel and feature losses to invert VGG-16 conv5.1 features (size $14 \times 14 \times 512$). Similarly to the AlexNet analysis, generators inverting both standard and AR features correspond to the mirror network of the encoder. We evaluate the reconstruction accuracy of both models and report their level of adversarial robustness (Tab. 2 and Fig. 5).

Quantitatively, reconstructions from AR VGG-16 features are more accurate than those of standard features in PSNR, SSIM and LPIPS by a large margin. Specifically, inverting AR VGG-16 features gives an average PSNR improvement of 2.7 dB. Qualitatively, reconstructions from AR VGG-16 features are more similar to the original images, reducing artifacts and preserving object boundaries.

Furthermore, the reconstruction accuracy attained by the AR VGG-16 autoencoder improves over that of the AR AlexNet model. This suggests that the benefits of inverting AR features are not constrained to shallow models such as AlexNet, but generalize to models with larger capacity.

Inverting ResNet features. To analyze the effect of inverting AR features from classifiers trained on different datasets, we evaluate the reconstruction accuracy obtained by inverting WideResNet-28-10 trained on CIFAR-10. We use the AR pre-trained classifier from the recent work by Zhang et al. [47]. This model obtains State-of-the-art AR classification accuracy via a novel weighted adversarial training regime. Specifically, the model is adversarially trained via PGD by ranking the importance of each sample based on how close it is to the decision boundary (how *attackable* the sample is).

AR training is performed using ℓ_∞ attacks of radius $\varepsilon = \frac{8}{255}$ and 10 steps of size $\frac{2}{255}$. Classification training is performed using 100 epochs (with a burn-in period of 30 epochs) via SGD with a learning rate of 0.1 reduced 10 times every 30 epochs. On the other hand, its standard version is trained on natural images via CE loss using the same SGD setup as in the AR case.

Generators for standard and AR WideResNet-28-10 models are trained to invert features from its 3rd residual block (size $8 \times 8 \times 640$) via pixel and feature losses. Similarly to our previous analysis, both generators correspond to the mirror architecture of the encoder. We evaluate their reconstruction via PSNR, SSIM and LPIPS, and their robustness via AutoAttack [48] (Tab. 3 and Fig. 6).

Similarly to previous scenarios, inverting WideResNet-28-10 AR features shows a large improvement over standard ones in all metrics. Specifically, invert-

Table 3: AR WideResNet-28-10 [47] feature inversion on CIFAR-10. Inverting AR features via our generator trained on pixel and feature losses significantly improves reconstruction.

	Standard Model	AR Model (ours)
Standard Accuracy	93.8	89.36
AutoAttack [48]	0	59.64
PSNR (dB) \uparrow	17.38 ± 2.039	22.14 ± 1.626
SSIM \uparrow	0.59 ± 0.1	0.81 ± 0.067
LPIPS \downarrow	0.2547 ± 0.055	0.2318 ± 0.0833




Fig. 6: AR WideResNet-28-10 reconstruction on CIFAR-10.

Table 4: Reconstruction vs. Robustness. Experiments on ImageNet show that learning to invert AlexNet features with different AR levels can significantly improve the reconstruction accuracy.

	ℓ_2 PGD Attack (ε)				
	0	0.5	2	3	4
Standard Accuracy	53.69	49.9	43.8	39.83	36.31
AutoAttack [48]	8.19 ($\varepsilon = 0.5$)	48.0 ($\varepsilon = 0.5$)	28.0 ($\varepsilon = 2$)	22.27 ($\varepsilon = 3$)	14.9 ($\varepsilon = 4$)
PSNR (dB) \uparrow	13.12	14.41	15.5	15.53	15.61
SSIM \uparrow	0.20	0.26	0.3	0.26	0.25
LPIPS \downarrow	0.657	0.625	0.614	0.629	0.644

ing AR features increases PSNR in 4.8 dB on average over standard features. Visually, the AR WideResNet-28-10 autoencoder reduces bogus components and preserves object contours on CIFAR-10 test samples.

Overall, results enforce our claim that the **benefits of inverting AR features extend to different models, datasets and training strategies.**

5.2 Robustness Level vs. Reconstruction Accuracy

We complement the reconstruction analysis by exploring the relation between adversarial robustness and inversion quality. We train five AlexNet classifiers on ImageNet, one on natural images (standard) and four via ℓ_2 -PGD attacks with $\varepsilon \in \{0.5, 2, 3, 4\}/255$. All other training parameters are identical across models.

For each classifier, an image generator is trained on an ImageNet subset via pixel, feature and GAN losses to invert conv5 features. Similar to Sec. 5.1, all five generators correspond to the mirror network of the encoder. To reliably measure the impact of adversarial robustness, reconstruction accuracy is evaluated in terms of PSNR, SSIM and LPIPS. We also report the effective robustness level achieved by each model via AutoAttack (Tab. 4).

Results show LPIPS and SSIM improve almost monotonically until a maximum value is reached at $\varepsilon = 2$, while PSNR keeps increasing. This implies that just by changing ε from 0.5 to 4 while keeping the exact same architecture and training regime, a reconstruction improvement of 1.2 dB PSNR is obtained.

Based on this, we use an AR AlexNet model trained with $\varepsilon = 3$ in our experiments, which gives the best tradeoff between PSNR, SSIM and LPIPS. Overall, our analysis suggests that, while all four AR models outperform the

Table 5: Reconstructing upscaled ImageNet samples. Images upscaled by a factor L are reconstructed from their standard and AR AlexNet features. In contrast to the degraded standard reconstructions, AR reconstructions show an outstanding accuracy that improves for large scaling factors.

L	Standard AlexNet			Robust AlexNet		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
1	15.057	0.3067	0.5473	17.2273	0.3580	0.5665
4	15.4258	0.4655	0.4136	22.575	0.5892	0.4012
7	13.8922	0.4852	0.4587	23.5778	0.6588	0.3898
10	13.1013	0.4969	0.486	23.9566	0.7244	0.3892

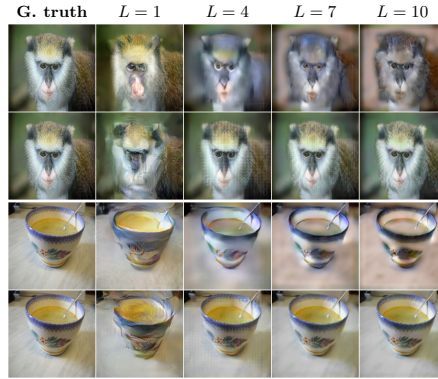


Fig. 7: Upscaled ImageNet samples reconstructed from their standard (top row) and AR (bottom row) features.

Table 6: High-resolution images inverted using our AR AlexNet model (trained on low resolution images) show improved quality over standard inversions.

Encoder	PSNR (dB)↑	SSIM↑	LPIPS↓
Standard	14.266± 1.9015	0.3874± 0.151	0.5729± 0.0465
AR (ours)	18.3606± 2.6012	0.4388± 0.1508	0.5673± 0.0337

inversion accuracy of the standard model, the reconstruction improvement is not proportional to the robustness level. Instead, it is maximized at a particular level. Please refer to Sec. A2.3 for additional robustness level vs. reconstruction accuracy experiments on ResNet-18 pointing to the same conclusion.

5.3 Reconstructing Images at Unseen Resolutions

Unlike extracting shift-invariant representations, image scaling is difficult to handle for standard CNN-based models [49,50]. Following previous work suggesting AR features are more generic and transferable than standard ones [51,52], we test whether our proposed AR autoencoder generalizes better to scale changes. We explore this property and show that our model trained on low-resolution samples improves reconstruction of images at unseen scales without any fine-tuning.

Scenario 1: Reconstructing Upscaled Images. Upscaled ImageNet samples are reconstructed from their AR AlexNet conv5 representations. For a fair comparison across scales, each image is normalized to 224×224 px. and then enlarged by an integer factor $L > 1$. Experiments show a higher accuracy obtained from AR features in terms of PSNR, SSIM and LPIPS (Tab. 5). All metrics improve almost monotonically with L . In contrast, accuracy using standard features degrades with L . Inversion from AR features show almost perfect reconstruction for large scales, while those of standard features show severe distortions (Fig. 7).

Scenario 2: Reconstructing High-Resolution Images. Standard and AR feature inversion is performed on the DIVERse 2K resolution dataset (DIV2K)



Fig. 8: At a resolution of 2040×1536 px., 10 times larger than training samples, standard reconstructions on DIV2K show color and structure degradation. In contrast, reconstructions from our AR model do not suffer from distortions.

[31], containing objects at multiple scales. AR feature reconstructions show a significant PSNR, SSIM and LPIPS improvement over standard ones, despite not being explicitly trained to handle such large-scale objects (Tab. 6).

Qualitatively, reconstructions from AR AlexNet features preserve sharp edges, reduces color degradation and diminishes checkerboard effects induced by standard inversion (Fig. 8). Thus, for unseen scales and without finetuning, AR features better preserve structure without penalizing the perceptual similarity.

5.4 Comparison against State-of-the-Art Inversion Techniques

The inversion accuracy of our AR autoencoder is compared against two alternative techniques: Optimization-based robust representation inversion (RI) [23] and DeePSiM [32]. For a fair comparison, all methods reconstruct images from AlexNet features. We begin by highlighting the differences between them.

While RI is a model-based approach that searches in the pixel domain for an image that matches a set of target AR features, we use a CNN-based generator trained on a combination of natural-image priors (Sec. 4.2). On the other hand, while DeePSiM is also a CNN-based technique trained under multiple priors, its generator has approximately 63% more trainable parameters than ours (Tab. 7).

Experimental Setup. All inversion methods are evaluated on ImageNet. Our standard and AR models are trained using pixel, feature and GAN losses using the training setup described in Sec. A4. DeePSiM is evaluated using its official Caffe implementation without any changes. RI is evaluated using its official PyTorch implementation modified to invert conv5 AR features. Input samples are rescaled to 224×224 px. (227×227 px. for DeePSiM).

Results. Our AR AlexNet autoencoder obtains the best accuracy in terms of PSNR and the second best in terms of SSIM (Tab. 7). While it outperforms its standard version in PSNR and SSIM, it gets a marginally worse LPIPS. Moreover, our AR model outperforms RI in all metrics. Also, despite DeePSiM having more layers and using larger inputs, our model achieves a large PSNR improvement over it. Results highlight the improvement obtained by inverting AR features and how this fundamental change allows competitive reconstruction quality using three times less trainable parameters.

Table 7: Comparison against state-of-the-art inversion techniques. By inverting AR features, our autoencoder outperforms the optimization-based RI method by a large margin. Despite having 63% less parameters, we also obtain favorable results against DeepSiM, showing a significant PSNR improvement.

Algorithm	Encoder	Trainable Pars.	PSNR (dB) \uparrow	SSIM \uparrow	LPIPS \downarrow
RI [23]	AR AlexNet	—	16.724 ± 2.434	0.181 ± 0.071	0.63 ± 0.04
Standard Autoencoder	Standard AlexNet	4,696,026	15.057 ± 2.392	0.307 ± 0.158	0.547 ± 0.055
AR Autoencoder (ours)	AR AlexNet	4,696,026	17.227 ± 2.725	0.358 ± 0.163	0.567 ± 0.056
DeepSiM [19]	Standard CaffeNet	12,702,307	15.321 ± 2.011	0.417 ± 0.158	0.531 ± 0.059

Table 8: Universal Style Transfer. Our AR AlexNet autoencoder outperforms both its standard counterpart and the original VGG-19 model in terms of Gram loss, the latter using more layers, larger feature maps and feature blending.

Encoder	Stylization Levels	Smallest Feature Map	Feature Blending	Gram Loss \downarrow (x_{cs}, x_s)	SSIM \uparrow (x_{cs}, x_c)
Standard AlexNet	3	$6 \times 6 \times 256$	X	1.694	0.226
AR AlexNet (ours)	3	$6 \times 6 \times 256$	X	1.186	0.259
VGG-19 [27]	5	$14 \times 14 \times 512$	✓	1.223	0.459

6 Downstream Tasks

We further evaluate the benefits of incorporating AR autoencoders into two downstream tasks: style transfer and image denoising. To assess the benefits of AR autoencoders, in each task, we simply replace the standard autoencoders by the AR versions without incorporating any additional task-specific priors or tuning. Despite not tailoring our architecture to each scenario, it obtains on-par or better results than well-established methods. Refer to Sec. A3 and Sec. A4 for more results and full implementation details.

6.1 Style Transfer via Robust Feature Alignment

Motivated by the perceptual properties of AR features [23], we analyze their impact on style transfer using our AR AlexNet autoencoder as backbone and measure their improvement in both structure and texture preservation.

Experimental Setup. Stylization is evaluated on 75 random content images and 100 random style images, leading to 7,500 image pairs. Content and style preservation is evaluated via the SSIM between content and stylized images and the VGG-19 Gram loss between style and stylized images, respectively. Conv1 and conv2 models use nearest neighbor interpolation instead of transposed convolution layers to improve reconstruction and avoid checkerboard effects, while the conv5 model remains unaltered. We also include results using Universal Style Transfer’s (UST) official implementation, using a VGG-19 backbone.

Results. Our AR autoencoder improves both texture and structure preservation over its standard version (Tab. 8). Stylization via AR features removes artifacts in flat areas, reducing blurry outputs and degraded structure (Fig. 9). Besides, our AR model gets a lower Gram loss with respect to UST. This implies




Fig. 9: Universal Style Transfer: By inverting AR features, our autoencoder improves both content and style preservation, obtaining a better image stylization.

Table 9: Image denoising ($\sigma = \frac{50}{255}$): Our AR denoiser outperforms its standard version on multiple datasets. On the largest one (CBSD68), it also outperforms alternative learn-based techniques. On smaller sets (Kodak24, McMaster), it improves in SSIM and gets comparable PSNR and LPIPS performance.

	PSNR (dB) \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR (dB) \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR (dB) \uparrow	SSIM \uparrow	LPIPS \downarrow
Encoder	CBSD68			Kodak24			McMaster		
TNRD [55]	24.75	0.662	0.445	25.994	0.695	0.461	25.01	0.66	0.387
MLP [56]	25.184	0.663	0.46	26.31	0.691	0.478	26.039	0.693	0.402
Standard	22.6297	0.6178	0.567	23.1868	0.6001	0.4968	23.1493	0.6072	0.4458
AR (ours)	25.258	0.7095	0.4043	25.4946	0.701	0.447	25.3527	0.6914	0.3965

that, despite matching less feature maps than the VGG-19 model (three instead of five), stylizing via our AR AlexNet autoencoder better preserves the style.

As expected, UST obtains a better SSIM since VGG-19 has more complexity and uses less contracted feature maps than our AlexNet model (e.g. $14 \times 14 \times 512$ vs. $6 \times 6 \times 256$). Also, UST *blends* stylized and content features to better preserve shapes. Overall, a comparison between our AR model and UST shows a tradeoff between content and style preservation.

6.2 Image Denoising via AR Autoencoder

Similarly to the robustness imposed by regularized autoencoders [40,53,54], we harness the manifold learned by AR models to obtain noise-free reconstructions. We evaluate our AR AlexNet denoising model and compare its restoration properties with alternative learn-based methods.

Experimental Setup. Our image denoising model consists of an AR autoencoder equipped with skip connections in `conv1`, `conv2` and `conv5` layers to better preserve image details. Skip connections follow the Wavelet Pooling approach [3]. Generators are trained on ImageNet via pixel and feature losses.

Accuracy is evaluated on the Kodak24, McMaster [57] and Color Berkeley Segmentation Dataset 68 (CBSD68) [58] for clipped additive Gaussian noise ($\sigma = 50/255$). We compare our AR model against two learn-based methods, Trainable Nonlinear Reaction Diffusion (TNRD) [55] and Multi Layer Perceptron-based model (MLP) [56], often included in real-noise denoising benchmarks [59,60].

Results. Our AR model improves over its standard version in all metrics across all datasets (Tab. 9). While standard predictions include color distortions




Fig. 10: Image denoising ($\sigma = \frac{50}{255}$): While inverting standard features introduces artifacts and degrades color, limiting their use for restoration tasks, our AR denoiser reduces the artifacts and better preserves the original texture.

and texture artifacts, AR predictions show a better texture preservation and significantly reduce the distortions introduced by the denoising process (Fig. 10).

Our AR model obtains the best PSNR, SSIM and LPIPS scores on CBSD68, the most diverse of all datasets. While it is outperformed in PSNR by MLP in the two remaining datasets, it improves in SSIM and LPIPS, getting best or second best performance. For the McMaster dataset, SSIM and LPIPS values obtained by our model are slightly below the best values. Overall, our model consistently preserves the perceptual and structural similarity across all datasets, showing competitive results with alternative data-driven approaches.

7 Conclusions

A novel encoding-decoding model for synthesis tasks is proposed by exploiting the perceptual properties of AR features. We show the reconstruction improvement obtained by generators trained on AR features and how it generalizes to models of different complexity. We showcase our model on style transfer and image denoising tasks, outperforming standard approaches and attaining competitive performance against alternative methods. A potential limitation of our model is the loss of details due to its contracted features. Yet, experiments show that using shortcut connections allow preserving these, enabling enhancement and restoration tasks. Our method also requires pre-training an AR encoder prior to training the generator, which may increase its computational requirements.

Learning how to invert AR features may be interestingly extended to conditional GANs for image-to-image translation tasks [65] and to VAEs as a latent variable regularizer [19]. Our AR autoencoder can also be seen as an energy-based model [5] for artificial and biological neural networks visualization [7,14,15].

Acknowledgements. AN was supported by NSF Grant No. 1850117 & 2145767, and donations from NaphCare Foundation & Adobe Research. We are grateful for Kelly Price’s tireless assistance with our GPU servers at Auburn University.

References

1. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2016) 2414–2423 [1](#), [3](#), [5](#)
2. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Universal style transfer via feature transforms. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. NIPS’17, Red Hook, NY, USA, Curran Associates Inc. (2017) 385–395 [1](#), [2](#), [5](#), [27](#)
3. Yoo, J., Uh, Y., Chun, S., Kang, B., Ha, J.W.: Photorealistic style transfer via wavelet transforms. In: Proceedings of the IEEE International Conference on Computer Vision. (2019) 9036–9045 [1](#), [6](#), [13](#), [30](#)
4. Yang, C., Lu, X., Lin, Z., Shechtman, E., Wang, O., Li, H.: High-resolution image inpainting using multi-scale neural patch synthesis. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2017) 6721–6729 [1](#)
5. Nguyen, A., Clune, J., Bengio, Y., Dosovitskiy, A., Yosinski, J.: Plug & play generative networks: Conditional iterative generation of images in latent space. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2017) 4467–4477 [1](#), [2](#), [14](#)
6. Shocher, A., Gandelsman, Y., Mosseri, I., Yarom, M., Irani, M., Freeman, W.T., Dekel, T.: Semantic pyramid for image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020) 7457–7466 [1](#), [2](#)
7. Nguyen, A., Dosovitskiy, A., Yosinski, J., Brox, T., Clune, J.: Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. In: Advances in neural information processing systems. (2016) 3387–3395 [1](#), [2](#), [14](#)
8. Rombach, R., Esser, P., Ommer, B.: Network-to-network translation with conditional invertible neural networks. *Advances in Neural Information Processing Systems* **33** (2020) 2784–2797 [1](#)
9. Santurkar, S., Ilyas, A., Tsipras, D., Engstrom, L., Tran, B., Madry, A.: Image synthesis with a single (robust) classifier. In: Advances in Neural Information Processing Systems. (2019) 1262–1273 [1](#), [2](#)
10. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2018) 586–595 [1](#), [2](#)
11. Goodfellow, I., Bengio, Y., Courville, A.: Deep learning book. MIT Press **521** (2016) 800 [1](#), [3](#)
12. Deecke, L., Vandermeulen, R., Ruff, L., Mandt, S., Kloft, M.: Image anomaly detection with generative adversarial networks. In: Joint european conference on machine learning and knowledge discovery in databases, Springer (2018) 3–17 [1](#), [2](#), [5](#), [19](#), [20](#), [21](#)
13. Golan, I., El-Yaniv, R.: Deep anomaly detection using geometric transformations. arXiv preprint arXiv:1805.10917 (2018) [1](#), [19](#), [20](#)
14. Nguyen, A., Yosinski, J., Clune, J.: Understanding neural networks via feature visualization: A survey. In: Explainable AI: interpreting, explaining and visualizing deep learning. Springer (2019) 55–76 [1](#), [14](#)
15. Ponce, C.R., Xiao, W., Schade, P.F., Hartmann, T.S., Kreiman, G., Livingstone, M.S.: Evolving images for visual neurons using a deep generative network reveals coding principles and neuronal preferences. *Cell* **177** (2019) 999–1009 [1](#), [14](#)

16. Rombach, R., Esser, P., Blattmann, A., Ommer, B.: Invertible neural networks for understanding semantics of invariances of cnn representations. In: *Deep Neural Networks and Data for Automated Driving*. Springer (2022) 197–224 [1](#), [2](#)
17. Donahue, J., Simonyan, K.: Large scale adversarial representation learning. (2019) [1](#), [2](#)
18. Dosovitskiy, A., T.Brox: Inverting visual representations with convolutional networks. In: *CVPR*. (2016) [1](#)
19. Dosovitskiy, A., Brox, T.: Generating images with perceptual similarity metrics based on deep networks. In: *Advances in neural information processing systems*. (2016) 658–666 [1](#), [2](#), [3](#), [4](#), [12](#), [14](#), [23](#), [26](#), [27](#)
20. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. (2021) 12873–12883 [2](#)
21. Esser, P., Rombach, R., Blattmann, A., Ommer, B.: Imagebart: Bidirectional context with multinomial diffusion for autoregressive image synthesis. *Advances in Neural Information Processing Systems* **34** (2021) 3518–3532 [2](#)
22. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. *iclr*. arXiv preprint arXiv:1706.06083 (2018) [2](#), [3](#), [7](#), [26](#)
23. Engstrom, L., Ilyas, A., Santurkar, S., Tsipras, D., Tran, B., Madry, A.: Adversarial robustness as a prior for learned representations. arXiv preprint arXiv:1906.00945 (2019) [2](#), [4](#), [11](#), [12](#), [23](#), [26](#)
24. Razavi, A., van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with vq-vae-2. In: *Advances in neural information processing systems*. (2019) 14866–14876 [2](#)
25. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. *Advances in neural information processing systems* **30** (2017) [2](#)
26. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*. (2012) 1097–1105 [2](#)
27. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014) [2](#), [12](#)
28. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2016) 770–778 [2](#)
29. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. (2009) [2](#)
30. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* **115** (2015) 211–252 [2](#)
31. Agustsson, E., Timofte, R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. (2017) 126–135 [2](#), [11](#)
32. Dosovitskiy, A., Brox, T.: Inverting convolutional networks with convolutional networks. arXiv preprint arXiv:1506.02753 **4** (2015) [2](#), [3](#), [11](#)
33. Pan, S.J., Yang, Q.: A survey on transfer learning. *IEEE Transactions on knowledge and data engineering* **22** (2009) 1345–1359 [3](#)
34. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *European conference on computer vision*, Springer (2016) 694–711 [3](#)

35. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: Visualising image classification models and saliency maps. arXiv preprint arXiv:1312.6034 (2013) [3](#)
36. Mahendran, A., Vedaldi, A.: Understanding deep image representations by inverting them. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2015) 5188–5196 [2](#), [3](#)
37. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572 (2014) [3](#)
38. Athalye, A., Carlini, N., Wagner, D.: Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. arXiv preprint arXiv:1802.00420 (2018) [3](#)
39. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2018) 9446–9454 [4](#)
40. Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., Manzagol, P.A., Bottou, L.: Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research* **11** (2010) [5](#), [6](#), [13](#)
41. Kessy, A., Lewin, A., Strimmer, K.: Optimal whitening and decorrelation. *The American Statistician* **72** (2018) 309–314 [5](#)
42. Mao, X.J., Shen, C., Yang, Y.B.: Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. arXiv preprint arXiv:1603.09056 (2016) [6](#)
43. El Helou, M., Süssstrunk, S.: Blind universal bayesian image denoising with gaussian noise level learning. *IEEE Transactions on Image Processing* **29** (2020) 4885–4897 [6](#)
44. Zhang, K., Zuo, W., Zhang, L.: Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing* **27** (2018) 4608–4622 [6](#)
45. Moeller, M., Diebold, J., Gilboa, G., Cremers, D.: Learning nonlinear spectral filters for color image reconstruction. In: Proceedings of the IEEE International Conference on Computer Vision. (2015) 289–297 [6](#)
46. Liu, X., Li, Y., Wu, C., Hsieh, C.J.: Adv-bnn: Improved adversarial defense through robust bayesian neural network. arXiv preprint arXiv:1810.01279 (2018) [7](#), [8](#)
47. Zhang, J., Zhu, J., Niu, G., Han, B., Sugiyama, M., Kankanhalli, M.: Geometry-aware instance-reweighted adversarial training. arXiv preprint arXiv:2010.01736 (2020) [8](#), [9](#)
48. Croce, F., Hein, M.: Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In: International conference on machine learning, PMLR (2020) 2206–2216 [8](#), [9](#)
49. Sosnovik, I., Szmaja, M., Smeulders, A.: Scale-equivariant steerable networks. arXiv preprint arXiv:1910.11093 (2019) [10](#)
50. Fan, Y., Yu, J., Liu, D., Huang, T.S.: Scale-wise convolution for image restoration. In: Proceedings of the AAAI Conference on Artificial Intelligence. Volume 34. (2020) 10770–10777 [10](#)
51. Chen, P., Agarwal, C., Nguyen, A.: The shape and simplicity biases of adversarially robust imagenet-trained cnns. arXiv preprint arXiv:2006.09373 (2020) [10](#)
52. Salman, H., Ilyas, A., Engstrom, L., Kapoor, A., Madry, A.: Do adversarially robust imagenet models transfer better? arXiv preprint arXiv:2007.08489 (2020) [2](#), [10](#)
53. Rifai, S., Vincent, P., Muller, X., Glorot, X., Bengio, Y.: Contractive auto-encoders: Explicit invariance during feature extraction. In: *Icml*. (2011) [13](#)

54. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013) [3](#), [13](#)
55. Chen, Y., Pock, T.: Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE transactions on pattern analysis and machine intelligence* **39** (2016) 1256–1272 [13](#)
56. Burger, H.C., Schuler, C.J., Harmeling, S.: Image denoising: Can plain neural networks compete with bm3d? In: 2012 IEEE conference on computer vision and pattern recognition, IEEE (2012) 2392–2399 [13](#)
57. Zhang, L., Wu, X., Buades, A., Li, X.: Color demosaicking by local directional interpolation and nonlocal adaptive thresholding. *Journal of Electronic imaging* **20** (2011) 023016 [13](#)
58. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: *Proc. 8th Int’l Conf. Computer Vision*. Volume 2. (2001) 416–423 [13](#)
59. Anwar, S., Barnes, N.: Real image denoising with feature attention. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. (2019) 3155–3164 [13](#)
60. Guo, S., Yan, Z., Zhang, K., Zuo, W., Zhang, L.: Toward convolutional blind denoising of real photographs. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. (2019) 1712–1722 [13](#)
61. Mahendran, A., Vedaldi, A.: Visualizing deep convolutional neural networks using natural pre-images. *International Journal of Computer Vision* **120** (2016) 233–255 [2](#)
62. Zhang, Y., Jia, R., Pei, H., Wang, W., Li, B., Song, D.: The secret revealer: Generative model-inversion attacks against deep neural networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. (2020) 253–261 [3](#)
63. Jun, H., Child, R., Chen, M., Schulman, J., Ramesh, A., Radford, A., Sutskever, I.: Distribution augmentation for generative modeling. In: *International Conference on Machine Learning*, PMLR (2020) 5006–5019 [3](#)
64. Ng, A., et al.: Sparse autoencoder. *CS294A Lecture notes* **72** (2011) 1–19 [3](#)
65. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2017) 1125–1134 [14](#)
66. Ruff, L., Vandermeulen, R.A., Görnitz, N., Binder, A., Müller, E., Müller, K.R., Kloft, M.: Deep semi-supervised anomaly detection. *arXiv preprint arXiv:1906.02694* (2019) [19](#)
67. Wang, S., Zeng, Y., Liu, X., Zhu, E., Yin, J., Xu, C., Kloft, M.: Effective end-to-end unsupervised outlier detection via inlier priority of discriminative network. In: *NeurIPS*. (2019) 5960–5973 [19](#)
68. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.: Cats and dogs. In: 2012 IEEE conference on computer vision and pattern recognition, IEEE (2012) 3498–3505 [20](#)
69. Engstrom, L., Ilyas, A., Salman, H., Santurkar, S., Tsipras, D.: Robustness (python library) (2019) [22](#), [27](#)
70. Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957* (2018) [28](#)