

Does Robustness on ImageNet Transfer to Downstream Tasks?

Yutaro Yamada
 Yale University

yutaro.yamada@yale.edu

Mayu Otani
 CyberAgent, Inc.

otani_mayu@cyberagent.co.jp

Abstract

As clean ImageNet accuracy nears its ceiling, the research community is increasingly more concerned about robust accuracy under distributional shifts. While a variety of methods have been proposed to robustify neural networks, these techniques often target models trained on ImageNet classification. At the same time, it is a common practice to use ImageNet pretrained backbones for downstream tasks such as object detection, semantic segmentation, and image classification from different domains. This raises a question: Can these robust image classifiers transfer robustness to downstream tasks? For object detection and semantic segmentation, we find that a vanilla Swin Transformer, a variant of Vision Transformer tailored for dense prediction tasks, transfers robustness better than Convolutional Neural Networks that are trained to be robust to the corrupted version of ImageNet. For CIFAR10 classification, we find that models that are robustified for ImageNet do not retain robustness when fully fine-tuned. These findings suggest that current robustification techniques tend to emphasize ImageNet evaluations. Moreover, network architecture is a strong source of robustness when we consider transfer learning.

1. Introduction

ImageNet [7] serves as an important benchmark in the field of computer vision. Numerous models and training techniques have emerged out of this benchmark [11, 17]. A newly proposed vision architecture, including recent Vision Transformer [8], is first tested against ImageNet to demonstrate a good performance before it gains popularity within the community. While accuracy on ImageNet has been considered as a surrogate for measuring progress in machine vision systems, the research community is now aware of the lack of robustness of vision models towards small input perturbations. [33] first reported that imperceptible adversarial perturbations can easily fool image classifiers. Recent studies show that even simpler, more natural noises such as blur, contrast change, and snow can significantly degrade the per-

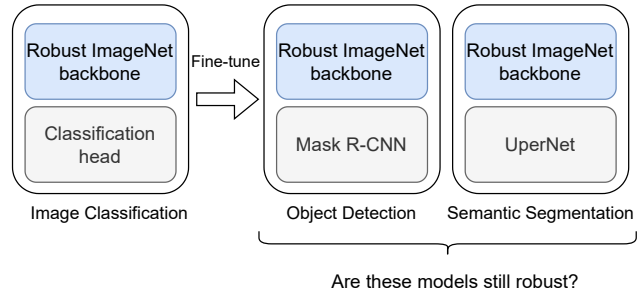


Figure 1. It is a common practice to use ImageNet classifiers as initialization for downstream tasks such as object detection and semantic segmentation. When we fine-tune robust ImageNet classifiers for downstream tasks, should we expect that the resulting vision system still maintains robustness? We tackle this question in the settings of fixed-feature and full-network transfer learning.

formance of models [13]. A typical strategy to increase robustness is data augmentation, where a vision model is trained with additional data, which are artificially corrupted during training. Examples include ANT [29], AugMix [14], and DeepAug [12]. However, these techniques often focus on improving robust accuracy for ImageNet classification. In fact, there are now a variety of ImageNet-scale robustness benchmarks, and the community is striving to improve accuracy on these benchmarks [2, 12, 15].

Due to the scale of ImageNet, it is a common practice to use ImageNet pretrained weights for downstream tasks such as object detection [16] and image segmentation [5, 10]. This practice of using pretrained ImageNet weights for transfer learning raises a fundamental question from a robustness perspective: When we use pretrained weights that are made to be robust to ImageNet benchmarks, do these models necessarily show robustness for downstream tasks as well? (See Figure 1 for the problem setting we consider.)

Contributions.

We find that when we freeze the backbone of ImageNet models, robustified Convolutional Neural Networks (CNNs) maintain robustness for object detection and semantic segmentation. These robustified CNNs continue to

demonstrate higher robustness than the regular model even when we fully fine-tune the whole network, which is practically more relevant. However, perhaps more notably, we observe that Swin Transformer [23], a variant of Vision Transformer tailored to dense prediction tasks, transfers robustness better than robustified CNNs in this fully-finetuned setting. Moreover, it seems difficult to transfer corruption robustness from ImageNet to CIFAR10 [21]. In fact, we find that a non-robustified ImageNet pretrained ResNet performs the best when fine-tuned for CIFAR10. We hope these findings encourage the community to reconsider how we evaluate the robustness of vision systems, as existing data augmentation techniques for robustifying neural networks might be overfitting to ImageNet benchmarks. Furthermore, it is noteworthy that, for robustness transfer, the robustness contribution from Swin Transformer architecture is more significant than the existing robustification methods.

Scope. While there are various kinds of distributional shifts and robustness that the vision community studies, we focus on common corruption robustness in this paper, because we are interested in robustness transfer from ImageNet classification to downstream tasks such as object detection and segmentation. See Section 3.1 for more details about why we specifically choose common corruptions as a topic of our study.

2. Background

Ensuring robustness in downstream tasks such as object detection and semantic segmentation is equally, if not more, important than achieving robustness in image classification. Especially for safety-critical applications such as self-driving cars, vision systems that are vulnerable to image perturbations can lead to dire consequences. In such real-world applications, classification is only the first step of the pipeline, and ensuring robustness through the entire system of object detection and segmentation needs further care.

When we consider how to ensure robustness for downstream tasks, there are two viable approaches. One is to transfer robustness effectively from a pretrained, robustified classifier backbone to each downstream task, which is our focus of this paper. The other approach is to apply an existing robust data augmentation technique during transfer learning. While applying robustification techniques during finetuning for downstream tasks is an option, a naive application of these methods can decrease downstream task performances¹ and often requires further modifications tailored for downstream tasks to maintain good accuracy while achieving robustness [6], partly because object detection and semantic segmentation systems tend to be more com-

¹See Table 4 in the Appendix as an example

plex than image classification. Therefore, rather than entirely resorting to data augmentation during fine-tuning, it is critical to better understand robustness transfer to achieve both robustness and good clean accuracy in downstream tasks.

2.1. Vision Transformer for Dense Prediction Tasks

While image classification only requires a single feature map typically extracted from the last layer, object detection and semantic segmentation benefits a lot from multiresolution feature maps. These feature maps provide richer information that helps object detection at different scales and pixel-level semantic prediction. Most object detection and semantic segmentation systems use a CNN as their backbone and exploit hierarchical feature maps that are extracted from different blocks of the model.

Motivated by the success of Transformer architecture in the Natural Language Processing (NLP) community, Vision Transformer (ViT) [8] was proposed. While the original ViT excels at image classification, it is not amenable to dense prediction tasks such as object detection and semantic segmentation. This is because the original ViT processes tokens at fixed scale, producing single low-resolution feature maps. Recently, a variant of ViT called Swin Transformer was proposed to address this limitation [23]. Swin Transformer uses a hierarchical architecture to build multiresolution feature maps, while achieving linear-time complexity with respect to the image size. Because of this, Swin Transformer achieves the state-of-the-art performance in both object detection and semantic segmentation. In this work, we use Swin Transformer for our ViT architecture.

2.2. Source of Robustness: Data augmentation and Architecture

Vanilla CNNs are vulnerable to image corruptions, as extensively studied by the vision community in the past. [13] shows that state-of-the-art ImageNet classifiers fail when naturally occurring image corruptions are applied to the ImageNet test set, which are introduced as ImageNet-C. To tackle this problem, the community develops many approaches relying on data augmentation [12, 37]. On the other hand, recent studies show that ViT is more robust to ImageNet-C than vanilla CNNs [3, 27] without resorting to data augmentation. These findings suggest that robustness arises both from data augmentation techniques and architecture itself. In terms of robustness transfer, it is unclear which source of robustness is more important. The following sections explore this question in depth.

3. Fixed-Feature Transfer Learning

When we consider transfer learning from image classifiers to object detection or segmentation, we can freeze the backbone, while only training the head of the detection or

Method	Noise	Blur	Digital	Weather
Regular	36.09	44.00	21.17	17.59
ANT	21.90	39.25	17.70	16.22
DeepAug+	16.39	29.25	15.49	11.27
Swin-T	18.01	38.18	17.90	10.12

Table 1. Performance drops across models and noise types are presented for fixed-feature transfer learning from ImageNet to COCO Object Detection. Regular represents a regular ImageNet-pretrained ResNet50, while DeepAug+ and ANT are ResNet50s that are robustified during ImageNet pretraining. Swin-T is a Swin Transformer (Tiny), where the model size is similar to ResNet50.

Method	Noise	Blur	Digital	Weather
Regular	48.98	29.42	17.01	25.68
ANT	17.78	23.41	12.99	25.62
DeepAug+	20.07	19.47	13.01	19.12
Swin-T	13.57	23.50	16.42	14.28

Table 2. Performance drops across models and noise types are presented for fixed-feature transfer learning from ImageNet to ADE10K Semantic Segmentation.

segmentation system. We refer to this approach as fixed-feature transfer learning. On the other hand, we can use pre-trained image classifiers as initialization to train object detection or segmentation models, which we call full-network transfer learning.

Fixed-feature transfer learning from ImageNet to object detection and semantic segmentation is not a common practice because full-network transfer learning generally performs better [16]. However, for *robustness transfer*, fixed-feature transfer learning is an important setup to consider because it allows us to directly leverage robustified ImageNet backbones and measure how much robustness the model carries over to downstream tasks after fine-tuning only the head of the entire model. Full-network transfer learning, on the other hand, potentially erases the robustness property of backbones during fine-tuning, which can confound our analysis of robustness transfer.

From earlier work [27] on robustness of image classifiers, a vanilla Swin Transformer is known to perform better than CNNs on ImageNet-C. At the same time, we can robustify these CNNs by data augmentation, so that they perform well on ImageNet-C. Should we expect that robustified CNNs transfer their robustness automatically when we only fine-tune the head while fixing the backbone? Moreover, which source of robustness (architecture vs. data augmentation) is better suited in terms of robustness transfer?

To resolve this question, we prepare two CNNs that are robustified during ImageNet-1k pretraining using ANT [29] and DeepAug+AugMix [12] respectively, and a Swin Transformer, also pretrained on ImageNet-1k but without applying any robustification technique. To control for the

model size, we use ResNet50 and Swin-T, where the parameter counts are 25M and 28M, respectively. For object detection, we use Mask-RCNN [10] and for semantic segmentation, we use UperNet [38] as the head.

3.1. Robustness Transfer Benchmark

To measure how well a model transfers robustness from ImageNet classification to downstream tasks, we have to prepare the same set of distributional shifts that can be applied to both classification and downstream tasks. While there are a variety of ImageNet-related benchmarks to measure robustness against distributional shifts (e.g. adversarial [15], viewpoint change [2], and background shift [37]), most of these distributional shifts are not adoptable to our setting, because they are specifically designed for ImageNet classification. To measure the performance of robustness transfer to downstream tasks, we focus on 15 synthetic image corruption types, grouped into 4 categories: “noise”, “blur”, “weather”, and “digital”, introduced in ImageNet-C [13]. They measure corruption robustness of ImageNet classifiers by computing how much the original accuracy drops when these models are evaluated on corrupted images of the ImageNet test set. Since these image corruptions are algorithmically generated, they can be applied to images in both classification and downstream tasks such as object detection and segmentation.² Therefore, these image corruptions allow us to compare the accuracy drop in classification with accuracy drop in downstream tasks, which is useful to measure the degree of robustness transfer across different models.

Formally, we take ImageNet models and fine-tune the head of these models for downstream tasks. We calculate model performance on the clean test set in downstream tasks, and compute the performance drop after we apply image corruptions. We then compare the accuracy drop for classification and downstream tasks. We report the mean performance drop across the 15 image corruptions as our metric. The benchmark performance is computed in terms of mean performance under corruption:

$$mPC = \frac{1}{N_c} \sum_{c=1}^{N_c} P_c, \quad (1)$$

where N_c is 15, and P_c is the task-specific performance measure evaluated under corruption c on the test set. We then compute the relative performance under corruption:

$$rPC = \frac{mPC}{P_{clean}} \quad (2)$$

where P_{clean} is the task-specific performance measure evaluated on the clean test set. We use $1 - rPC$ as one of our

²We use the following python library to generate synthetic image corruptions <https://github.com/bethgelab/imagecorruptions> introduced by [25].

main metrics to report and refer to this metric as Accuracy Drop or Performance Drop depending on the context. rPC allows us to compare the degree of robustness transfer from ImageNet to downstream tasks such as object detection and semantic segmentation.

Dataset. For object detection, we choose MS-COCO [22] and use the COCO 2017 validation set for COCO as our test split, following the convention. For semantic segmentation, we choose ADE20K [41]. ADE20K consists of 20210 train, 2000 validation images, and 150 semantic classes. For downstream-task specific performance measures, we use the following metrics:

Object Detection. We use the COCO Average Precision metric, which averages over Intersection-over-Unions (IoUs) between 50% and 95%.

Semantic Segmentation. We use the mean IoU, which indicates the intersection-over-union between the predicted and ground truth pixels, averaged over all the classes.

Table 1 and 2 summarize the results for the fixed feature transfer learning experiment. While ANT and DeepAug+ transfer robustness well across both downstream tasks, we also notice that for some noise types, Swin-T outperforms the robust CNNs (e.g. Noise, Weather in Table 2 and Weather in Table 1.) This suggests that, to our surprise, a vanilla Swin Transformer has a potential to transfer robustness better than robust CNNs. In the next section, we investigate to what extent these phenomena can be observed in the full-network transfer learning setting.

4. Full-Network Transfer Learning

A more common practice to perform transfer learning is to use ImageNet pretrained weights as initialization and fine-tune the entire network for downstream tasks. Even though it takes more computational resources than the fixed-feature case, full-network transfer learning generally performs better [16].

However, when we take robustness into consideration, full-network transfer learning can be detrimental, because gradient updates during fine-tuning can erase robustified features acquired during ImageNet pretraining. This possibility is especially concerning for robustification techniques that rely on data augmentation during pretraining such as DeepAug, AugMix, and ANT. Thus, one may argue that robustness arising from these data augmentation techniques might be less effective when we fine-tune the entire network for downstream tasks. On the other hand, robustness arising from the architecture itself can be more resistant to full-network fine-tuning, because the robustness property is not directly encoded into weights, but rather stems from the topology of architecture. Thus, we do not need to worry

Model	#params	Pre-train Data	Input size	window
Tiny	29M	IN-1k	224	7
Small	50M	IN-1k	224	7
Base''	88M	IN-22k	224	7
Base'	88M	IN-22k	384	12
Base	88M	IN-1k	224	7

Table 3. Swin Transformer architectures we use to test common corruption robustness on object detection and semantic segmentation. The pre-training data is either the ImageNet-1K or ImageNet-22k training set.

about erasing robustness that arises from architecture during transfer learning. As we see that a vanilla Swin Transformer outperforms robustified CNNs for some noise types in the Section 3, architecture indeed plays some role in transferring robustness. Therefore, we hypothesize that in the setting of full-network transfer learning, Transformer architectures might be more effective than CNNs that are robustified via data augmentation.

To test this hypothesis, we repeat the same set of experiments as in the Section 3, but now train all weights for object detection, semantic segmentation, and image classification. For downstream image classification tasks, we choose CIFAR10. The results are shown in Figure 2. As a reference, we also plot the original ImageNet accuracy as well as the Top-1 Accuracy Drop on ImageNet-C for all ImageNet models we use. We can confirm that the two robust CNNs (DeepAug+ and ANT) indeed demonstrate higher robustness than Regular. It is noteworthy that a vanilla Swin-T shows slightly higher robustness than ANT (represented as a lower accuracy drop in the blue bar). More surprisingly, Swin-T performs best in object detection and semantic segmentation. This shows that DeepAug+ and ANT are less successful to transfer their ImageNet-C robustness to downstream tasks than Swin-T, supporting our hypothesis. Moreover, when we test robust transfer from ImageNet-C to CIFAR10, we find that these robust models fail to outperform Regular. This shows that robustness from ImageNet for downstream image classification seems to be harder to transfer than object detection and semantic segmentation.

5. Do Larger Models Transfer Robustness Better?

Having established that the Swin Transformer architecture is a strong source of robustness transfer for object detection and semantic segmentation, especially in full-network transfer learning, we now explore whether or not the size of Transformer architecture affects the performance of model robustness for downstream tasks. In this section, we compare Tiny, Small, and Base Swin Transformers in full-network transfer learning, where the detailed configu-

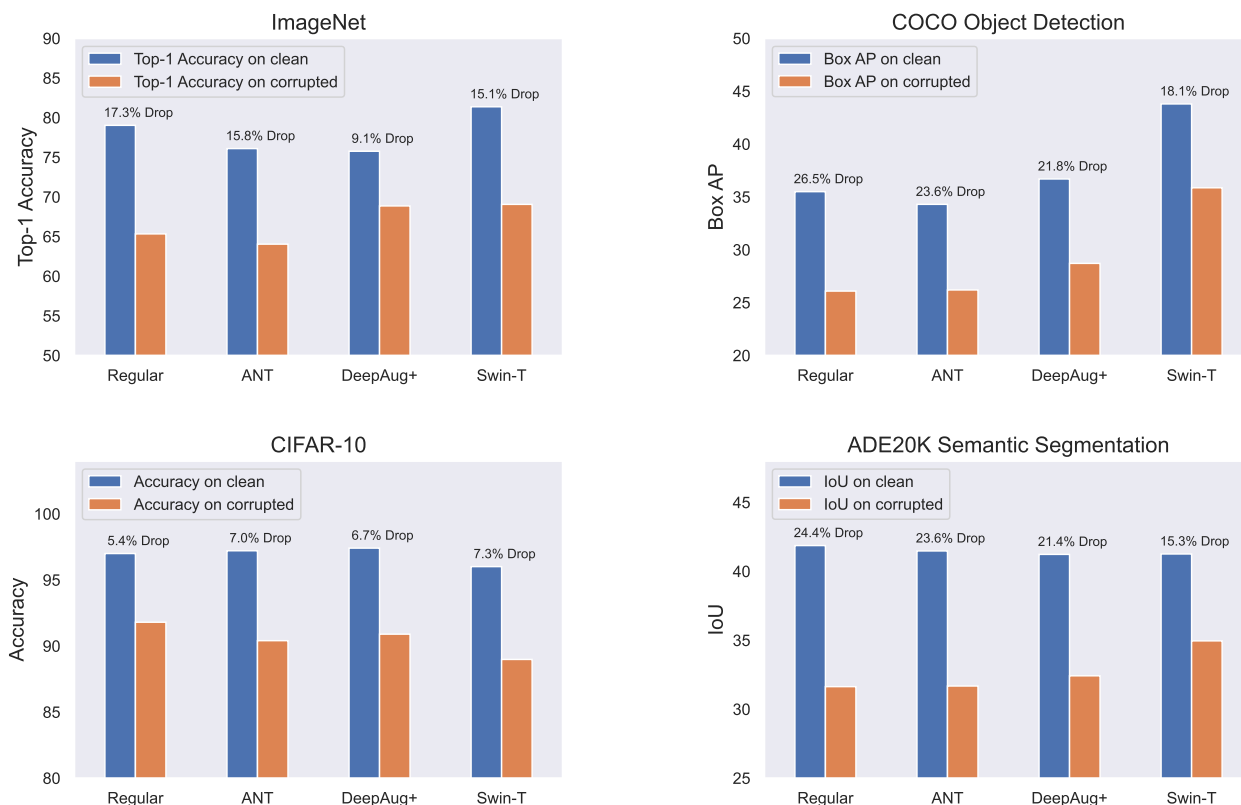


Figure 2. Robust models vs. mean performance drop under 15 corruption types in full-network fine-tuning. The lower the performance drop, the more robust models are to these image corruptions. Regular is a vanilla ResNet50. DeepAug+ and ANT refer to ResNet50 models robustified via DeepAug+AugMix and ANT, which are all data augmentation techniques to increase robustness against common corruptions [12, 29]. Swin-T is a vanilla Tiny Swin Transformer, where the parameter counts are similar to ResNet50. If robustness on ImageNet is transferable to other downstream tasks, we would see a similar pattern of ImageNet-C in object detection and semantic segmentation as well. However, we see that Swin-T performs much better than DeepAug+, the most robust model against ImageNet-C. This shows that the Swin Transformer as architecture is a stronger source of robustness transfer than robustification techniques that are used (e.g. DeepAug, AugMix, or ANT). Moreover, for CIFAR-10, Regular appears to be the most robust model, highlighting the difficulty of transferring ImageNet robustness effectively.

rations are shown in Table 3.

Figure 3 shows the mean performance drop after we apply image corruptions as well as the original performance of each model for ADE20K semantic segmentation, COCO object detection, and COCO instance segmentation. We see that in general the larger the model size is, the smaller the performance drop becomes. This suggests that larger models tend to have more robustness. However, we also observe that there are a few exceptions to this general trend. For instance, Base in ADE20K Semantic Segmentation and Base in COCO Object Detection and Instance Segmentation demonstrate larger performance drop compared to Small. Here we note that the original performances of these large models are similar to Small. We hypothesize that the failure of these large models can be attributed to the pretrained Swin Transformer models, which only use ImageNet-1k for

pretraining. Indeed, when Base is pretrained on ImageNet-22k instead of ImageNet-1k, we see that the IoU Performance Drop is smaller than Small. Similar phenomena are also reported in [3], where large models tend to require more training data to outperform smaller models on clean test sets.

6. Adversarially-trained Networks do not Transfer Robustness to Downstream Tasks

Recent studies [30, 35] find that adversarial robustness is a good prior for transfer learning. Adversarial robustness refers to a model’s stability against small worst-case input perturbations, called adversarial examples [33]. Robustness is typically induced by training a model on adversarial ex-




Figure 3. Swin Transformers varying model size vs. performance drop for downstream tasks under image corruptions. Tiny, Small, and Base are all pretrained on ImageNet-1k while Base' and Base'' are pretrained on ImageNet-22k. See Table 3 for more details about the difference in configurations of models. We see that the larger the models, the more robust in general. However, when we compare Small and Base, it is clear that Base underperforms Small in terms of both robustness to corruption and clean performance. This can be attributed to the pretraining dataset size, where Base requires larger training data to regularize the model than Small.

amples via the following robust optimization objective [24]:

$$\min_{\theta} \mathbb{E}_{x,y \sim \mathcal{D}} \left[\max_{\|\delta\|_2 \leq \epsilon} \mathcal{L}(x + \delta, y; \theta) \right],$$

where θ is the model parameter, the expectation is taken over the data distribution \mathcal{D} , and ϵ controls the magnitude of adversarial perturbation vector δ . Therefore, the larger the ϵ is, the more robust the adversarially-trained models become. Their hypothesis was that adversarially-trained networks maintain better-behaved gradients [34, 40], which might help transfer learning.

While [30, 35] demonstrate that adversarially-trained networks can transfer better than standard models for downstream image classification tasks (without any image corruption), it is unclear how these networks perform in terms of robustness transfer when we consider the performance under image corruptions. In this section, we investigate if adversarial prior is helpful for robustness transfer from image classification to object detection and semantic segmentation. A limitation is that preparing adversarially-trained

models from scratch is difficult since adversarial training is resource intensive. Fortunately, ResNet50 models that are adversarially-trained on ImageNet are made publicly available by [30]. Here, we focus our study on these ResNet50 pretrained models, and will leave for future work how adversarial prior affects Swin Transformer's robustness transfer.

We use four ℓ_2 -robust models that are trained using $\epsilon = 0.1, 1.0, 3.0,$ and 5.0 , respectively, and fine-tune the whole network for COCO object detection and ADE20K semantic segmentation. The original ImageNet clean accuracy as well as mean accuracy drop on ImageNet-C are shown in the top-left panel of Figure 4. We see a clear trend that more robust models tend to perform worse on both clean ImageNet and ImageNet-C. Therefore, the model with $\epsilon = 0.1$ is optimal in terms of both clean accuracy and corruption robustness. For downstream tasks, we see that more robust models perform worse on clean data, but the performance drop is less severe. In fact, for all downstream tasks we consider, the $\epsilon = 0.1$ model performs worst in terms of per-




Figure 4. Adversarial robustness of ImageNet classifiers vs. performance drop for downstream tasks under image corruptions. We vary the ℓ_2 -robustness budget ϵ that is used for adversarial training. The higher the ϵ , the more robust the trained models are towards adversarial attacks in image classification. We can see that the model with $\epsilon = 0.1$ shows the strongest robustness on ImageNet-C, while the same model reveals the worst robustness when fine-tuned for downstream tasks. This shows that adversarial prior is not helpful for robustness transfer as opposed to the findings in [30, 35], where they show adversarial prior is important for transfer learning of image classification.

formance drop on both COCO and ADE20K. This suggests that adversarial prior of ImageNet classifiers is not helpful for robustness transfer to downstream tasks, as opposed to the regular transfer learning setting, where they evaluate clean performance on downstream image classification tasks.

7. Related Works and Discussions

Transfer learning to image classification tasks. [20] performs a large-scale study of transfer learning from ImageNet to other image classification tasks. While they only test CNN architectures, they demonstrate that architectures that perform better on ImageNet are capable of learning better features that are transferable across different classification tasks. On the other hand, they also show that ImageNet pretrained weights do not necessarily transfer well to small fine-grained classification datasets. Our findings add to their results in that architecture is not only beneficial for regular transfer learning but also can be a good source of

robustness transfer.

[5, 16] demonstrate that ImageNet models with higher accuracy transfer better to object detection and semantic segmentation. While these works offer important insights regarding regular transfer learning, our study is orthogonal to these works, because we are specifically interested in how well ImageNet pretrained classifiers can transfer their robustness, instead of clean performance transfer.

Robustness of Vision Transformer. Our work is inspired by recent findings that Vision Transformers show more robustness than CNNs to common image corruptions [27]. [3] shows that larger Vision Transformers require more training data to be robust to ImageNet-C. This is in line with our finding that larger Swin Transformers need ImageNet-21K pretraining to increase robustness in object detection and semantic segmentation. It is hypothesized that Vision Transformers generally require more training data than CNNs since they do not have the inductive bias like convolutions that are useful for image domains [8].

Robust object detection and segmentation. There are several studies that attempt to increase adversarial robustness of object detection systems. [39] relies on adversarial training. However, adversarial perturbations are artificially crafted examples. A more natural situation is object detection under occlusion. [36] addresses such an issue by developing specifically designed architectures to handle occlusion. There are several works that attempt to increase robustness for semantic segmentation. [1] proposes a specialized student-teacher architecture for robust semantic segmentation. [18] relies on increasing shape bias of networks to build the robust semantic segmentation system, inspired by the success of image classification using a similar shape-bias approach [9]. Instead of inducing robustness directly in object detection or semantic segmentation, we study robustness transfer from ImageNet-pretrained models to these downstream tasks. While developing robust detection or segmentation systems is important, we think it is also beneficial to tackle building robust systems from the point of view of transfer learning from robust image classifiers.

[19] benchmarks the robustness of various CNN models for semantic segmentation under image corruptions similar to ImageNet-C. They find that, within the CNN models they tested, models with higher accuracy show stronger robustness on semantic segmentation. Our work is orthogonal to their finding as we start from robust models, and how much robust transfer occurs via fine-tuning.

Robustness Transfer. While there exist prior works on robustness transfer, they focus on transferring adversarial robustness from source models to target models in image classification settings. For instance, [31] proposes lifelong learning strategies to transfer adversarial robustness effectively. [4] shows that adversarial robustness transfer can be achieved by input gradient adversarial matching in the form of student-teacher framework. While these works are important, more practically relevant is the task of robustness transfer for common image corruptions such as noise, blur, and weather change. Furthermore, as opposed to these works, we focus on robustness transfer from image classification to object detection and semantic segmentation.

Overfitting to ImageNet. Our results from fixed-feature transfer learning suggest that robustness of ImageNet-pretrained backbones can be maintained if we freeze the weights of the backbones, but this sacrifices the validation accuracy on downstream tasks. A more ideal scenario would be to preserve robustness even in full-network fine-tuning. However, as Section 4 shows, robustness from weights of CNNs were less effective than robustness of the Swin Transformer architecture. That is, the robustness performance on ImageNet-C is not perfectly transferable to downstream tasks. This result is reminiscent of several re-

ports regarding overfitting to ImageNet. For instance, [28] studies that ImageNet models do not generalize well to additional test data generated using a data collection process similar to that of ImageNet. Another work [32] shows that ImageNet pretrained models do not generalize to videos. Our robustness transfer results add to these works, suggesting again that over-reliance on ImageNet benchmarks can be misleading.

8. Limitations

While we claim that the architecture is a strong source of robustness for transfer learning, this statement is limited in a sense that we only compare ResNet and Swin Transformer. We encourage future work to study broader types of architectures and what properties of models can be well-preserved during transfer learning. We also note that our goal is to study the effect of robust transfer, and therefore we did not necessarily aim for achieving the state-of-the-art performance on downstream tasks. Developing a general recipe for achieving good clean accuracy while maintaining robustness on downstream tasks remains an important future work.

9. Conclusions

In this work, we study the problem of *robustness transfer* from ImageNet pretrained classifiers to downstream tasks such as object detection and semantic segmentation. Our study is motivated by the two observations: 1. Even though there are many proposals to robustify neural networks, these methods target ImageNet classifiers. 2. It is common to use ImageNet pretrained weights for object detection and semantic segmentation. This leads to our central question of this paper: Do robustified ImageNet classifiers necessarily transfer robustness to downstream tasks? In the fixed-feature transfer learning setting, we find that robustness of ImageNet backbones is partially preserved on downstream tasks. However, in full-network transfer learning, which is more practically relevant, we see that the contribution from the Transformer architecture is more significant than the specific robustification techniques that are applied to CNNs. We also test if the adversarial prior, which is shown to be important for regular transfer learning, is also important for robustness transfer. We find that, as opposed to previous findings, the adversarial prior does not help robustness transfer. We hope that our findings encourage the community to reconsider how we evaluate corruption robustness of vision systems.

References

- [1] Andreas Bar, Marvin Klingner, Serin Varghese, Fabian Huger, Peter Schlicht, and Tim Fingscheidt. Robust Semantic Segmentation by Redundant Networks With a Layer-

- Specific Loss Contribution and Majority Vote. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1348–1358, Seattle, WA, USA, June 2020. IEEE.
- [2] Andrei Barbu, David Mayo, Julian Alverio, William Luo, Christopher Wang, Dan Gutfreund, Josh Tenenbaum, and Boris Katz. ObjectNet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. In *Advances in Neural Information Processing Systems, 2019*, volume 32, pages 9453–9463, 2019.
- [3] Srinadh Bhojanapalli, Ayan Chakrabarti, Daniel Glasner, Daliang Li, Thomas Unterthiner, and Andreas Veit. Understanding Robustness of Transformers for Image Classification. *ICCV*, page 11, 2021.
- [4] Alvin Chan, Yi Tay, and Yew-Soon Ong. What It Thinks Is Important Is Important: Robustness Transfers Through Input Gradients. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 329–338, Seattle, WA, USA, June 2020. IEEE.
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(04):834–848, Apr. 2018.
- [6] Xiangning Chen, Cihang Xie, Mingxing Tan, Li Zhang, Chojui Hsieh, and Boqing Gong. Robust and Accurate Object Detection via Adversarial Learning. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16617–16626, Nashville, TN, USA, June 2021. IEEE.
- [7] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, June 2009.
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*, Sept. 2020.
- [9] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, Sept. 2018.
- [10] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *arXiv:1512.03385 [cs]*, Dec. 2015.
- [12] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The Many Faces of Robustness: A Critical Analysis of Out-of-Distribution Generalization. *ICCV*, 2021.
- [13] Dan Hendrycks and Thomas Dietterich. Benchmarking Neural Network Robustness to Common Corruptions and Perturbations. In *International Conference on Learning Representations*, Sept. 2018.
- [14] Dan Hendrycks, Norman Mu, E. D. Cubuk, Barret Zoph, J. Gilmer, and Balaji Lakshminarayanan. AugMix: A Simple Data Processing Method to Improve Robustness and Uncertainty. *ICLR*, 2020.
- [15] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural Adversarial Examples. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15257–15266, Nashville, TN, USA, June 2021. IEEE.
- [16] Jonathan Huang, Vivek Rathod, Chen Sun, Menglong Zhu, Anoop Korattikara, Alireza Fathi, Ian Fischer, Zbigniew Wojna, Yang Song, Sergio Guadarrama, and Kevin Murphy. Speed/Accuracy Trade-Offs for Modern Convolutional Object Detectors. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3296–3297, Honolulu, HI, July 2017. IEEE.
- [17] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 448–456. PMLR, June 2015.
- [18] Christoph Kamann, Burkhard Güssefeld, Robin Huttmacher, Jan Hendrik Metzen, and Carsten Rother. Increasing the Robustness of Semantic Segmentation Models with Painting-by-Numbers. *ECCV*, 12355:369–387, 2020.
- [19] Christoph Kamann and Carsten Rother. Benchmarking the Robustness of Semantic Segmentation Models. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8825–8835, Seattle, WA, USA, June 2020. IEEE.
- [20] Simon Kornblith, Jonathon Shlens, and Quoc V. Le. Do Better ImageNet Models Transfer Better? In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2656–2666, Long Beach, CA, USA, June 2019. IEEE.
- [21] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollar, and Larry Zitnick. Microsoft COCO: Common Objects in Context. In *ECCV*, Sept. 2014.
- [23] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. *International Conference on Computer Vision (ICCV)*, 2021.
- [24] Aleksander Madry, Aleksandar Makelev, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*, Feb. 2018.
- [25] Claudio Michaelis, Benjamin Mitzkus, Robert Geirhos, Evgenia Rusak, Oliver Bringmann, Alexander S. Ecker,

- Matthias Bethge, and Wieland Brendel. Benchmarking Robustness in Object Detection: Autonomous Driving when Winter is Coming. *arXiv:1907.07484 [cs, stat]*, Mar. 2020.
- [26] Apostolos Modas, Rahul Rade, Guillermo Ortiz-Jiménez, Seyed-Mohsen Moosavi-Dezfooli, and Pascal Frossard. PRIME: A few primitives can boost robustness to common corruptions. *arXiv:2112.13547 [cs]*, Mar. 2022.
- [27] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Munawar Hayat, Fahad Khan, and Ming-Hsuan Yang. Intriguing Properties of Vision Transformers. In *Thirty-Fifth Conference on Neural Information Processing Systems*, May 2021.
- [28] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. Do ImageNet Classifiers Generalize to ImageNet? In *Proceedings of the 36th International Conference on Machine Learning, 2019*, pages 5389–5400. PMLR, May 2019.
- [29] Evgenia Rusak, Lukas Schott, Roland S. Zimmermann, Julian Bitterwolf, Oliver Bringmann, Matthias Bethge, and Wieland Brendel. A Simple Way to Make Neural Networks Robust Against Diverse Image Corruptions. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, volume 12348, pages 53–69. Springer International Publishing, Cham, 2020.
- [30] Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do Adversarially Robust ImageNet Models Transfer Better? *NeurIPS*, 2020.
- [31] Ali Shafahi, Parsa Saadatpanah, Chen Zhu, Amin Ghiasi, Christoph Studer, David Jacobs, and Tom Goldstein. Adversarially robust transfer learning. In *International Conference on Learning Representations*, Sept. 2019.
- [32] Vaishal Shankar, Achal Dave, Rebecca Roelofs, Deva Ramanan, Benjamin Recht, and Ludwig Schmidt. Do Image Classifiers Generalize Across Time? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9661–9669, 2021.
- [33] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014.
- [34] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness May Be at Odds with Accuracy. In *International Conference on Learning Representations*, Sept. 2018.
- [35] Francisco Utrera, Evan Kravitz, N. Benjamin Erichson, Rajiv Khanna, and Michael W. Mahoney. Adversarially-Trained Deep Nets Transfer Better: Illustration on Image Classification. In *International Conference on Learning Representations*, Sept. 2020.
- [36] Angtian Wang, Yihong Sun, Adam Kortylewski, and Alan Yuille. Robust Object Detection Under Occlusion With Context-Aware CompositionalNets. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12642–12651, Seattle, WA, USA, June 2020. IEEE.
- [37] Kai Xiao, Logan Engstrom, Andrew Ilyas, and Aleksander Ma. Noise or Signal: The Role of Image Backgrounds in Object Recognition. *ICLR*, page 28, 2021.
- [38] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified Perceptual Parsing for Scene Understanding. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, volume 11209, pages 432–448. Springer International Publishing, Cham, 2018.
- [39] Haichao Zhang and Jianyu Wang. Towards Adversarially Robust Object Detection. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 421–430, Seoul, Korea (South), Oct. 2019. IEEE.
- [40] Tianyuan Zhang and Zhanxing Zhu. Interpreting Adversarially Trained Convolutional Neural Networks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 7502–7511. PMLR, May 2019.
- [41] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene Parsing through ADE20K Dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5122–5130, July 2017.