CheXtransfer: Performance and Parameter Efficiency of ImageNet Models for Chest X-Ray Interpretation

Alexander Ke*
alexke@cs.stanford.edu
Stanford University
USA

William Ellsworth*
willells@cs.stanford.edu
Stanford University
USA

Oishi Banerjee*
oishi.banerjee@cs.stanford.edu
Stanford University
USA

Andrew Y. Ng
ang@cs.stanford.edu
Stanford University
USA

Pranav Rajpurkar
pranavsr@cs.stanford.edu
Stanford University
USA

Figure 1: Visual summary of our contributions. From left to right: scatterplot and best-fit line for 16 pretrained models showing no relationship between ImageNet and CheXpert performance, CheXpert performance relationship varies across architecture families much more than within, average CheXpert performance improves with pretraining, models can maintain performance and improve parameter efficiency through truncation of final blocks. Error bars show one standard deviation.

ABSTRACT

Deep learning methods for chest X-ray interpretation typically rely on pretrained models developed for ImageNet. This paradigm assumes that better ImageNet architectures perform better on chest X-ray tasks and that ImageNet-pretrained weights provide a performance boost over random initialization. In this work, we compare the transfer performance and parameter efficiency of 16 popular convolutional architectures on a large chest X-ray dataset (CheXpert) to investigate these assumptions. First, we find no relationship between ImageNet performance and CheXpert performance for both models without pretraining and models with pretraining. Second, we find that, for models without pretraining, the choice of model family influences performance more than size within a family for medical imaging tasks. Third, we observe that ImageNet

pretraining yields a statistically significant boost in performance across architectures, with a higher boost for smaller architectures. Fourth, we examine whether ImageNet architectures are unnecessarily large for CheXpert by truncating final blocks from pretrained models, and find that we can make models 3.25x more parameter-efficient on average without a statistically significant drop in performance. Our work contributes new experimental evidence about the relation of ImageNet to chest x-ray interpretation performance.

CCS CONCEPTS

• **Applied computing** → **Health informatics**; • **Computing methodologies** → **Computer vision**; *Neural networks*.

KEYWORDS

generalization, efficiency, pretraining, chest x-ray interpretation, ImageNet, truncation

ACM Reference Format:

Alexander Ke, William Ellsworth, Oishi Banerjee, Andrew Y. Ng, and Pranav Rajpurkar. 2021. CheXtransfer: Performance and Parameter Efficiency of ImageNet Models for Chest X-Ray Interpretation. In *ACM Conference on Health, Inference, and Learning (ACM CHIL '21)*, April 8–10, 2021, Virtual Event, . ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3450439.3451867>

* Authors contributed equally to this research.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ACM CHIL '21, April 8–10, 2021, Virtual Event,
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8359-2/21/04.
<https://doi.org/10.1145/3450439.3451867>

1 INTRODUCTION

Deep learning models for chest X-ray interpretation have high potential for social impact by aiding clinicians in their workflow and increasing access to radiology expertise worldwide [22, 23]. Transfer learning using pretrained ImageNet [9] models has been the standard approach for developing models not only on chest X-rays [1, 20, 29] but also for many other medical imaging modalities [8, 10, 16, 17, 31]. This transfer assumes that better ImageNet architectures perform better and pretrained weights boost performance on their target medical tasks. However, there has not been a systematic investigation of how ImageNet architectures and weights both relate to performance on downstream medical tasks.

In this work, we systematically investigate how ImageNet architectures and weights both relate to performance on chest X-ray tasks. Our primary contributions are:

- (1) For models without pretraining and models with pretraining, we find *no relationship between ImageNet performance and CheXpert performance* (Spearman $\rho = 0.08$, $\rho = 0.06$ respectively). This finding suggests that architecture improvements on ImageNet may not lead to improvements on medical imaging tasks.
- (2) For models without pretraining, we find that within an architecture family, the largest and smallest models have small differences (ResNet 0.005, DenseNet 0.003, EfficientNet 0.004) in CheXpert AUC, but different model families have larger differences in AUC (> 0.006). This finding suggests that *the choice of model family influences performance more than size within a family* for medical imaging tasks.
- (3) We observe that *ImageNet pretraining yields a statistically significant boost in performance* (average boost of 0.016 AUC) across architectures, with a higher boost for smaller architectures (Spearman $\rho = -0.72$ with number of parameters). This finding supports the ImageNet pretraining paradigm for medical imaging tasks, especially for smaller models.
- (4) We find that by truncating final blocks of pretrained models, we can make models *3.25x more parameter-efficient on average without a statistically significant drop in performance*. This finding suggests model truncation may be a simple method to yield lighter pretrained models by preserving architecture design features while reducing model size.

Our study, to the best of our knowledge, contributes the first systematic investigation of the performance and efficiency of ImageNet architectures and weights for chest X-ray interpretation. Our investigation and findings may be further validated on other datasets and medical imaging tasks.

2 RELATED WORK

2.1 ImageNet Transfer

Kornblith et al. [15] examined the performance of 16 convolutional neural networks (CNNs) on 12 image classification datasets. They found that using these ImageNet pretrained architectures either as feature extractors for logistic regression or fine tuning them on the target dataset yielded a Spearman $\rho = 0.99$ and $\rho = 0.97$ between ImageNet accuracy and transfer accuracy respectively. However,

they showed ImageNet performance was less correlated with transfer accuracy for some fine-grained tasks, corroborating He et al. [11]. They found that without ImageNet pretraining, ImageNet accuracy and transfer accuracy had a weaker Spearman $\rho = 0.59$. We extend Kornblith et al. [15] to the medical setting by studying the relationship between ImageNet and CheXpert performance.

Raghu et al. [19] explored properties of transfer learning onto retinal fundus images and chest X-rays. They studied ResNet50 and InceptionV3 and showed pretraining offers little performance improvement. Architectures composed of just four to five sequential convolution and pooling layers achieved comparable performance on these tasks as ResNet50 with less than 40% the parameters. In our work, we find pretraining does not boost performance for ResNet50, InceptionV3, InceptionV4, and MNASNet but does boost performance for the remaining 12 architectures. Thus, we were able to replicate Raghu et al. [19]’s results, but upon studying a broader set of newer and more popular models, we reached the opposite conclusion that ImageNet pretraining yields a statistically significant boost in performance.

2.2 Medical Task Architectures

Irvin et al. [13] compared the performance of ResNet152, DenseNet121, InceptionV4, and SEResNeXt101 on CheXpert, finding that DenseNet121 performed best. In a recent analysis, all but one of the top ten CheXpert competition models used DenseNets as part of their ensemble, even though they have been surpassed on ImageNet [21]. Few groups design their own networks from scratch, preferring to use established ResNet and DenseNet architectures for CheXpert [3]. This trend extends to retinal fundus and skin cancer tasks as well, where Inception architectures remain popular [8, 16, 17, 31]. The popularity of these older ImageNet architectures hints that there may be a disconnect between ImageNet performance and medical task performance for newer architectures generated through architecture search. We verify that these newer architectures generated through search (EfficientNet, MobileNet, MNASNet) underperform older architectures (DenseNet, ResNet) on CheXpert, suggesting that search has overfit to ImageNet and explaining the popularity of these older architectures in the medical imaging literature.

Bressem et al. [3] postulated that deep CNNs that can represent more complex relationships for ImageNet may not be necessary for CheXpert, which has greyscale inputs and fewer image classes. They studied ResNet, DenseNet, VGG, SqueezeNet, and AlexNet performance on CheXpert and found that ResNet152, DenseNet161, and ResNet50 performed best on CheXpert AUC. In terms of AUPRC, they showed that smaller architectures like AlexNet and VGG can perform similarly to deeper architectures on CheXpert. Models such as AlexNet, VGG, and SqueezeNet are no longer popular in the medical setting, so in our work, we systematically investigate the performance and efficiency of 16 more contemporary ImageNet architectures with and without pretraining. Additionally, we extend [3] by studying the effects of pretraining, characterizing the relationship between ImageNet and CheXpert performance, and drawing conclusions about architecture design.

2.3 Truncated Architectures

The more complex a convolutional architecture becomes, the more computational and memory resources are needed for its training and deployment. Model complexity thus may impede the deployment of CNNs to clinical settings with less resources. Therefore, efficiency, often reported in terms of the number of parameters in a model, the number of FLOPS in the forward pass, or the latency of the forward pass, has become increasingly important in model design. Low-rank factorization [7, 14], transferred/compact convolutional filters [6], knowledge distillation [12], and parameter pruning [27] have all been proposed to make CNNs more efficient.

Layer-wise pruning is a type of parameter pruning that locates and removes layers that are not as useful to the target task [24]. Through feature diagnosis, a linear classifier is trained using the feature maps at intermediate layers to quantify how much a particular layer contributes to performance on the target task [5]. In this work, we propose model truncation as a simple method for layer-wise pruning where the final pretrained layers after a given point are pruned off, a classification layer is appended, and this whole model is finetuned on the target task.

3 METHODS

3.1 Training and Evaluation Procedure

We train chest X-ray classification models with different architectures with and without pretraining. The task of interest is to predict the probability of different pathologies from one or more chest X-rays. We use the CheXpert dataset consisting of 224,316 chest X-rays of 65,240 patients [13] labeled for the presence or absence of 14 radiological observations. We evaluate models using the average of their AUROC metrics (AUC) on the five CheXpert-defined competition tasks (Atelectasis, Cardiomegaly, Consolidation, Edema, Pleural Effusion) as well as the No Finding task to balance clinical importance and prevalence in the validation set.

We select 16 models pretrained on ImageNet from public checkpoints implemented in PyTorch 1.4.0: DenseNet (121, 169, 201) and ResNet (18, 34, 50, 101) from Paszke et al. [18], Inception (V3, V4) and MNASNet from Cadene [4], and EfficientNet (B0, B1, B2, B3) and MobileNet (V2, V3) from Wightman [30]. We finetune and evaluate these architectures with and without ImageNet pretraining.

For each model, we finetune all parameters on the CheXpert training set. If a model is pretrained, inputs are normalized using mean and standard deviation learned from ImageNet. If a model is not pretrained, inputs are normalized with mean and standard deviation learned from CheXpert. We use the Adam optimizer ($\beta_1 = 0.9$, $\beta_2 = 0.999$) with learning rate of 1×10^{-4} , a batch size of 16, and a cross-entropy loss function. We train on up to four Nvidia GTX 1080 with CUDA 10.1 and Intel Xeon CPU ES-2609 running Ubuntu 16.04. For one run of an architecture, we train for three epochs and evaluate each model every 8192 gradient steps. We train each model and create a final ensemble model from the ten checkpoints, which achieved the best average CheXpert AUC across the six tasks on the validation set. We report all our results on the CheXpert test set.

We use the nonparametric bootstrap to estimate 95% confidence intervals for each statistic. 1,000 replicates are drawn from the test set, and the statistic is calculated on each replicate. This procedure

produces a distribution for each statistic, and we report the 2.5 and 97.5 percentiles as a confidence interval. Significance is assessed at the $p = 0.05$ level.

3.2 Truncated Architectures

We study truncated versions of DenseNet121, MNASNet, ResNet18, and EfficientNetB0. DenseNet121 and MNASNet were chosen because we found they have the greatest efficiency (by AUC per parameters) on CheXpert of the models we profile, ResNet18 was chosen because of its popularity as a compact model for medical tasks, and EfficientNetB0 was chosen because it is the smallest current-generation model of the 16 we study. DenseNet121 contains four dense blocks separated by transition blocks before the classification layer. By pruning the final dense block and associated transition block, the model now only contains three dense blocks, yielding DenseNet121Minus1. Similarly, pruning two dense blocks and associated transition blocks yields DenseNet121Minus2, and pruning three dense blocks and associated transition blocks yields DenseNet121Minus3. For MNASNet, we remove up to the four of the final MBConv blocks to produce MNASNetMinus1 through MNASNetMinus4. For ResNet18, we remove up to the three of the final residual blocks with a similar method to produce ResNet18Minus1 through ResNet18Minus3. For EfficientNet, we remove up to two of the final MBConv6 blocks to produce EfficientNetB0Minus1 and EfficientNetB0Minus2.

After truncating a model, we append a classification block containing a global average pooling layer followed by a fully connected layer to yield outputs of the correct shape. We initialize the model with ImageNet pretrained weights, except the randomly initialized classification block, and finetune using the same training procedure as the 16 ImageNet models.

3.3 Class Activation Maps

We compare the class activation maps (CAMs) among a truncated DenseNet121 family to visualize their higher resolution CAMs. We generate CAMs using the Grad-CAM method [25], using a weighted combination of the model's final convolutional feature maps, with weights based on the positive partial derivatives with respect to class score. This averaged map is scaled by the outputted probability so more confident predictions appear brighter. Finally, the map is upsampled to the input image resolution and overlain onto the input image, highlighting image regions that had the greatest influence on a model's decision.

4 EXPERIMENTS

4.1 ImageNet Transfer Performance

We investigate whether higher performance on natural image classification translates to higher performance on chest X-ray classification. We display the relationship between the CheXpert AUC, with and without ImageNet pretraining, and ImageNet top-1 accuracy in Figure 2

When models are trained without pretraining, we find no monotonic relationship between ImageNet top-1 accuracy and average CheXpert AUC, with Spearman $\rho = 0.082$ at $p = 0.762$. Model performance without pretraining would describe how a given architecture would perform on the target task, independent of any

Figure 2: Average CheXpert AUC vs. ImageNet Top-1 Accuracy. The left plot shows results obtained without pretraining, while the right plot shows results with pretraining. There is no monotonic relationship between ImageNet and CheXpert performance without pretraining (Spearman $\rho = 0.08$) or with pretraining (Spearman $\rho = 0.06$).

pretrained weights. When models are trained with pretraining, we again find no monotonic relationship between ImageNet top-1 accuracy and average CheXpert AUC with Spearman $\rho = 0.059$ at $p = 0.829$.

Overall, we find no relationship between ImageNet and CheXpert performance, so models that succeed on ImageNet do not necessarily succeed on CheXpert. These relationships between ImageNet performance and CheXpert performance are much weaker than the relationships between ImageNet performance and performance on various natural image tasks reported by Kornblith et al. [15].

We compare the CheXpert performance within and across architecture families. Without pretraining, we find that ResNet101 performs only 0.005 AUC greater than ResNet18, which is well within the confidence interval of this metric (Figure 2). Similarly, DenseNet201 performs 0.004 AUC greater than DenseNet121 and EfficientNetB3 performs 0.003 AUC greater than EfficientNetB0. With pretraining, we continue to find minor performance differences between the largest model and smallest model that we test in each family. We find AUC increases of 0.002 for ResNet, 0.004 for DenseNet and -0.006 for EfficientNet. Thus, increasing complexity

within a model family does not yield increases in CheXpert performance as meaningful as the corresponding increases in ImageNet performance.

Without pretraining, we find that the best model studied performs significantly better than the worst model studied. Among models trained without pretraining, we find that InceptionV3 performs best with 0.866 (0.851, 0.880) AUC, while MobileNetV2 performs worst with 0.814 (0.796, 0.832) AUC. Their difference in performance is 0.052 (0.043, 0.063) AUC. InceptionV3 is also the third largest architecture studied and MobileNetV2 the smallest. We find a significant difference in the CheXpert performance of these models. This difference again hints at the importance of architecture design.

4.2 CheXpert Performance and Efficiency

We examine whether larger architectures perform better than smaller architectures on chest X-ray interpretation, where architecture size is measured by number of parameters. We display these relationships in Figure 3 and Table 1.

Without ImageNet pretraining, we find a positive monotonic relationship between the number of parameters of an architecture and CheXpert performance, with Spearman $\rho = 0.791$ significant

Figure 3: Average CheXpert AUC vs. Model Size. The left plot shows results obtained without pretraining, while the right plot shows results with pretraining. The logarithm of the model size has a near linear relationship with CheXpert performance when we omit pretraining (Spearman $\rho = 0.79$). However once we incorporate pretraining, the monotonic relationship is weaker (Spearman $\rho = 0.56$).

Model	CheXpert AUC	#Params (M)
DenseNet121	0.859 (0.846, 0.871)	6.968
DenseNet169	0.860 (0.848, 0.873)	12.508
DenseNet201	0.864 (0.850, 0.876)	18.120
EfficientNetB0	0.859 (0.846, 0.871)	4.025
EfficientNetB1	0.858 (0.844, 0.872)	6.531
EfficientNetB2	0.866 (0.853, 0.880)	7.721
EfficientNetB3	0.853 (0.837, 0.867)	10.718
InceptionV3	0.862 (0.848, 0.876)	27.161
InceptionV4	0.861 (0.846, 0.873)	42.680
MNASNet	0.858 (0.845, 0.871)	5.290
MobileNetV2	0.854 (0.839, 0.869)	2.242
MobileNetV3	0.859 (0.847, 0.872)	4.220
ResNet101	0.863 (0.848, 0.876)	44.549
ResNet18	0.862 (0.847, 0.875)	11.690
ResNet34	0.863 (0.849, 0.875)	21.798
ResNet50	0.859 (0.843, 0.871)	25.557

Table 1: CheXpert AUC (with 95% Confidence Intervals) and Number of Parameters for 16 ImageNet-Pretrained Models.

at $p = 2.62 \times 10^{-4}$. With ImageNet pretraining, there is a weaker positive monotonic relationship between the number of parameters and average CheXpert AUC, with Spearman $\rho = 0.565$ at $p = 0.023$.

Although there exists a positive monotonic relationship between the number of parameters of an architecture and average CheXpert AUC, the Spearman ρ does not highlight the increase in parameters necessary to realize marginal increases in CheXpert AUC. For example, ResNet101 is 11.1x larger than EfficientNetB0, but with only increase of 0.005 in CheXpert AUC with pretraining.

Within a model family, increasing the number of parameters does not lead to meaningful gains in CheXpert AUC. We see this relationship in all families studied without pretraining (EfficientNet, DenseNet, and ResNet) in Figure 3. For example, DenseNet201 has an AUC 0.003 greater than DenseNet121, but is 2.6x larger. EfficientNetB3 has an AUC 0.004 greater than EfficientNetB0, but is 1.9x larger. Despite the positive relationship between model size and CheXpert performance across all models, bigger does not necessarily mean better within a model family.

Since within a model family there is a weaker relationship between model size and CheXpert performance than across all models, we find that CheXpert performance is influenced more by the

macro architecture design than by its size. Models within a family have similar architecture design choices but different sizes, so they perform similarly on CheXpert. We observe large discrepancies in performance between architecture families. For example DenseNet, ResNet, and Inception typically outperform EfficientNet and MobileNet architectures, regardless of their size. EfficientNet, MobileNet, and MNASNet were all generated through neural architecture search to some degree, a process that optimized for performance on ImageNet. Our findings suggest that this search could have overfit to the natural image objective to the detriment of chest X-ray tasks.

4.3 ImageNet Pretraining Boost

We study the effects of ImageNet pretraining on CheXpert performance by defining the pretraining boost as the CheXpert AUC of a model initialized with ImageNet pretraining minus the CheXpert AUC of its counterpart without pretraining. The pretraining boosts of our architectures are reported in Figure 4.

We find that ImageNet pretraining provides a meaningful boost for most architectures (on average 0.015 AUC). We find a Spearman $\rho = -0.718$ at $p = 0.002$ between the number of parameters of a given model and the pretraining boost. Therefore, this boost tends to be larger for smaller architectures such as EfficientNetB0 (0.023), MobileNetV2 (0.040) and MobileNetV3 (0.033) and smaller for larger architectures such as InceptionV4 (-0.002) and ResNet101 (0.013). Further work is required to explain this relationship.

Within a model family, the pretraining boost also does not meaningfully increase as model size increases. For example, DenseNet201 has a pretraining boost only 0.002 AUC greater than DenseNet121 does. This finding supports our earlier conclusion that model families perform similarly on CheXpert regardless of their size.

4.4 Truncated Architectures

We truncate the final blocks of DenseNet121, MNASNet, ResNet18, and EfficientNetB0 with pretrained weights and study their CheXpert performance to understand whether ImageNet models are unnecessarily large for the chest X-ray task. We express efficiency gains in terms of Times-Smaller, or the number of parameters of the original architecture divided by the number of parameters of the truncated architecture: intuitively, how many times larger the original architecture is compared to the truncated architecture. The efficiency gains and AUC changes of model truncation on DenseNet121, MNASNet, ResNet18, and EfficientNetB0 are displayed in Table 2.

For all four model families, we find that truncating the final block leads to no significant decrease in CheXpert AUC but can save 1.4x to 4.2x the parameters. Notably, truncating the final block of ResNet18 yields a model that is not significantly different (difference -0.002 ($-0.008, 0.004$)) in CheXpert AUC, but is 4.2x smaller. Truncating the final two blocks of an EfficientNetB0 yields a model that is not significantly different (difference 0.004 ($-0.003, 0.009$)) in CheXpert AUC, but is 4.7x smaller. However, truncating the second block and beyond in each of MNASNet, DenseNet121, and ResNet18 yields models that have statistically significant drops in CheXpert performance.

Model	AUC Change	Times-Smaller
EfficientNetB0	0.00%	1x
EfficientNetB0Minus1	0.15%	1.4x
EfficientNetB0Minus2	-0.45%	4.7x
MNASNet	0.00%	1x
MNASNetMinus1	-0.07%	2.5x
MNASNetMinus2*	-2.30%	11.2x
MNASNetMinus3*	-2.51%	20.0x
MNASNetMinus4*	-6.40%	112.9x
DenseNet121	0.00%	1x
DenseNet121Minus1	-0.04%	1.6x
DenseNet121Minus2*	-1.33%	5.3x
DenseNet121Minus3*	-4.73%	20.0x
ResNet18	0.00%	1x
ResNet18Minus1	0.24%	4.2x
ResNet18Minus2*	-3.70%	17.1x
ResNet18Minus3*	-8.33%	73.8x

Table 2: Efficiency Trade-Off of Truncated Models. Pre-trained models can be truncated without significant decrease in CheXpert AUC. Truncated models with significantly different AUC from the base model are denoted with an asterisk.

Model truncation effectively compresses models performant on CheXpert, making them more parameter efficient while still using pretrained weights to capture the pretraining boost. Parameter efficient models are able to lighten the computational and memory burdens for deployment to low-resource environments such as portable devices. In the clinical setting, the simplicity of our model truncation method encourages its adoption for model compression.

This finding corroborates Raghu et al. [19] and Bressem et al. [3], which show simpler models can achieve performance comparable to more complex models on CheXpert. Our truncated models can use readily-available pretrained weights, which may allow these models to capture the pretraining boost and speed up training. However, we do not study the performance of these truncated models without their pretrained weights.

As an additional benefit, architectures that truncate pooling layers will also produce higher-resolution class activation maps, as shown in Figure 5. The higher-resolution class activation maps (CAMs) may more effectively localize pathologies with little to no decrease in classification performance. In clinical settings, improved explainability through better CAMs may be useful for validating predictions and diagnosing mispredictions. As a result, clinicians may have more trust in models that provide these higher-resolution CAMs.

5 DISCUSSION

In this work, we study the performance and efficiency of ImageNet architectures for chest x-ray interpretation.

Is ImageNet performance correlated with CheXpert? No. We show no statistically significant relationship between ImageNet and CheXpert performance. This finding extends Kornblith et al. [15]—which

Figure 4: Pretraining Boost vs. Model Size. We define pretraining boost as the increase in the average CheXpert AUCs achieved with pretraining vs. without pretraining. Most models benefit significantly from ImageNet pretraining. Smaller models tend to benefit more than larger models (Spearman $\rho = -0.72$).

found a significant correlation between ImageNet performance and transfer performance on typical image classification datasets—to the medical setting of chest x-ray interpretation. This difference could be attributed to unique aspects the chest X-ray interpretation task and data attributes. The chest X-ray interpretation task differs from natural image classification in that (1) disease classification may depend on abnormalities in a small number of pixels, (2) chest X-ray interpretation is a multi-task classification setup, and (3) there are far fewer classes than in many natural image classification datasets. Second, the data attributes for chest X-rays differ from natural image classification in that X-rays are greyscale and have similar spatial structures across images (always either anterior-posterior, posterior-anterior, or lateral).

Does model architecture matter? Yes. For models without pretraining, we find that the choice of architecture family may influence performance more than model size. Our findings extend Raghu et al. [19] beyond the effect of ImageNet weights, since we show that architectures that succeed on ImageNet do not necessarily succeed on medical imaging tasks. A notable finding of our work is that newer architectures generated through search on ImageNet (EfficientNet, MobileNet, MNASNet) underperform older architectures

(DenseNet, ResNet) on CheXpert. This finding suggests that search may have overfit to ImageNet to the detriment of medical task performance, and ImageNet may not be an appropriate benchmark for selecting architectures for medical imaging tasks. Instead, medical imaging architectures could be benchmarked on CheXpert or other large medical datasets. Architectures derived from selection and search on CheXpert and other large medical datasets may be applicable to similar medical imaging modalities including other x-ray studies, or CT scans. Thus architecture search directly on CheXpert or other large medical datasets may allow us to unlock next generation performance for medical imaging tasks.

Does ImageNet pretraining help? Yes. We find that ImageNet pretraining yields a statistically significant boost in performance for chest x-ray classification. Our findings are consistent with Raghu et al. [19], who find no pretraining boost on ResNet50 and InceptionV3, but we find pretraining does boost performance for 12 out of 16 architectures. Our findings extend He et al. [11]—who find models without pretraining had comparable performance to models pretrained on ImageNet for object detection and image segmentation of natural images—to the medical imaging setting.

Figure 5: Comparison of Class Activation Maps Among Truncated Model Family. CAMs yielded by models, from left to right, DenseNet121, DenseNet121Minus1, and DenseNet121Minus2. Displays frontal chest X-ray demonstrating Atelectasis (top) and Edema (bottom). Further truncated models more effectively localize the Atelectasis, as well as tracing the hila and vessel branching for Edema.

Future work may investigate the relationship between network architectures and the impact of self-supervised pre-training for chest x-ray interpretation as has recently been developed by Azizi et al. [2], Sowrirajan et al. [26], Sriram et al. [28].

Can models be smaller? Yes. We find that by truncating final blocks of ImageNet-pretrained architectures, we can make models 3.25x more parameter-efficient on average without a statistically significant drop in performance. This method preserves the critical components of architecture design while cutting its size. This observation suggests model truncation may be a simple method to yield lighter models, using ImageNet pretrained weights to boost CheXpert performance. In the clinical setting, truncated models may provide value through improved parameter-efficiency and higher resolution CAMs. This change may enable deployment to low-resource clinical environments and further develop model trust through improved explainability.

In closing, our work contributes to the understanding of the transfer performance and parameter efficiency of ImageNet models for chest X-ray interpretation. We hope that our new experimental evidence about the relation of ImageNet to medical task performance will shed light on potential future directions for progress.

REFERENCES

- [1] Ioannis D Apostolopoulos and Tzani A Mpesiana. 2020. Covid-19: automatic detection from x-ray images utilizing transfer learning with convolutional neural networks. *Physical and Engineering Sciences in Medicine* (2020), 1.
- [2] Shekoofeh Azizi, Basil Mustafa, Fiona Ryan, Zachary Beaver, Jan Freyberg, Jonathan Deaton, Aaron Loh, Alan Karthikesalingam, Simon Kornblith, Ting Chen, Vivek Natarajan, and Mohammad Norouzi. 2021. Big Self-Supervised Models Advance Medical Image Classification. arXiv:2101.05224 [eess.IV]
- [3] Keno K. Bressen, Lisa Adams, Christoph Erxleben, Bernd Hamm, Stefan Niehues, and Janis Vahldiek. 2020. Comparing Different Deep Learning Architectures for Classification of Chest Radiographs. arXiv:2002.08991 [cs.LG]
- [4] Remi Cadene. 2018. pretrainedmodels 0.7.4. <https://pypi.org/project/pretrainedmodels/>.
- [5] S. Chen and Q. Zhao. 2019. Shallowing Deep Networks: Layer-Wise Pruning Based on Feature Representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 12 (2019), 3048–3056.
- [6] Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. 2017. A Survey of Model Compression and Acceleration for Deep Neural Networks. CoRR abs/1710.09282 (2017). arXiv:1710.09282 <http://arxiv.org/abs/1710.09282>
- [7] François Chollet. 2016. Xception: Deep Learning with Depthwise Separable Convolutions. CoRR abs/1610.02357 (2016). arXiv:1610.02357 <http://arxiv.org/abs/1610.02357>
- [8] Jeffrey De Fauw, Joseph R. Ledsam, Bernardino Romera-Paredes, Stanislaw Nikolov, Nenad Tomasev, Sam Blackwell, Harry Askham, Xavier Glorot, Brendan O’Donoghue, Daniel Visentin, George van den Driessche, Balaji Lakshminarayanan, Clemens Meyer, Faith Mackinder, Simon Bouton, Kareem Ayoub, Reena Chopra, Dominic King, Alan Karthikesalingam, Cian O. Hughes, Rosalind Raine, Julian Hughes, Dawn A. Sim, Catherine Egan, Adnan Tufail, Hugh Montgomery, Demis Hassabis, Geraint Rees, Trevor Back, Peng T. Khaw, Mustafa Suleyman, Julien Cornebise, Pearse A. Keane, and Olaf Ronneberger. 2018. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature Medicine* 24, 9 (01 Sep 2018), 1342–1350. <https://doi.org/10.1038/s41591-018-0107-6>
- [9] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 248–255.
- [10] Andre Esteve, Brett Kuperl, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542, 7639 (2017), 115–118. <https://doi.org/10.1038/nature21056>
- [11] Kaiming He, Ross B. Girshick, and Piotr Dollár. 2018. Rethinking ImageNet Pre-training. CoRR abs/1811.08883 (2018). arXiv:1811.08883 <http://arxiv.org/abs/1811.08883>
- [12] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the Knowledge in a Neural Network. arXiv:1503.02531 [stat.ML]
- [13] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghighi, Robyn L. Ball, Katie S. Shpanskaya, Jayne Seekins, David A. Mong, Safwan S. Halabi, Jesse K. Sandberg, Ricky Jones, David B. Larson, Curtis P. Langlotz, Bhavik N. Patel, Matthew P. Lungren, and Andrew Y. Ng. 2019. CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison. CoRR abs/1901.07031 (2019). arXiv:1901.07031 <http://arxiv.org/abs/1901.07031>
- [14] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. 2014. Speeding up Convolutional Neural Networks with Low Rank Expansions. CoRR abs/1405.3866 (2014). arXiv:1405.3866 <http://arxiv.org/abs/1405.3866>
- [15] Simon Kornblith, Jonathon Shlens, and Quoc V. Le. 2018. Do Better ImageNet Pre-training Transfer Better? CoRR abs/1805.08974 (2018). arXiv:1805.08974 <http://arxiv.org/abs/1805.08974>
- [16] Feng Li, Zheng Liu, Hua Chen, Minshan Jiang, Xuedian Zhang, and Zhizheng Wu. 2019. Automatic Detection of Diabetic Retinopathy in Retinal Fundus Photographs Based on Deep Learning Algorithm. *Translational Vision Science & Technology* 8, 6 (11 2019), 4–4. <https://doi.org/10.1167/tvst.8.6.4> arXiv:https://arxiv.org/abs/1901.07031
- [17] Akinoori Mitani, Abigail Huang, Subhashini Venugopalan, Greg S. Corrado, Lily Peng, Dale R. Webster, Naama Hammel, Yun Liu, and Avinash V. Varadarajan. 2020. Detection of anaemia from retinal fundus images via deep learning. *Nature Biomedical Engineering* 4, 1 (01 Jan 2020), 18–27. <https://doi.org/10.1038/s41551-019-0487-z>
- [18] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035. <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [19] Maithra Raghu, Chiyuan Zhang, Jon M. Kleinberg, and Samy Bengio. 2019. Transfusion: Understanding Transfer Learning with Applications to Medical Imaging. CoRR abs/1902.07208 (2019). arXiv:1902.07208 <http://arxiv.org/abs/1902.07208>
- [20] Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Yi Ding, Aarti Bagul, Curtis Langlotz, Katie S. Shpanskaya, Matthew P. Lungren, and Andrew Y. Ng. 2017. CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning. CoRR abs/1711.05225 (2017). arXiv:1711.05225 <http://arxiv.org/abs/1711.05225>
- [21] Pranav Rajpurkar, Anirudh Joshi, Anuj Pareek, Phil Chen, Amirhossein Kiani, Jeremy Irvin, Andrew Y. Ng, and Matthew P. Lungren. 2020. CheXpedition: investigating generalization challenges for translation of chest x-ray algorithms to the clinical setting. arXiv preprint arXiv:2002.11379 (2020).

- [22] Pranav Rajpurkar, Anirudh Joshi, Anuj Pareek, Andrew Y. Ng, and Matthew P. Lungren. 2021. CheXternal: Generalization of Deep Learning Models for Chest X-ray Interpretation to Photos of Chest X-rays and External Clinical Settings. [arXiv:2102.08660](https://arxiv.org/abs/2102.08660) [eess.IV]
- [23] Pranav Rajpurkar, Chloe O'Connell, Amit Schechter, Nishit Asnani, Jason Li, Amirhossein Kiani, Robyn L Ball, Marc Mendelson, Gary Maartens, Daniël J van Hoving, et al. 2020. CheXaid: deep learning assistance for physician diagnosis of tuberculosis using chest x-rays in patients with HIV. *NPJ digital medicine* 3, 1 (2020), 1–8.
- [24] Youngmin Ro and Jin Young Choi. 2020. Layer-wise Pruning and Auto-tuning of Layer-wise Learning Rates in Fine-tuning of Deep Networks. [arXiv:2002.06048](https://arxiv.org/abs/2002.06048) [cs.CV]
- [25] Ramprasaath R. Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra. 2016. Grad-CAM: Why did you say that? Visual Explanations from Deep Networks via Gradient-based Localization. *CoRR* abs/1610.02391 (2016). [arXiv:1610.02391](https://arxiv.org/abs/1610.02391) [http://arxiv.org/abs/1610.02391](https://arxiv.org/abs/1610.02391)
- [26] Hari Sowrirajan, Jingbo Yang, Andrew Y. Ng, and Pranav Rajpurkar. 2020. MoCo Pretraining Improves Representation and Transferability of Chest X-ray Models. [arXiv:2010.05352](https://arxiv.org/abs/2010.05352) [cs.CV]
- [27] Suraj Srinivas and R. Venkatesh Babu. 2015. Data-free parameter pruning for Deep Neural Networks. *CoRR* abs/1507.06149 (2015). [arXiv:1507.06149](https://arxiv.org/abs/1507.06149) [http://arxiv.org/abs/1507.06149](https://arxiv.org/abs/1507.06149)
- [28] Anuroop Sriram, Matthew Muckley, Koustuv Sinha, Farah Shamout, Joelle Pineau, Krzysztof J. Geras, Lea Azour, Yindalon Aphinyanaphongs, Nafissa Yakubova, and William Moore. 2021. COVID-19 Deterioration Prediction via Self-Supervised Representation Learning and Multi-Image Prediction. [arXiv:2101.04909](https://arxiv.org/abs/2101.04909) [cs.CV]
- [29] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M. Summers. 2017. ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases. *CoRR* abs/1705.02315 (2017). [arXiv:1705.02315](https://arxiv.org/abs/1705.02315) [http://arxiv.org/abs/1705.02315](https://arxiv.org/abs/1705.02315)
- [30] Ross Wightman. 2020. timm 0.2.1. <https://pytorch.org/project/timm/>.
- [31] Li Zhang, Mengya Yuan, Zhen An, Xiangmei Zhao, Hui Wu, Haibin Li, Ya Wang, Beibei Sun, Huijun Li, Shibin Ding, Xiang Zeng, Ling Chao, Pan Li, and Weidong Wu. 2020. Prediction of hypertension, hyperglycemia and dyslipidemia from retinal fundus photographs via deep learning: A cross-sectional study of chronic diseases in central China. *PLOS ONE* 15, 5 (05 2020), 1–11. <https://doi.org/10.1371/journal.pone.0233166>