




Contents lists available at ScienceDirect

# Computer Methods and Programs in Biomedicine

journal homepage: [www.sciencedirect.com/journal/computer-methods-and-programs-in-biomedicine](http://www.sciencedirect.com/journal/computer-methods-and-programs-in-biomedicine)



## Why does my medical AI look at pictures of birds? Exploring the efficacy of transfer learning across domain boundaries

Frederic Jonske<sup>a,\*</sup> , Moon Kim<sup>a</sup> , Enrico Nasca<sup>a</sup> , Janis Evers<sup>a</sup>, Johannes Haubold<sup>b</sup>, René Hosch<sup>a,b</sup>, Felix Nensa<sup>a,b</sup>, Michael Kamp<sup>a,c,d</sup> , Constantin Seibold<sup>a</sup>, Jan Egger<sup>a,e</sup> , Jens Kleesiek<sup>a,e,f</sup> 

<sup>a</sup> Institute of AI in Medicine (IKIM), University Medicine Essen (AöR), University Duisburg-Essen, Germany (Address: Institut für künstliche Intelligenz in der Medizin (IKIM), Girardetstr. 2 45131 Essen, NRW, Germany)

<sup>b</sup> Institute of Diagnostic and Interventional Radiology and Neuroradiology, University Medicine Essen (AöR), Germany

<sup>c</sup> Institute for Neuroinformatics, Ruhr University Bochum, Germany

<sup>d</sup> Department of Data Science & AI, Monash University, Australia

<sup>e</sup> Cancer Research Center Cologne Essen (CCCE), University Medicine Essen (AöR), Germany

<sup>f</sup> Department of Physics, TU Dortmund University, Dortmund, Germany

### ARTICLE INFO

#### Keywords:

Deep learning  
Domain adaptation  
Foundation model  
Transfer learning  
Unsupervised Pretraining

### ABSTRACT

**Purpose:** In medical deep learning, models not trained from scratch are typically fine-tuned based on ImageNet-pretrained models. We posit that pretraining on data from the domain of the downstream task should almost always be preferable.

**Materials and methods:** We leverage RadNet-12M and RadNet-1.28M, datasets containing >12 million/1.28 million acquired CT image slices from 90,663 individual scans, and explore the efficacy of self-supervised, contrastive pretraining on the medical and natural image domains. We compare the respective performance gains for five downstream tasks. For each experiment, we report accuracy, AUC, or DICE score and uncertainty estimations based on four separate runs. We quantify significance using Welch's *t*-test. Finally, we perform feature space analysis to characterize the nature of the observed performance gains.

**Results:** We observe that intra-domain transfer (RadNet pretraining and CT-based tasks) compares favorably to cross-domain transfer (ImageNet pretraining and CT-based tasks), generally achieving comparable or improved performance –  $\Delta = +0.44\%$  ( $p = 0.541$ ) when fine-tuned on RadNet-1.28M,  $\Delta = +2.07\%$  ( $p = 0.025$ ) when linearly evaluating on

RadNet-1.28M, and  $\Delta = +1.63\%$  ( $p = 0.114$ ) when fine-tuning on 1 % of RadNet-1.28M data. This intra-domain advantage extends to LiTS 2017, another CT-based dataset, but not to other medical imaging modalities. A corresponding intra-domain advantage was also observed for natural images. Outside the CT image domain, ImageNet-pretrained models generalized better than RadNet-pretrained models.

We further demonstrate that pretraining on medical images yields domain-specific features that are preserved during fine-tuning, and which correspond to macroscopic image properties and structures.

**Conclusion:** We conclude that intra-domain pretraining generally outperforms cross-domain pretraining, but that very narrow domain definitions apply. Put simply, pretraining on CT images instead of natural images yields an advantage when fine-tuning on CT images, and only on CT images. We further conclude that ImageNet pretraining remains a strong baseline, as well as the best choice for pretraining if only insufficient data from the target domain is available. Finally, we publish our pretrained models and pretraining guidelines as a baseline for future research.

\* Corresponding author at: Hohenstein 120, 42283 Wuppertal, NRW, Germany.

E-mail addresses: [frederic.jonske@uk-essen.de](mailto:frederic.jonske@uk-essen.de) (F. Jonske), [moon.kim@uk-essen.de](mailto:moon.kim@uk-essen.de) (M. Kim), [enrico.nasca@uk-essen.de](mailto:enrico.nasca@uk-essen.de) (E. Nasca), [janis.evers@stud.uni-due.de](mailto:janis.evers@stud.uni-due.de) (J. Evers), [johannes.haubold@uk-essen.de](mailto:johannes.haubold@uk-essen.de) (J. Haubold), [rene.hosch@uk-essen.de](mailto:rene.hosch@uk-essen.de) (R. Hosch), [felix.nensa@uk-essen.de](mailto:felix.nensa@uk-essen.de) (F. Nensa), [michael.kamp@uk-essen.de](mailto:michael.kamp@uk-essen.de) (M. Kamp), [constantin.seibold@uk-essen.de](mailto:constantin.seibold@uk-essen.de) (C. Seibold), [jan.egger@uk-essen.de](mailto:jan.egger@uk-essen.de) (J. Egger), [jens.kleesiek@uk-essen.de](mailto:jens.kleesiek@uk-essen.de) (J. Kleesiek).

<https://doi.org/10.1016/j.cmpb.2025.108634>

Received 6 August 2024; Received in revised form 23 January 2025; Accepted 31 January 2025

Available online 31 January 2025

0169-2607/© 2025 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

## 1. Introduction

Is there a guaranteed-to-be-best dataset to pretrain on, given a task? One could argue that any task's dataset is naturally best represented by features learned during pretraining on itself, or on a dataset from the same underlying distribution. One could alternatively argue that any image should be representable by a generic set of features, no matter where that set is learned. After all, humans also perform transfer learning by first being "pretrained" on natural images while growing up, before "getting fine-tuned" on medical images to become radiologists. In this paper, we will systematically explore this question regarding medical and natural image domains, using contrastive, self-supervised pretraining.

In the literature, several works explore transfer learning applications from medical pretraining to medical downstream tasks. Some have reported that pretraining on one domain improves fine-tuning on another. Others, however, have found that unsupervised pretraining offered no substantial improvements. This implies that both lines of reasoning provided by us have at least some merits.

Cherti & Jitsev [1] explored the scalability of supervised pretraining and transfer learning in medical-medical, natural-natural and cross-domain scenarios. Their findings indicated that in some cases, upscaling the available pretraining data led to proportionally increased performance, while for few-shot tasks the effect disappeared. They additionally reported that models pretrained on ImageNet-21 k [2] outperformed models pretrained on the largest available X-ray datasets when fine-tuning on X-ray image downstream tasks. Notably, the combined number of labeled X-rays available for this comparison was about 16 times less than the images contained in ImageNet-21 k (approx. 873k vs. approx. 14 M).

Leveraging around 100 million medical images from multiple modalities for self-supervised pretraining, Ghesu et al. reported that their self-supervised algorithm, together with the significant dataset upscaling, yielded a significant advantage over both training from scratch and pretraining with other methods [3].

Mei et al. created RadImageNet [4], a dataset composed of 1.4 million medical images. They reported that models trained on RadImageNet achieved superior performance on medical tasks, compared to ImageNet-pretrained models. They further observed a correlation between the size of the dataset in the downstream task and the transfer learning performance, noting that RadImageNet yielded greater improvements on smaller datasets.

Azizi et al. created REMEDIS [5], a multi-supervision level approach for transfer learning on medical tasks. REMEDIS combines large-scale supervised learning on natural images with self-supervised learning on medical images, and achieves significant performance gains on 15 medical tasks, compared to supervised baselines.

While noting that ImageNet-pretraining was often a *de facto* method for medical image classification problems, Raghu et al. observed across a range of different model architectures and medical imaging tasks, that pretraining with ImageNet yielded no significant performance boost and that meaningful feature reuse did not always occur after pretraining [6].

Newell & Deng experimentally probed the efficacy of self-supervised training using synthetic data and various pretraining algorithms [7]. They observed an anti-proportional correlation between pretraining utility and the number fine-tuning images and noted that in some cases, pretraining offered only accelerated convergence or even no performance gain.

In a study on cross-domain generalization [8], Cohen et al. found that many of their models could perform well after fine-tuning on the same downstream task, yet disagree in their predictions. Moreover, models that were in strong agreement nonetheless performed poorly, implying that generalization to other tasks is often a nuanced process.

A study by Gururangan et al. [9] explored intra-domain transfer learning on Natural Language Processing (NLP) tasks, using multi-phased pretraining on different datasets. They found that

spending at least some time training on data from the domain of the downstream task led to consistent performance gains.

Based on these previous findings and our initial lines of reasoning, we formulate the following hypotheses:

- I) When applying transfer learning, any difference between the pretraining and fine-tuning task domains produces a *generalization gap* [10]. This occurs because the two underlying data distributions, and consequently learned feature subsets, are likely to be different at least to some degree [11,12]. Intuitively, we expect this gap to shrink or disappear in an intra-domain transfer learning scenario, since the data distributions are highly similar or identical. In essence, we propose that **"there is no better data to train medical questions on than medical data"**.
- II) The generalization gap effect is not symmetrical. A model trained on a specific pretraining dataset may learn a superset of the features it would learn from a different dataset. While this is most likely not immediately true for every downstream task, it may well be true for a significant fraction of them. We posit that **"some pretraining datasets are principally more useful than others"**.

Investigating these hypotheses not only provides fundamental insights into the nature of fine-tuning, domain adaptation and transfer learning, but will, more importantly, yield practical strategies to improve our medical AI models.

Our main contributions are as follows:

- To validate these hypotheses, we leverage RadNet-1.28M and RadNet-12M, two pretraining datasets comprising over 12 million CT image slices from our local hospital, the largest dataset on which this kind of analysis has been performed to date. We also leverage four publicly available datasets from the natural and medical image domain.
- We design an experimental setup (cf. Fig. 1 and the Methods section) to compare unsupervised pretraining efficacy depending on the fine-tuning dataset domain. We additionally extend previous work by exploring the dependency of the downstream performance on task complexity, and in a linear evaluation vs fine-tuning scenario.
- We present our results and uncertainty estimations, discuss them with respect to our hypotheses, and perform novel feature space analysis for our models. Here, we demonstrate the existence of unique, explainable image features learned only during intra-domain pretraining that are preserved during fine-tuning.
- Finally, we discuss practical guidelines for the choice of pretraining datasets, depending on downstream tasks, and publish our code framework and all pretrained/fine-tuned models as a baseline for future research or the extension of our work.

## 2. Methods

For the purposes of this paper, we define a domain as the collection of all data points that satisfy some condition. For example, a very general domain would be the domain of medical images, containing all images with any medical context. A smaller, more distinct domain would be the CT image domain, or the domain of abdominal CTs. Each of the datasets below comes from one of these domains.

### 2.1. Datasets and tasks

#### 2.1.1. RadNet-X

The RadNet-X datasets are composed purely of CT slices. We opted for a single modality because we intend to test our hypotheses on a narrow, clearly defined domain. The CT data stems from the University Hospital Essen, Germany, and was retrospectively collected for this study. The criteria are as follows: There were no exclusion criteria based

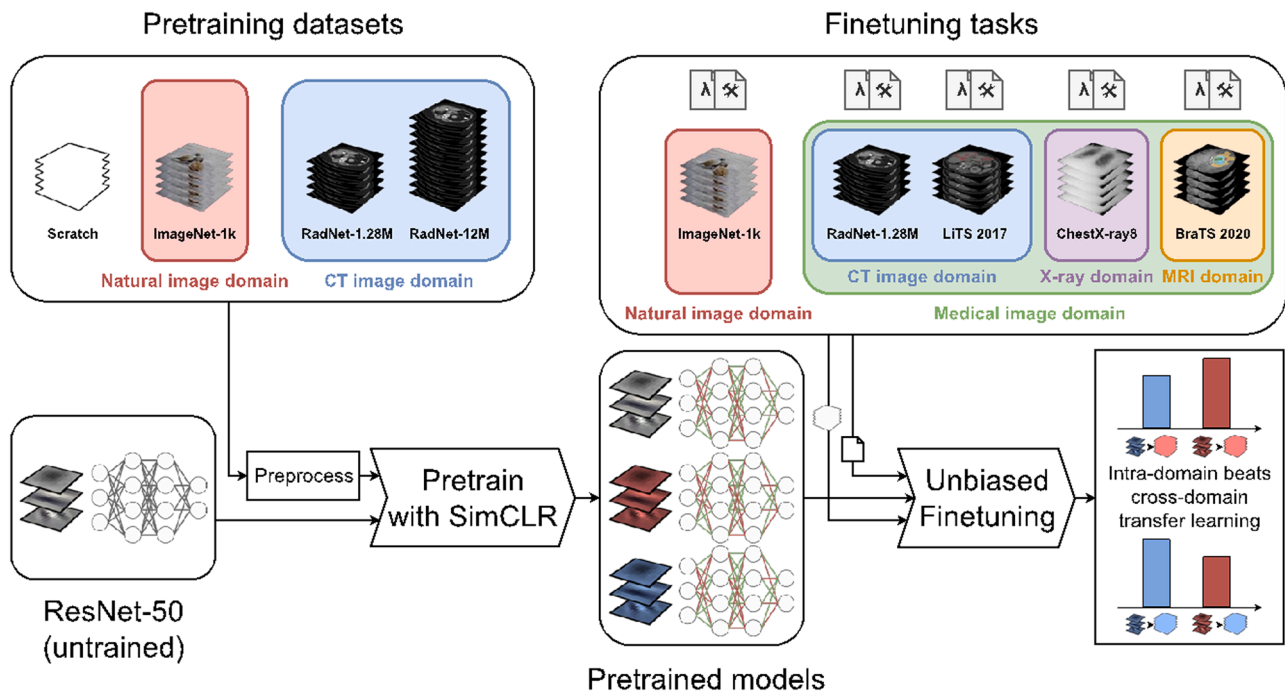


Fig. 1. Schematic visualization of the transfer learning experiment.

During the transfer learning experiment, models are pretrained on medical, natural, or no images (denoted by five stacked images drawn from the actual dataset, top left). The value of the pretraining in terms of performance gains is measured across several downstream tasks (top right), which also come from the medical or natural image domain. Each downstream task has a hyperparameter set and build instructions (denoted by the  $\lambda$  and symbols) that are shared between executions. This makes extending the framework straightforward and more easily reproducible.

on patients or dates. Scans with fewer than 30 slices in the z-direction were excluded to filter for high resolution scans. Slices with  $<25\%$  coverage were removed. This process yielded 12,034,617 slices (approx. 84.8 % the size of ImageNet-21k) from 90,663 CTs. Data used by Koitka et al. [13] and Jonske et al. [14] was used in this study. <https://github.com/UMEssen/DicomDeidentify> was used to deidentify all data. The slices were automatically labeled based on metadata, dividing them into 7 anatomical classes – “head”, “thorax”, “abdomen”, “extremities”, “pelvis”, “spine”, and “other”. From these 12 million slices, two datasets were created:

- The larger dataset, called RadNet-12M, contains all 12 million data. It is used as one of the pretraining options.
- RadNet-1.28M contains 1.28 million data and is thus the size of ImageNet-1k. It is a subset of RadNet-12M. RadNet-1.28M maintains the same class ratio as RadNet-12M, except for instances of the “other” class, which are excluded to improve fine-tuning stability. RadNet-1.28M is used both as a pretraining and as a fine-tuning dataset in our experiments.

In both datasets, slices from the same volume can only appear in either the train, validation, or test set, but not in multiple ones.

Due to human anatomy, neighboring CT slices are correlated, effectively reducing the amount of information contained in the dataset compared to one in which every image slice comes from a different scan. To facilitate a fair comparison with ImageNet, we offset this effect by interpolating between the results for RadNet-1.28M and –12M, thus simulating a larger dataset. An upper and lower bound for the factor by which RadNet-1.28M must be upscaled are calculated (see supplementary materials) as  $S_{LB} = 1.136$  and  $S_{UB} = 3.807$ . As discussed in the supplementary materials, we consider RadNet-1.28M (UB-adjusted) as the closest radiological imaging analog to ImageNet-1k.

While we cannot make the dataset itself publicly available, the pre-trained model weights for both RadNet-X datasets are made publicly

available (consent waived and permission granted by the responsible IRB, cn. 20–9745-BO).

### 2.1.2. ImageNet-1k

ImageNet-1k [2] consists of 1.28 million RGB images of 1000 classes. The associated data domain for these images is called “natural images”. Human-created labels exist for every point of data and have a one-to-one relation with images. We make an even, stratified split of the official validation set to obtain a validation and test set.

### 2.1.3. LiTS 2017

LiTS (Liver Tumor Segmentation) 2017 [15] is a segmentation dataset composed of abdominal CT images. The data comes from 131 distinct patients and totals 41,561 individual 2D slices, all of which are professionally segmented into background, liver, and tumor regions. The official validation set is evenly split to obtain a validation and test set. Slices from the same scan do not appear in different sets.

### 2.1.4. BraTS 2020

BraTS (Brain Tumor Segmentation) 2020 [16] is a segmentation dataset composed of cranial MRT scans from 369 glioma patients. For each patient, T1, T1ce, T2, and FLAIR modalities are available. Tumor regions are divided into Tumor Core (TC), Active Tumor (AT), and Whole Tumor (WT). We remove empty slices from the dataset. We use 30 patients each from the official training dataset to construct a validation and test set. Slices from the same scan do not appear in different sets.

### 2.1.5. ChestX-Ray8

ChestX-Ray8 [17] is a medical imaging dataset containing 108,948 frontal-view X-ray images from 32,717 unique patients. Classification labels for each X-ray were automatically acquired via text mining. We construct a training and validation set from the official “train\_val” set using a 90–10 split. To avoid bias from an uneven split, training and

validation sets are randomized between runs of an experiment.

## 2.2. Experimental setup

To test our hypotheses, we combine all four pretraining options (training from scratch, ImageNet-1k, RadNet-1.28M, and RadNet-12M) and all five fine-tuning tasks (ImageNet classification, RadNet-1.28M classification, ChestX-Ray8 multi-label classification, LiTS 2017 segmentation, and BraTS 2020 segmentation) in several different settings.

For implementation details, consult the supplementary materials or our documentation at <https://github.com/TIO-IKIM/Transfer-learning-across-domain-boundaries>.

### 2.2.1. Experiment 1 (Task dependency)

According to Hypothesis I, we expect to see a generalization gap when we the pretraining and fine-tuning dataset domains are different. If Hypothesis I is correct, we expect RadNet-pretrained models to outperform ImageNet-pretrained models on the RadNet and LiTS 2017 downstream tasks. Simultaneously, we test Hypothesis II - if the relative performance loss in the cross-domain transfer scenarios is asymmetrical, this is a strong indication that ImageNet is generally better suited for transfer learning than RadNet.

### 2.2.2. Experiment 2 (Fine-tuning vs. linear evaluation)

During regular fine-tuning, such as in Experiment 1, all parameters of a pretrained model are learnable and will be adjusted. When all parameters except those of the last layer of the model are frozen (cannot be updated), that training process is termed “linear evaluation”. This effectively increases the difficulty of the task and enforces the reuse of pretrained features. For segmentation tasks, the added U-Net decoder is randomly initialized at the start of the linear evaluation and thus needs to be trained to achieve any meaningful results. In this case, we allow

training of the weights of half the model. The expectation according to Hypothesis I is that any observed intra-domain advantage from Experiment 1 becomes more pronounced, because the model is forced to rely entirely on the features learned during pretraining.

### 2.2.3. Experiment 3 (Task complexity)

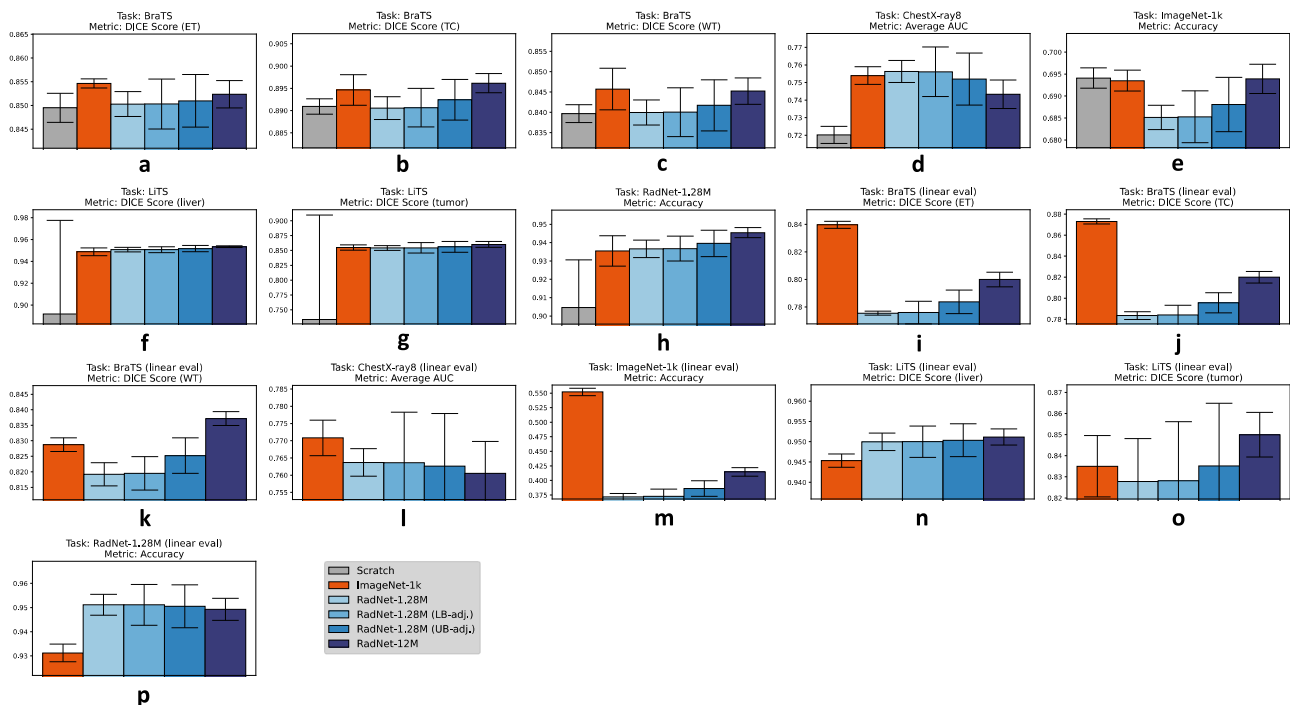
We define the complexity of a task as the ratio between number of labeled examples and number of target classes. We increase this complexity by decreasing the training dataset size to 50 %, 10 %, 1 %, or 0.1 % in a stratified manner, or we decrease it by removing classes. According to Hypothesis I, we expect that in this low-data scenario, pretraining features learned on the task’s domain should be significantly more useful than those learned in a different domain.

## 2.3. Reported metrics

For classification tasks, we report the accuracy (ImageNet, RadNet) and average AUC (ChestX-Ray8), while for segmentation tasks (LiTS 2017, BraTS 2020) we report DICE scores per class and per slice. We provide estimates of the aleatoric uncertainty of our results based on four separate runs of the same experiment, each using different seeds. Finally, we report relative performance comparisons and their significance based on Welch’s *t*-test [18].

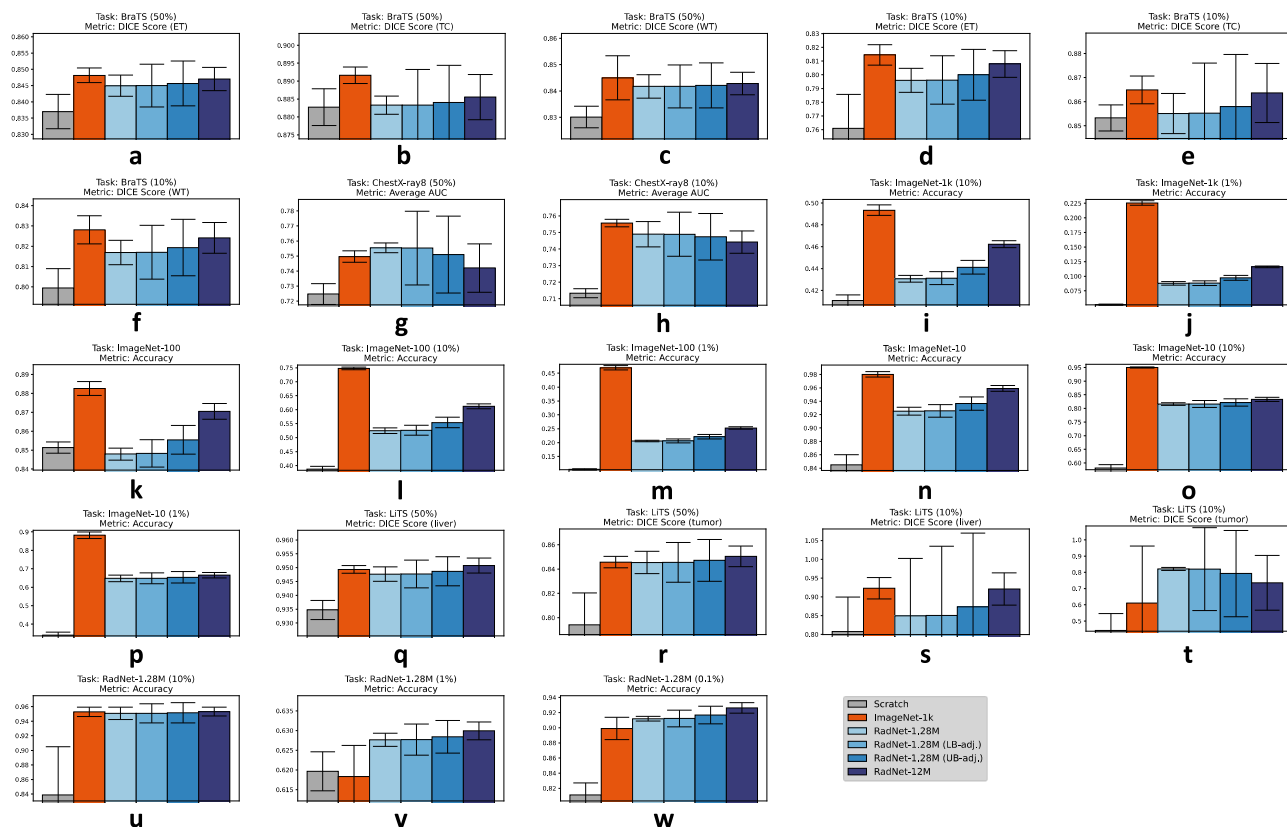
## 3. Results

Firstly, we observe that pretraining on any dataset yields some degree of improved performance in almost every fine-tuning scenario we tested (cf. Figs. 2 and 3, or Supplementary Tables ST1–3). We further observe that intra-domain transfer learning (e.g. RadNet→RadNet) generally offers equal or slightly improved performance compared to cross-domain transfer learning (e.g. ImageNet→RadNet), although this



**Fig. 2.** Intra-domain outperforms cross-domain transfer learning.

Performance across four separate runs for every combination of pretraining dataset and fine-tuning task of Experiments #1 and #2 are on display: a, Fine-tuning on BraTS 2020 (Enhancing Tumor). b, Fine-tuning on BraTS 2020 (Tumor Core). c, Fine-tuning on BraTS 2020 (Whole Tumor). d, Fine-tuning on ChestX-Ray8. e, Fine-tuning on ImageNet-1k. f, Fine-tuning on LiTS 2017 (Liver). g, Fine-tuning on LiTS 2017 (Lesion). h, Fine-tuning on RadNet-1.28M. i, Linear evaluation on BraTS 2020 (Enhancing Tumor). j, Linear evaluation on BraTS 2020 (Tumor Core). k, Linear evaluation on BraTS 2020 (Whole Tumor). l, Linear evaluation on ChestX-Ray8. m, Linear evaluation on ImageNet-1k. n, Linear evaluation on LiTS 2017 (Liver). o, Linear evaluation on LiTS 2017 (Lesion). p, Linear evaluation on RadNet-1.28M.



**Fig. 3.** The intra-domain advantage scales with task complexity.

Performance across four separate runs for every combination of pretraining dataset, fine-tuning task of Experiment #3 are on display: **a**, Fine-tuning on BraTS 2020 (50 %, Enhancing Tumor). **b**, Fine-tuning on BraTS 2020 (50 %, Tumor Core). **c**, Fine-tuning on BraTS 2020 (50 %, Whole Tumor). **d**, Fine-tuning on BraTS 2020 (10 %, Enhancing Tumor). **e**, Fine-tuning on BraTS 2020 (10 %, Tumor Core). **f**, Fine-tuning on BraTS 2020 (10 %, Whole Tumor). **g**, Fine-tuning on ChestX-Ray8 (50 %). **h**, Fine-tuning on ChestX-Ray8 (10 %). **i**, Fine-tuning on ImageNet-1k (10 %). **j**, Fine-tuning on ImageNet-1k (1 %). **k**, Fine-tuning on ImageNet-100. **l**, Fine-tuning on ImageNet-100 (10 %). **m**, Fine-tuning on ImageNet-100 (1 %). **n**, Fine-tuning on ImageNet-10. **o**, Fine-tuning on ImageNet-10 (10 %). **p**, Fine-tuning on ImageNet-10 (1 %). **q**, Fine-tuning on LiTS 2017 (50 %, Liver). **r**, Fine-tuning on LiTS 2017 (50 %, Tumor). **s**, Fine-tuning on LiTS 2017 (10 %, Liver). **t**, Fine-tuning on LiTS 2017 (10 %, Tumor). **u**, Fine-tuning on RadNet-1.28M (10 %). **v**, Fine-tuning on RadNet-1.28M (1 %). **w**, Fine-tuning on RadNet-1.28M (0.1 %).

is not unambiguously observed everywhere. If the complexity of the task is increased by reducing the number of images per class during fine-tuning, or linear evaluation is performed, this intra-domain advantage grows substantially, both in absolute terms and in statistical significance. The relative performance of each pairing of pretraining datasets, as well as the corresponding significance can be found in the supplementary materials (Figs. SF1–3).

We further note that the observed intra-domain advantage is not symmetrical - for example, pretraining on RadNet (abbreviated *R-1.28M*) instead of ImageNet (abbreviated *I-1k*), followed by linear evaluation on RadNet, offers a statistically less significant and smaller relative performance gain than the reverse case ( $\Delta_{R-1.28M} = +0.44\%$ ,  $p = 0.541$  vs.  $\Delta_{I-1k} = +0.79\%$ ,  $p = 0.231$ ). This asymmetry appears to intensify substantially during linear evaluation ( $\Delta_{R-1.28M}(\text{lineval}) = +2.07\%$ ,  $p = 0.025$  vs.  $\Delta_{I-1k}(\text{lineval}) = +42.92\%$ ,  $p < 0.001$ ) or for increasing task complexity ( $\Delta_{R-1.28M(1\%)} = +1.63\%$ ,  $p = 0.114$  vs.  $\Delta_{I-1k(1\%)} = +131.62\%$ ,  $p < 0.001$ ). Because of this asymmetry, the ImageNet-pretrained model displays a competitive performance in most experiments.

Notably, the observed intra-domain advantage does not apply across all medical imaging tasks - RadNet-1.28M pretraining yields slightly superior or comparable performance only on RadNet and LiTS ( $\Delta_{R-1.28M} = +0.30\%$ ,  $p = 0.320$ ,  $\#x0394_{R-1.28M} = +0.53\%$ ,  $p = 0.119$ ) - both CT imaging datasets - but inferior performance on BraTS, an MRI dataset, for example.

Additionally, we note that even between target classes in the same task, different behaviors can be observed. For example, RadNet-

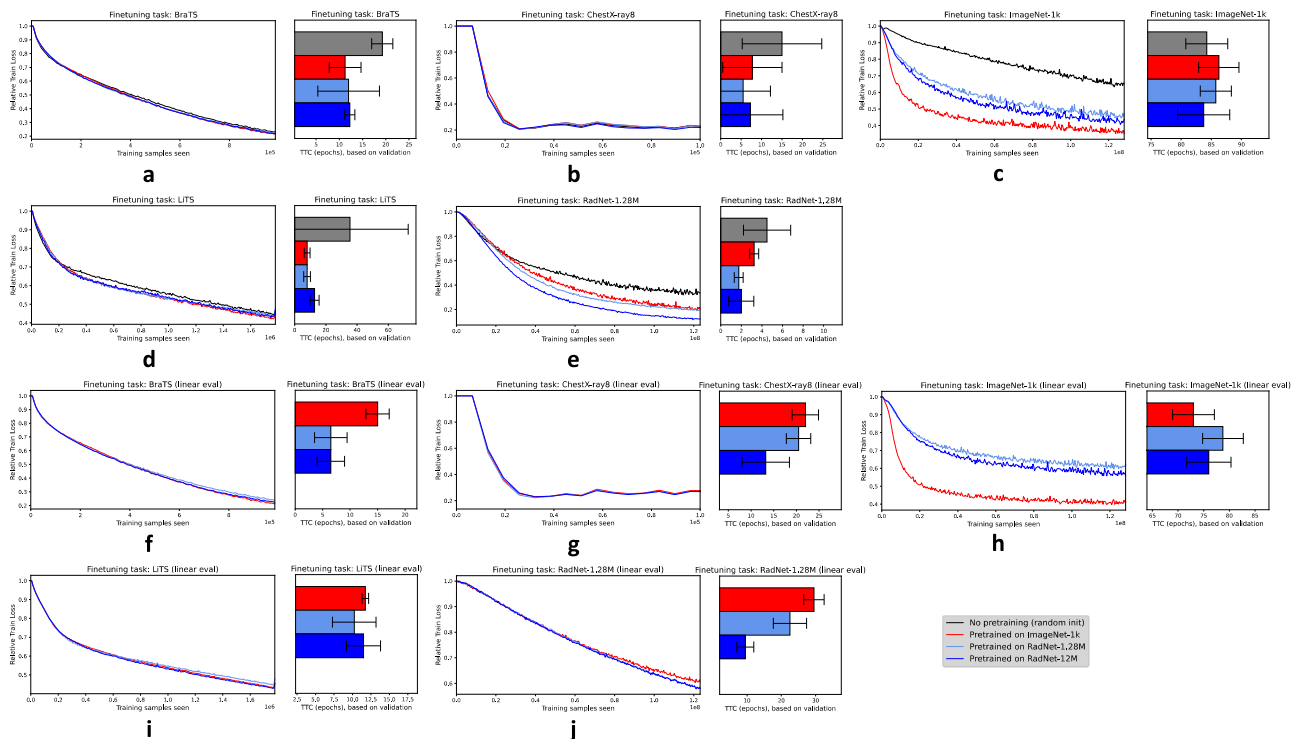
pretrained models slightly outperformed ImageNet-pretrained models in liver segmentation, but not tumor segmentation, where RadNet-1.28M (UB-adjusted) performs similarly to ImageNet-1k. This effect is conserved when linearly evaluating, suggesting that the effect is not simply a statistical coincidence. Finally, we observe that an increase in the scale of available pretraining data almost universally significantly improved transfer performance of RadNet-pretrained models, often even beyond ImageNet performance (cf. Figs. 2 and 3, or Supplementary Tables ST1–3).

#### 4. Discussion

From our results, we conclude that cross-domain transfer indeed degrades performance compared to intra-domain transfer learning in accordance with Hypothesis I, particularly in few-shot or linear evaluation scenarios. For CT-based tasks, this intra-domain advantage is relatively small, compared to natural image-based tasks. In order to observe it with high statistical significance, multiple runs of the same experiment can be necessary. However, despite similar test-time performance, the train-time loss falls at markedly different speeds (cf. Fig. 4) between intra-domain, cross-domain and no pretraining. This suggests that the proposed intra-domain advantage, while probably extant, is suppressed. This explanation is also corroborated by the fact that the expected intra-domain advantage for both data domains is visible and amplified substantially for most few-shot fine-tuning and linear evaluation scenarios.

The asymmetrical generalization gap we observe follows an intuitive





**Fig. 4.** Models generally converge more quickly after intra-domain transfer. Convergence speed is estimated in two ways for each combination of pretraining dataset and task. Firstly, via training loss decrease (left graph), and secondly via time-to-convergence (TTC), which is approximated as the first epoch in which validation performance has not improved by at least 0.5 % for a patience duration of 10 epochs, analogously to common early stopping methods. The figure displays these estimates for fine-tuning on BraTS 2020 (a), ChestX-ray8 (b), ImageNet-1k (c), LiTS 2017 (d) and RadNet-1.28M (e). It also displays these estimates for linear evaluation on BraTS 2020 (f), ChestX-ray8 (g), ImageNet-1k (h), LiTS 2017 (i) and RadNet-1.28M (j). Performance improvement still sometimes occurs after this point, although only with barely visible changes in validation loss. A pattern of inverse correlation between performance and convergence speed estimates is observed, although not across all experiments.

logic: information encoded in the color channels has different meanings between RadNet and ImageNet-pretraining. Whereas models pretrained on ImageNet can reuse any texture-encoding feature and “discard” color-encoding features in a cross-domain transfer scenario, RadNet-pretrained models must learn an entirely new concept in the reverse case. The overall competitive performance of ImageNet, and the asymmetry in generalization gaps which we observe, substantiate our second Hypothesis; ImageNet pretraining delivers a set of features that on average transfer much better than RadNet features, though they do not generally outperform intra-domain transfer.

Another implication of our results is that data domains possess very narrow definitions. Judging by their performance, RadNet’s CT images are part of the same subdomain as the LiTS CT images, but not as the BraTS MRI images. If there exists such a thing as a medical image domain, its subdomains (CTs, MRIs, X-rays, etc.) are distinct entities divided by generalization gaps. As CTs and MRIs rely on different physics and display the same anatomical structure somewhat differently, this distinction is quite intuitive. Consequentially, if the choice is between a capable generalist model and a specialist medical model from a different domain, the generalist model, i.e. an ImageNet-pretrained model, appears to be the better choice.

Curiously, the convergence speed in terms of training loss and validation performance both appear to correlate only partially with the pretraining domain, sometimes offering a similar and sometimes offering a slightly improved convergence speed. Generally, any pretraining appears to increase convergence speed on any task. However, more performant models occasionally require more time to converge or decrease their training loss at a slower or similar rate, compared to less performant ones (for example during fine-tuning on ImageNet-1k). The rate of training loss decrease and convergence speed measured by validation performance also appear to be only loosely correlated to each

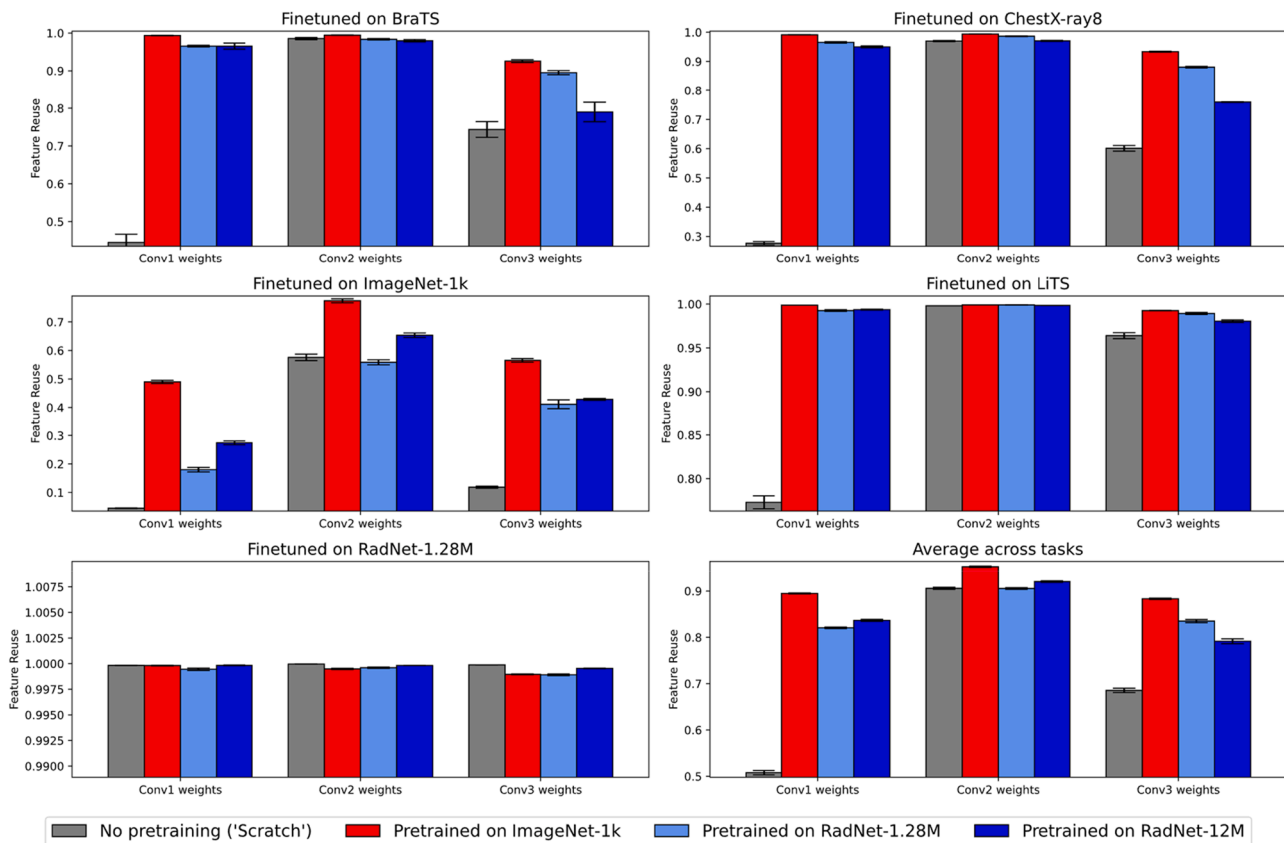
other. This implies that eventual test-time performance should not be estimated by observing either of these quantities.

#### 4.1. Additional feature space analysis

Fig. 5 depicts the degree of feature reuse during transfer learning, approximating the similarity of lowest-level convolution filters before and after fine-tuning using Centered Kernel Alignment [19]. We find that feature reuse is greater in intra-domain than cross-domain transfer scenarios for the ImageNet (significantly) and RadNet (slightly) classification tasks. Like our observations on performance, reduced feature reuse reaffirms that CT and MRI scans constitute separate data domains. Additionally, we observe a higher average feature reuse across tasks for ImageNet pretraining compared to RadNet pretraining further corroborating Hypothesis II. Interestingly, this feature reuse remains high even when the color channels no longer encode color (CTs) or encode different information (MRIs). This implies that for early convolution kernels, ImageNet pretraining encodes shape and texture more strongly than color, which is in agreement with previous observations in literature [20].

#### 4.2. The nature of the intra-domain transfer performance gains

To understand where the performance gains of RadNet pretraining are found in practice, we investigated two possibilities: I) Unique low-level features which models acquire during in-domain pretraining are responsible for the intra-domain advantage. In this case we expect to see many data points which only the model with the intra-domain advantage correctly identifies and few or no data points which only the other models correctly identify. II) Alternatively, the classification process could be so intractable that such features cannot be isolated. In that case,



**Fig. 5.** Intra-domain transfer leads to increased feature reuse. Increased feature reuse is evident for intra-domain transfer scenarios. Feature reuse is measured for the first three convolution layers of the pretrained and fine-tuned models (using RBF-kernel CKA). It is visibly increased in intra-domain transfer scenarios, although feature reuse after RadNet pretraining is only marginally higher than that of ImageNet pretraining when fine-tuning on RadNet-1.28M. Notably, even when channels encode different information, such as when pretraining happens on ImageNet and fine-tuning on BraTS, this feature reuse remains high – higher than for pretraining on CT images.

whether a data point is correctly identified should depend only on a model’s known performance but not the data point itself. Fig. 6 displays histograms of the fraction of data points correctly identified by the combination of models described on the X-axis. For example, the bar labeled “ImageNet” shows the number of data points identified correctly only by the ImageNet-pretrained model, but not by any other models. The corresponding hatched bars show the expected numbers, assuming the second explanation is correct.

We observe that two models with similar performance do not necessarily correctly classify the same data. Similarly, models with superior performance do not correctly classify a strict superset of data compared to models of inferior performance. This observation is in agreement with previous observations by Cohen et al. [8]. While there exists for every model at least a small number of data points which only that model could correctly identify, this number is always highest for the model(s) pretrained on the task domain (rightmost bars in Figs. 6a-d). This favors our first explanation that unique features are responsible for the intra-domain advantage. This effect is again substantially magnified during linear evaluation.

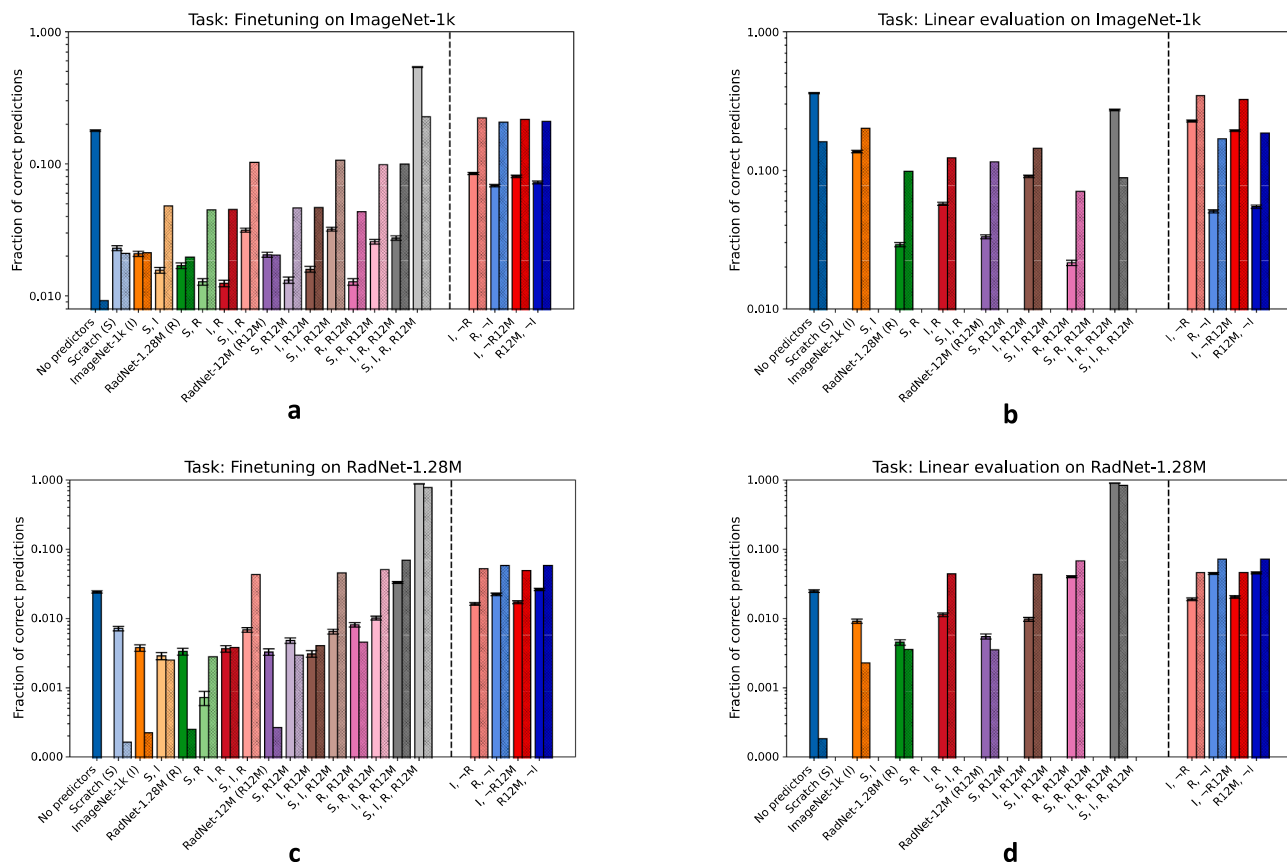
Through investigation of around 60 LiTS images with high prediction discrepancies between RadNet- and ImageNet-pretrained models, we found that an appreciable part of the performance gains for the LiTS liver segmentation corresponded to explainable, macroscopic image/region properties. Improvements from RadNet pretraining are often observed in border regions between classes and for (semi-)insular segmentation targets like seemingly unconnected pieces of liver tissue, whose connection to the liver is vertical. These conditions seem intuitively reasonable: insular areas of organ tissue are likely to appear in the training set during RadNet pretraining and less likely during ImageNet

pretraining. Furthermore, contrast-enhancing features must be learned during RadNet pretraining to identify structures, whereas ImageNet-pretraining can do the same using only texture or color information. This potentially gives RadNet-pretrained models an advantage near region borders on the colorless CT images of LiTS or RadNet, where the unenhanced contrast is low. Fig. 7 shows several representative examples of this advantage.

### 5. Conclusions & outlook

In this paper, we have systematically explored the efficacy of self-supervised pretraining and transfer learning across multiple datasets from the natural and medical image domains. Our experiments demonstrate the existence of a performance-degrading generalization gap that occurs during cross-domain transfer learning. They further highlight that different combinations of pretraining and fine-tuning datasets yield differently sized gaps, depending on task domain, complexity, and whether or not linear evaluation is performed. We confirmed our initial hypotheses that I) intra-domain pretraining is generally preferable, even though II) some datasets such as ImageNet generalize better on average. We have provided evidence that in-domain pretraining yields unique low-level features which are preserved during fine-tuning on data from the same domain, and have discussed the nature of the performance gains they provided. Finally, we have made available our pretrained and fine-tuned models as baselines for future work.

Several limitations apply to this work. Firstly, our conclusions may not generalize to different hyperparameters. Furthermore, models trained using other pretraining algorithms (e.g. generative or



**Fig. 6.** Intra-domain transfer leads to development of unique features and meaningful feature reuse. This figure encodes the fraction of the fine-tuning dataset was correctly predicted by which pretraining sources (i.e. the bar labeled “RadNet, -ImageNet” shows the fraction of the test set identified correctly only by the RadNet-pretrained model, but not (-) by the ImageNet-pretrained model). This is shown for: **a**, Fine-tuning on ImageNet-1k. **b**, Linear evaluation on ImageNet-1k. **c**, Fine-tuning on RadNet-1.28M. **d**, Linear evaluation on RadNet-1.28M. An intra-domain advantage exists, with a significant number of data points only predicted by the model pretrained on data of the same domain. This effect is amplified in the linear evaluation scenario.

generative-adversarial methods), supervision levels, or data domains different from ours may not support our conclusions, although there is no obvious reason why our a priori assumptions would not apply. The same limitation extends to other, yet to be tested model architectures. Modern Transformer [21] or State Space models [22], for example, may exhibit different behaviors. It should also be noted that our results for linear evaluation on segmentation tasks should be viewed with some degree of caution, as the addition of the trainable U-Net decoder still allowed half the model to be trained, potentially offsetting some intra-domain advantage and introducing a source of uncertainty. Finally, all data in the RadNet datasets comes from one hospital. This limitation may have somewhat reduced the generalizability of the dataset, potentially implying a larger intra-domain advantage than we report in this paper.

We summarize a general set of rules for optimal transfer learning performance, based on our experiments:

- Intra-domain transfer learning typically performs similarly or outperforms cross-domain transfer learning by at least a small margin. In short, **“My medical AI should be looking at pictures of tumors instead of birds to learn about tumors.”**
- Intra-domain pretraining is generally preferable in few-shot or linear evaluation settings.
- The definition of a data domain can be very narrow. Our experiments suggest that intra-domain transfer for MR images would require different pretraining data compared to CT images, for example.
- We observed no performance saturation for further upscaling of the pretraining in many scenarios, with RadNet-12M-pretrained models

often outperforming ImageNet-1k-pretrained models. Thus, **a wealth of previously unused, unlabeled medical images can in fact be leveraged to achieve scaling performance gains.**

- We observed that the generalization gap varies very nuancedly; Even the target class or metric, e.g. liver vs. lesion on the LiTS task, can be relevant for the choice of pretraining dataset.
- Using ImageNet as a default pretraining dataset is suitable. ImageNet pretraining gave models remarkable training stability and competitive performance across most tasks in our experiments. Hence, we recommend its use if in-domain pretraining data is not available in sufficient quantities, and always at least testing its performance.
- **There is no “one size fits all” solution.** If the best possible performance on a task is more relevant than a quick solution, there is no substitute for testing multiple pretraining datasets, particularly from the same data domain as the task.

In future work, we plan to extend RadNet-12M with additional data from other modalities and to pretrain additional architectures using RadNet. We believe that research extending our work to other pre-training algorithms and supervised scenarios would constitute a promising research direction, addressing some of the limitations of this work. Creating large-scale models pretrained on a previously untapped wealth of unlabeled clinical images would be highly valuable to the exponentially growing, yet notoriously data-starved [23] field of medical computer vision.






Fig. 7. RadNet's intra-domain advantage relates to explainable, macroscopic image properties.

Four example images from the LiTS 2017 segmentation task are shown, with colored areas highlighting pixels correctly classified only by the RadNet-pretrained model, but not the ImageNet-pretrained model (blue) and the reverse case (yellow). Each image corresponds to one or two of the three typical scenarios in which RadNet pretraining appears to outperform ImageNet pretraining: **a**, Border region between liver and tumor. **b**, Border region between liver and surrounding tissue. **c**, Semi-insular region of the liver. **d**, Insular region of the liver. Note that the images shown are some of the more extreme examples - most of the time, the differences in predictions are much more minute.

#### Data availability

We link the sources of the public datasets we used in our code repository. The RadNet datasets are not made public. All pretrained/fine-tuned models reported are linked on <https://github.com/TIO-IKIM/Transfer-learning-across-domain-boundaries> and may be used for further research. This permission includes models pretrained on the RadNet datasets.

#### Ethics statement

Permission for the retrospective use of the reported imaging data was granted by the IRB of the University Hospital Essen, on December 7th, 2020 (cn. 20-9745-BO). Informed consent was waived, as all data was fully anonymized before usage, and no images were made publicly available. All data was used in accordance with local data protection laws and ethical guidelines.

#### CRediT authorship contribution statement

**Frederic Jonske:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Moon Kim:** Software, Resources, Funding acquisition. **Enrico Nasca:** Software, Resources. **Janis Evers:** Software, Methodology, Data curation. **Johannes Haubold:** Writing – review & editing, Resources. **René Hosch:** Resources, Data curation. **Felix Nensa:** Software, Resources, Funding acquisition. **Michael Kamp:** Writing – review & editing, Validation, Methodology. **Constantin Seibold:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Jan Egger:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization. **Jens Kleesiek:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition, Conceptualization.

#### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Acknowledgements

We acknowledge KITE (Plattform für KI-Translation Essen) from the REACT-EU initiative (EFRE-0801977, <https://kite.ikim.nrw/>). We further acknowledge support by the Open Access Publication Fund of the University of Duisburg-Essen.

#### Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.cmpb.2025.108634](https://doi.org/10.1016/j.cmpb.2025.108634).

#### References

- [1] M. Cherti, J. Jitsev, Effect of pre-training scale on intra- and inter-domain, full and few-shot transfer learning for natural and X-ray chest images, in: 2022 international joint conference on neural networks (IJCNN), Padua, Italy, IEEE, Jul. 2022, pp. 1–9, <https://doi.org/10.1109/IJCNN55064.2022.9892393>.
- [2] J. Deng, W. Dong, R. Socher, L.J. Li, Kai Li, Li Fei-Fei, ImageNet: a large-scale hierarchical image database, in: 2009 IEEE conference on computer vision and pattern recognition, Miami, FL, IEEE, Jun. 2009, pp. 248–255, <https://doi.org/10.1109/CVPR.2009.5206848>.
- [3] F.C. Ghesu, et al., Contrastive self-supervised learning from 100 million medical images with optional supervision, *J. Med. Imag.* 9 (06) (Nov. 2022), <https://doi.org/10.1117/1.JMI.9.6.064503>.
- [4] X. Mei, et al., RadImageNet: an open radiologic deep learning research dataset for effective transfer learning, *Artif. Intellig.* 4 (5) (Sep. 2022), <https://doi.org/10.1148/ryai.210315> e210315.
- [5] S. Azizi, et al., Robust and data-efficient generalization of self-supervised machine learning for diagnostic imaging, *Nat. Biomed. Eng.* 7 (6) (Jun. 2023) 756–779, <https://doi.org/10.1038/s41551-023-01049-7>.

- [6] M. Raghu, C. Zhang, J. Kleinberg, S. Bengio, Transfusion: understanding transfer learning for medical imaging, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Eds.), *Transfusion: understanding transfer learning for medical imaging*, Adv. Neural Inf. Process. Syst. (2019) [Online]. Available, [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/eb1e78328c46506b46a4ac4a1e378b91-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/eb1e78328c46506b46a4ac4a1e378b91-Paper.pdf).
- [7] A. Newell, J. Deng, How useful is self-supervised pretraining for visual tasks?, in: 2020 IEEE/CVF conference on computer vision and pattern recognition (CVPR), Seattle, WA, USA IEEE, Jun. 2020, pp. 7343–7352, <https://doi.org/10.1109/CVPR42600.2020.00737>.
- [8] J.P. Cohen, M. Hashir, R. Brooks, H. Bertrand, On the limits of cross-domain generalization in automated X-ray prediction, in: T. Arbel, I. Ben Ayed, M. de Bruijne, M. Descoteaux, H. Lombaert, C. Pal (Eds.), *Proceedings of the third conference on medical imaging with deep learning 121*, PMLR, Jul. 2020, pp. 136–155. *Proceedings of Machine Learning Research*[Online]. Available, <https://proceedings.mlr.press/v121/cohen20a.html>.
- [9] S. Gururangan, et al., Don't stop pretraining: adapt language models to domains and tasks, in: *Proceedings of the 58th annual meeting of the association for computational linguistics*, Online: Association for Computational Linguistics, 2020, pp. 8342–8360, <https://doi.org/10.18653/v1/2020.acl-main.740>.
- [10] H. Petzka, M. Kamp, L. Adilova, C. Sminchisescu, M. Boley, *Relative flatness and generalization*, Adv. Neural Inf. Process. Syst. (2021) 18420–18432.
- [11] S. Ben-David, J. Blitzer, K. Crammer, F. Pereira, Analysis of representations for domain adaptation, in: B. Schölkopf, J. Platt, T. Hoffman (Eds.), *Analysis of representations for domain adaptation*, Adv. Neural Inf. Process. Syst. (2006) [Online]. Available, [https://proceedings.neurips.cc/paper\\_files/paper/2006/file/b1b0432ceafb0ce714426e9114852ac7-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2006/file/b1b0432ceafb0ce714426e9114852ac7-Paper.pdf).
- [12] Y. Mansour, M. Mohri, A. Rostamizadeh, Domain adaptation: learning bounds and algorithms, in: *Proceedings of The 22nd annual conference on learning theory (COLT2009)*, Canada, Montréal, 2009 [Online]. Available, <http://www.cs.nyu.edu/~mohri/postscript/nadap.pdf>.
- [13] S. Koitka, L. Kroll, E. Malamutmann, A. Oezcelik, F. Nensa, Fully automated body composition analysis in routine CT imaging using 3D semantic segmentation convolutional neural networks, *Eur. Radiol.* 31 (4) (Apr. 2021) 1795–1804, <https://doi.org/10.1007/s00330-020-07147-3>.
- [14] F. Jonske, et al., Deep Learning-driven classification of external DICOM studies for PACS archiving, *Eur. Radiol.* 32 (12) (Dec. 2022) 8769–8776, <https://doi.org/10.1007/s00330-022-08926-w>.
- [15] P. Bilic, et al., The liver tumor segmentation benchmark (LiTS), *Med. Image Anal.* 84 (2023) 102680, <https://doi.org/10.1016/j.media.2022.102680>.
- [16] R. Mehta, et al., QU-BraTS: MICCAI BraTS 2020 challenge on quantifying uncertainty in brain tumor segmentation – Analysis of ranking scores and benchmarking results, *Melba 1* (August 2022) 1–54, <https://doi.org/10.59275/j.melba.2022-354b>. Aug. 2022.
- [17] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, R.M. Summers, ChestX-Ray8: hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases, in: 2017 IEEE conference on computer vision and pattern recognition (CVPR), Honolulu, HI, IEEE, Jul. 2017, pp. 3462–3471, <https://doi.org/10.1109/CVPR.2017.369>.
- [18] B.L. Welch, The generalization of 'student's' Problem when several different population variances are involved, *Biometrika* 34 (1–2) (1947) 28–35, <https://doi.org/10.1093/biomet/34.1-2.28>.
- [19] S. Kornblith, M. Norouzi, H. Lee, G. Hinton, Similarity of neural network representations revisited, in: K. Chaudhuri, R. Salakhutdinov (Eds.), *Proceedings of the 36th International conference on machine learning 97*, PMLR, Jun. 2019, pp. 3519–3529. *Proceedings of Machine Learning Research*[Online]. Available, <https://proceedings.mlr.press/v97/kornblith19a.html>.
- [20] R. Geirhos, P. Rubisch, C. Michaelis, M. Bethge, F.A. Wichmann, W. Brendel, ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness, in: *International conference on learning representations, 2019* [Online]. Available, <https://openreview.net/forum?id=Bygh9j09KX>.
- [21] A. Dosovitskiy et al., "An image is worth 16x16 words: transformers for image recognition at scale," Jun. 03, 2021, *arXiv:2010.11929*. doi: 10.48550/arXiv.2010.11929.
- [22] A. Gu, K. Goel, and C. Ré, "Efficiently modeling long sequences with structured State spaces," 2021, *arXiv:2010.11929*. doi: 10.48550/ARXIV.2111.00396.
- [23] J. Egger, et al., Medical deep learning—A systematic meta-review, *Comput. Methods Programs Biomed.* 221 (Jun. 2022) 106874, <https://doi.org/10.1016/j.cmpb.2022.106874>.