

The Role of Pre-training Data in Transfer Learning

Rahim Entezari^{*1}, Mitchell Wortsman², Olga Saukh¹,
M.Moein Shariatnia³, Hanie Sedghi⁴, and Ludwig Schmidt²

¹TU Graz / CSH Vienna

²University of Washington

³Tehran University of Medical Sciences

⁴Google Research, Brain Team

Abstract

The transfer learning paradigm of model pre-training and subsequent fine-tuning produces high-accuracy models. While most studies recommend scaling the pre-training size to benefit most from transfer learning, a question remains: what data and method should be used for pre-training? We investigate the impact of pre-training data distribution on the few-shot and full fine-tuning performance using 3 pre-training methods (supervised, contrastive language-image and image-image), 7 pre-training datasets, and 9 downstream datasets. Through extensive controlled experiments, we find that the choice of the pre-training data source is essential for the few-shot transfer, but its role decreases as more data is made available for fine-tuning. Additionally, we explore the role of data curation and examine the trade-offs between label noise and the size of the pre-training dataset. We find that using 2000× more pre-training data from LAION can match the performance of supervised ImageNet pre-training. Furthermore, we investigate the effect of pre-training methods, comparing language-image contrastive vs. image-image contrastive, and find that the latter leads to better downstream accuracy¹².

1 Introduction

The best-performing computer vision models are produced by the transfer learning paradigm. While transfer learning is not new, the substantial improvement in the quality of the pre-trained models in recent years has brought transfer learning to the spotlight (e.g., CLIP (Radford et al., 2021), BASIC (Pham et al., 2021), and Flamingo (Alayrac et al., 2022)). These improvements are driven by new datasets for pre-training as well as better pre-training algorithms. This naturally leads to a question:

How do the dataset and the algorithm used for pre-training affect downstream performance?

In contrast to prior works (Kornblith et al., 2019; Abnar et al., 2021; You et al., 2021; Nguyen et al., 2020; Deshpande et al., 2021; Bolya et al., 2021), our main focus is on the role of the pre-training data distribution in downstream performance. We set up systematic experiments to explore our research questions and contributions as follow:

^{*}Work done while interning at the University of Washington.

¹Our comparisons include a smaller scale pre-training compared to SOTA ImageNet results using CLIP pre-training. SOTA CLIP pre-training lacks a comprehensive comparison to image-only pre-training on few-shot and full-finetune.

²Code is available at <https://github.com/rahimentezari/DataDistributionTransferLearning>




Figure 1: **Differences between various pre-training sources diminish as more data is available for the downstream tasks.** In the few-shot setting, different pre-training datasets lead to noticeable differences in downstream performance. However, if many samples are available for fine-tuning, the difference in absolute accuracy between models pre-trained on different sources largely evaporates. (see Figure 14 for a detailed view).

Do different pre-training distributions lead to different transfer learning performances? In practice, one has many options to download pre-training checkpoints and fine-tune the model on the target dataset. Should we expect different pre-training datasets to perform differently in the transfer setting? When controlling for the size of the pre-train model and the downstream dataset, but changing the pre-train dataset, we observe noticeable differences in downstream accuracy in the few-shot setting (only a few examples per class are available for fine-tuning). However as more samples are available for fine-tuning, the difference in absolute accuracy when varying the pre-training dataset largely evaporates. In the few-shot regime, we observe that certain pre-training datasets (*e.g.*, Shutterstock) consistently lead to a better transfer accuracy than the other (*e.g.*, WiT) across many downstream tasks. However, the ranking of the other pre-training datasets in our selection appears mixed. Moreover, even the pre-training dataset which leads to the worst transfer accuracy (WiT) still outperforms training from scratch (see Section 4.1, Figure 1 and Figure 14).

How much is expensive labeling worth compared to noisier but larger pre-training data? We compare different pre-training strategies: supervised pre-training on the carefully labeled ImageNet dataset and semi-supervised pre-training on language-image pairs from larger but noisier datasets. We find that pre-training on a well-curated dataset leads to better transfer accuracy than pre-training on a noisy dataset of a similar size. Our investigations also show that pre-training on a 15x-2000x larger but noisier dataset (LAION) can close the gap for some downstream tasks (see section 4.4, section 4.5, Figure 2 and Figure 3).

How does increasing the pre-training dataset size impact the transfer learning performance? Can we expect that more pre-training data implies a better performance, or can the pre-training effectiveness saturate at some point? Our findings indicate that, while larger pre-training datasets tend to improve transfer performance, such improvement depends on the dataset and the number of fine-tuning samples. We observe that some downstream tasks show little improvement even with large datasets, while others benefit

significantly. (see section 4.6 and Figure 4).

What is the impact of the pre-training method on downstream performance? We examine the difference between supervised pre-training with the popular CLIP and SimCLR semi-supervised algorithms. Overall the diminishing effect for pre-training distributions shown in Figure 1 is also observed when the pre-training method is switched from CLIP to SimCLR in Figure 5, meaning differences in downstream accuracy is largely observed in the few-shot setting. In addition, we find that the SimCLR pre-training leads to better transfer than CLIP pre-training in the few-shot regime. However, the difference is small if many images are used for fine-tuning (see section 4.7 and Figure 6).

We conduct an extensive empirical investigation in the context of transfer learning for computer vision tasks (See section B for details on 4000 experiments). Our study covers seven pre-training datasets including YFCC, LAION, Redcaps, Conceptual Captions-3m, Conceptual Captions-12m, WiT, Shutterstock, and ImageNet (Thomee et al., 2016; Schuhmann et al., 2021; Desai et al., 2021; Sharma et al., 2018; Changpinyo et al., 2021; Srinivasan et al., 2021; Deng et al., 2009), nine fine-tuning datasets including CIFAR100, DTD, Caltech-101, Oxford-PETS, REAL and CLIPART from DomainNet, EuroSAT, Cassava Leaf Disease, and Caltech Camera Traps (Krizhevsky et al., 2009; Cimpoi et al., 2014; Fei-Fei et al., 2004; Parkhi et al., 2012; Peng et al., 2019; Helber et al., 2019; Cas; Beery et al., 2018), and three pre-training methods: supervised, CLIP (Radford et al., 2021) and SimCLR (Chen et al., 2020). To evaluate downstream performance, we examine both few-shot fine-tuning and full fine-tuning.

The paper is structured as follows: we review closely related works in Section 2, followed by our experimental setup presented in Section 3. Section 4 details our observations relating to our research questions by measuring the downstream transfer accuracy of models pre-trained on various data sources, dataset sizes, and with different pre-training losses. We discuss our findings and conclude with future research directions in Section 5.

2 Related Work

Transfer learning is widely used in deep learning research and practice and has become a cornerstone in both computer vision and natural language processing. Through the years, there have been many questions on why transfer helps and how to choose a good pre-trained model to transfer from. Neyshabur et al. (2020) separated the effect of feature reuse from that of learning low-level pre-training data statistics. Raghu et al. (2019) investigate the similarity of the pre-training and downstream datasets by looking into medical datasets and found that transfer learning from ImageNet pre-trained models shows little benefit in performance. Ericsson et al. (2021) studied the downstream performance of self-supervised models and found that the best self-supervised models of that time could outperform supervised pre-training as an upstream source of knowledge transfer and that the performance of self-supervised models on ImageNet is indicative of downstream performance on natural image classification tasks. Similarly, Islam et al. (2021) found that contrastively trained models consistently outperform standard cross-entropy models in transfer learning. Goyal et al. (2021) showed that self-supervised models outperform supervised models on ImageNet, even when trained on random and uncurated images from the web. Moreover, they showed that these models are also good at few shot learning by achieving 77.9% top-1 accuracy using only 10% on ImageNet.

Building on contrastive techniques, Radford et al. (2021) introduced CLIP which learns a joint embedding space for both images and their descriptive captions, making it possible to effectively leverage a large-scale dataset from the Internet. Flamingo (Alayrac et al., 2022), a visual language model, is another successful example in the line of multimodal models and enables visual question answering and image captioning. CLIP and similar models like ALIGN (Jia et al., 2021), BASIC (Pham et al., 2021), and LiT (Zhai et al., 2022) demonstrated unprecedented robustness to challenging data distribution shifts. This accomplishment raised questions on the probable sources of such robustness—whether this robustness is caused by language supervision, the pre-training data distribution, size, or contrastive loss functions.

Fang et al. (2022) investigated this question and found that the diverse training distribution is the main cause of the robustness properties of CLIP. Nguyen et al. (2022) explored the role of the pre-training dataset

for CLIP with a testbed of six pre-training sources, finding that no single pre-training dataset consistently performs best. In recent work, [Santurkar et al. \(2022\)](#) carefully investigated the effect of language supervision in CLIP-like models, finding it an important factor if the pre-training dataset is large and the captions are descriptive enough. Unlike their work, we consider end-to-end fine-tuning which result in higher accuracy. [Cherti & Jitsev \(2021\)](#) study the effect of scaling the pre-training model and data for both in-domain and out-of-domain transfer. They conduct supervised pre-training while varying pre-training model size and data source (ImageNet-1k/21k or large medical chest XRay datasets), and transfer pre-trained models to different natural or medical targets. They find that, when performing transfer to large X-Ray targets, pre-training on natural ImageNet-21k is as good or better than pre-trained medical X-Ray data. [Djologna et al. \(2021\)](#) also investigate the impact of the pre-training data size and model scale, finding that increasing both the training set and model sizes significantly improve the distributional shift robustness.

This work is inspired by and closely related to [Kim et al. \(2022\)](#) and [Abnar et al. \(2021\)](#). [Kim et al. \(2022\)](#) conducted an in-depth study of the effect of the network architecture, pre-training dataset, supervised vs self-supervised learning objectives, and different domain transfer methods on the transferability of representations to new domains. They found that the transferability of the pre-trained representations depends on factors such as the target benchmark, adaptation method, and network depth. However, they do not study few-shot transfer (where we see the most impact of the pre-training distribution). They also did not provide a set of controlled experiments for some of the studied impacting factors because they are limited to available pre-trained models. For example, when comparing the role of data distribution (their Figure 2, ImageNet-22K vs. JFT-300m), they change the dataset size and also architecture, and the reader is left wondering if JFT has a better distribution for transfer or if the observed effects come from more data or a better architecture?

[Abnar et al. \(2021\)](#) also explored how different upstream training settings affect transfer accuracy for two upstream datasets and more than 20 downstream tasks. They showed that as the upstream accuracy increases, the transfer learning performance on downstream datasets saturates. However, the authors study only upstream models that are pre-trained with a supervised loss function on ImageNet-21K ([Deng et al., 2009](#)) or JFT-300M ([Sun et al., 2017](#)) (different size and distributions). In this work, we extend these results to more pre-training datasets and methods, with a special focus on data distribution and curation. [Abnar et al. \(2021\)](#) also lacks controlled comparison between different distributions in the pre-training datasets *e.g.*, they compare JFT and ImageNet with very different sample sizes. We consider full fine-tuning in addition to few-shot transfer. Moreover, [You et al. \(2021\)](#); [Nguyen et al. \(2020\)](#); [Deshpande et al. \(2021\)](#) develop metrics for predicting the transferability of a model. Their main focus is to develop a measure to predict the full fine-tune accuracy without actually fine-tuning on the downstream task. While we also cover full fine-tune accuracy, our main research question lies in studying the extent to which pre-training data affect transfer accuracy. Looking at few-shot and full-shot also gives us the ability to study the effect of transfer learning as more target data become available. Moreover, predictability of the transfer performance is mostly limited to supervised ImageNet-1K pretraining, while we scale both pre-training distributions, size, and pre-training loss functions. The transferability line of research also mainly focuses on Internet-crawled datasets, while we extended our results to domain-specific datasets (Camera Traps, Cassava Leaf Diseases, and EuroSAT).

3 Experimental setup

Model The main focus of this study is the CLIP model ([Radford et al., 2021](#)). This model has demonstrated unprecedented robustness to natural distribution shifts ([Taori et al., 2020](#); [Miller et al., 2021](#)), and transfers well to many downstream tasks ([Radford et al., 2021](#); [Wortsman et al., 2021](#)). Given an image-text pair, CLIP learns a joint embedding space for both images and their captions and tries to maximize the cosine similarity between the text and image embedding for an image relative to the cosine similarity of unaligned pairs. We use the CLIP implementation from the OpenCLIP GitHub repository ([Ilharco et al., 2021](#)).

Pre-training We mainly use ResNet-50 ([He et al., 2016](#)) as the image encoder unless stated otherwise. We vary the pre-training data distribution in section 4.1, curation method in section 4.4, and pre-training dataset size in section 4.6 to obtain different pre-trained models. We also change the contrastive loss function

	CIFAR100	DTD	CALTECH	PETS	DomainNet REAL	DomainNet CLIPART	EuroSAT	Cassava Leaf Disease	Camera Traps	average
Train from scratch (no pre-training)	72.82	44.62	55.32	67.96	77.76	55.20	58.50	70.30	89.10	65.73
Worst performing pre-training	81.62	61.06	82.38	83.15	81.87	68.98	70.72	87.05	98.85	79.52
Best performing pre-training	83.71	68.56	86.88	87.84	82.92	72.38	72.79	87.57	98.91	82.39

Table 1: Full fine-tune accuracy from six different pre-training datasets across all downstream datasets. Transfer learning from even the worst pre-training dataset outperforms training from scratch largely.

to SimCLR in section 4.7 to test the effect of the pre-training method on downstream transfer accuracy. Our CLIP models are trained from scratch on each of the pre-training datasets unless otherwise mentioned and follow the training code from the OpenCLIP GitHub repository (Ilharco et al., 2021). CLIP models are trained using AdamW optimizer (Loshchilov & Hutter, 2017) with default PyTorch parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, batch size 1024, and weight decay of 0.1. For learning rate, we start with a learning rate of 10^{-3} and apply a cosine-annealing learning rate schedule (Loshchilov & Hutter, 2016) with 5,000 steps warm-up. We use the same data augmentations as in (Radford et al., 2021).

Our SimCLR implementation closely follows the training code from the SLIP (Mu et al., 2021). SimCLR models are also trained for 16 epochs from scratch using AdamW optimizer (Loshchilov & Hutter, 2017) with $\beta_1 = 0.9$, $\beta_2 = 0.98$, $\epsilon = 10^{-8}$, batch size 1024, and weight decay of 0.1. we start with a learning rate of 10^{-3} and apply a cosine-annealing learning rate schedule (Loshchilov & Hutter, 2016) with 2 epochs of warm-up. The hidden dimension of SimCLR MLP projection head is set to 4,094 and the output embedding dimension of MLP projection head is set to 256.

Fine-tuning For most of the experiments we fine-tune the pre-trained model end-to-end on the target transfer dataset unless stated otherwise. For each pre-trained model and downstream transfer dataset, we used a large grid search over various fine-tuning hyperparameters including learning rate, batch size, and the number of epochs. We report the best-performing accuracy in the plots.

Each pretrained model is finetuned on the specific downstream task for 128 epochs while the learning rate is mostly from 0.0001, 0.0003, 0.001, 0.003 as starting and applying a cosine-annealing learning rate schedule (Loshchilov & Hutter, 2016) with 500 steps warm-up and batch size of 128. For each fine-tuning, we choose the best-performing result on the test set among the performed grid search. We use the implementation from the WiSE-FiT GitHub repository for fine-tuning, where we have only one model and $\alpha = 1$ (Wortsman et al., 2021).

Datasets Our large-scale experiments yield more than 4000 trained networks. Our pre-training datasets consist of million-size image and language pairs from multiple recent multi-modal datasets including YFCC, LAION, RedCaps, Shutterstock, Conceptual Captions, WiT (Thomee et al., 2016; Schuhmann et al., 2021; Desai et al., 2021; Sharma et al., 2018; Changpinyo et al., 2021; Srinivasan et al., 2021). Our pre-training datasets are crawled from different sources covering different data distributions. For downstream tasks, we use nine different datasets CIFAR100, DTD, Caltech-101, Oxford-PETS, REAL, and CLIPART from DomainNet, EuroSAT, Cassava Leaf Disease, and Caltech Camera Traps (Krizhevsky et al., 2009; Cimpoi et al., 2014; Fei-Fei et al., 2004; Parkhi et al., 2012; Peng et al., 2019; Helber et al., 2019; Cas; Beery et al., 2018). While the first six datasets are internet-crawled datasets (similar to pre-training datasets) and are more common in transfer learning in computer vision benchmarks, we include three new downstream datasets that are

domain-specific, i.e. the dataset is created after a specific challenge is defined in a specific domain.

3.1 Downstream tasks

We have used 9 different downstream datasets. Table 5 describes the first six datasets in Figure 14. While these six datasets are internet-crawled datasets and are more common in transfer learning in computer vision benchmarks, we include three new downstream datasets that are domain-specific, *i.e.*, the dataset is created after a specific challenge is defined in a specific domain.

- EuroSAT (Helber et al., 2019): The task is to classify land use and land cover based on Sentinel-2 satellite images. The dataset covers 13 spectral bands and consists of 10 classes within a total of 27,019 labeled and geo-referenced images. we create an 80%-20% random class-balanced split with the provided dataset.
- Cassava Leaf Disease Classification (Cas): The dataset contains 21,397 images from the Kaggle competition, to give farmers access to methods for diagnosing plant diseases. The images are labeled as healthy or as one of four different diseases. we split the dataset with 80%-20% random class-balanced ratio for train and test, respectively.
- Caltech Camera Traps-20 (Beery et al., 2018): CCT-20 contains 57,864 images in 15 classes, taken from camera traps deployed to monitor animal populations. Classes are either single species *e.g.*, "Coyote") or groups of species, *e.g.*, "Bird"). CCT-20 is a subset of the iWildCam Challenge 2018, whose yearly editions have been hosted on Kaggle. Here we study the subset of CCT-20 that come from the same locations ³, including 14,071 and 16,395 images for train and test respectively.

3.2 Pre-training datasets

Our study covers 7 pre-training datasets as follow:

- YFCC: Our experiments mostly include YFCC-2.7M, a random subset of YFCC-15M. The 15M subset of the YFCC-100M dataset (Thomee et al., 2016) was filtered to only include images with English titles or descriptions. The dataset contains 14,829,396 images with natural language captions associated with each image. The images and captions are collected from Flickr.
- LAION (Schuhmann et al., 2021): The images and corresponding alt-texts come from web pages collected by Common Crawl (Com) between 2014 and 2021. We randomly select a subset of 2.7M and 15M samples for our experiments.
- Redcaps (Desai et al., 2021): Redcaps contains 11,882,403 examples from 350 manually curated subreddit collected between 2008 and 2020. The subreddits are selected to contain a large number of image posts that are mostly photographs and not images of people.
- Shutterstock: 11,800,000 images and captions from the Shutterstock website.
- Conceptual Captions-3m (Sharma et al., 2018): The raw descriptions in Conceptual Captions are harvested from the alt-text HTML attribute associated with web images. This dataset contains 2,799,553 samples, denoted as CC_2.7m in the plots.
- Conceptual Captions-12m (Changpinyo et al., 2021): A dataset with 12 million image-text pairs. It is larger than CC_2.7m and covers a much more diverse set of visual concepts. We randomly select 2.7M samples from this dataset, denoted as CC_12_2.7m.

³"Cis" in the main dataset refers to images from locations seen during training, and "trans" refers to new locations not seen during training

Table 2: Details on the downstream datasets used in the experiments.

Downstream Task	Description
CIFAR100	The task consists in classifying natural images (100 classes, with 500 training images each). Some examples include apples, bottles, dinosaurs, and bicycles. The image size is 32x32.
DTD	The task consists in classifying images of textural patterns (47 classes, with 120 training images each). Some of the textures are banded, bubbly, meshed, lined, or porous. The image size ranges between 300x300 and 640x640 pixels.
CALTECH-101	The task consists in classifying images of objects (9144 images in 101 classes plus a background clutter class), including animals, airplanes, chairs, or scissors. The image size varies, but it typically ranges from 200-300 pixels per edge.
PETS	The task consists in classifying images of cat and dog breeds (7000 images in 37 classes). Images dimensions are typically 200 pixels or larger
REAL	The task is a subset of larger DomainNet from six distinct domains, including photos (real), painting, clipart, quickdraw, infograph, and sketch. Total size of 172,000
CLIPART	The task is a subset of larger DomainNet from six distinct domains, including photos (real), painting, clipart, quickdraw, infograph, and sketch. Total size of 172,000

- WIT (Srinivasan et al., 2021): Image-text pairs come from Wikipedia pages. We use reference description as the source of text data and obtain 5,038,295 examples in total after filtering to include only the English language.

Table 6 shows their main source and total size. We also show some examples of image-caption pairs randomly selected from Shutterstock in Figure 16, Redcaps in Figure 17, YFCC-15m in Figure 18, LAION-15m in Figure 19, Conceptual Captions in Figure 20, and WIT in Figure 21. Table 7 also shows the most common words in captions of these pre-training datasets.

Looking at Redcaps samples in Figure 17 and also the top 20 captions shows many samples of animals. This is showing why Redcaps perform better on PETS. Samples from WIT in Figure 21 and also its top 20 words mostly featuring geographical locations, which is rare in our downstream task, hence performing worst compared to other pre-training distributions. Shutterstock top 20 words also include words like "pattern", "texture", "and design" which are close to DTD classes, hence showing superior performance in this downstream task.

4 Experiments and Results

In this section, we ask a set of research questions to investigate the role of the pre-training data, the choice of the pre-training method (supervised vs. CLIP vs. SimCLR), and the impact of the ImageNet distribution on the downstream performance. To answer these questions, we carefully design experiments for each section while ablating other impacting factors.

Table 3: Details on pre-training datasets

Dataset	Source	Total size
YFCC	Flickr	14,826,000
LAION	Common Crawl	15,504,742
CC-12M	Unspecified web pages	9,594,338
RedCaps	Reddit	11,882,403
WIT	Wikipedia	5,038,295
Shutterstock	Shutterstock	11,800,000
IN1K-Captions	ImageNet	463,622

Table 4: Most common words in captions of pre-training distributions

Pre-training dataset	Top 20 words in 1M sample of captions
Shutterstock	background , vector, illustration, design , icon, pattern , texture , style, woman, concept, hand, color, flower, view, template, line, business, logo, card, symbol
Redcaps	day, today, year, time, cat, plant, friend, anyone, picture, baby, guy, week, dog, home, morning, night, month, way, boy, work
YFCC-15m	photo, day, park, street, city, picture, view, time, world, year, house, state, center, part, garden, shot, image, building, road, museum
LAION-15m	photo, stock, image, black, woman, design, set, vector, white, print, home, men, blue, dress, art, card, sale, gold, bag, cover
CC-12m	illustration, stock, art, design, photo, image, background, room, vector, house, home, woman, wedding, style, photography, royalty, car, fashion, girl, world
CC-3m	background, actor, artist, player, illustration, view, woman, man, football, team, tree, premiere, city, vector, day, girl, beach, game, hand, people
WIT	view, church, station, map, house, building, hall, museum, city, location, street, park, river, state, john, county, town, center, bridge, world

4.1 What is the impact of different pre-training data sources on transfer learning?

Do we expect different distributions to perform differently in the transfer setting? Figure 1 aggregates transfer performance from different pre-training datasets across all downstream datasets. To get each point, we (1) pre-train CLIP models using a set of seven large sources, (2) fine-tune each pre-trained model on all downstream datasets across different shots, and (3) for each downstream dataset, calculate the difference between the best and worst fine-tune performance among used pre-training sources, normalized by the maximum fine-tune performance. Figure 1 aggregates over all downstream datasets for each number of shots, highlighting as an example different pre-training models fine-tuned using 20 samples/class on all downstream datasets. We observe that changing the pre-training dataset leads to noticeable differences in the downstream performance in a few-shot setting. However, as more images are available for fine-tuning, the difference in absolute accuracy between different pre-training models is largely diminished. Figure 14 shows this diminishing effect in detail for different downstream datasets. The full fine-tuned models have very similar downstream performances despite different pre-training datasets (see the top-right point of CIFAR100 and REAL in Figure 14, and also the top-right point for CameraTraps, Cassava Leaf, and EuroSAT in Figure 14). However, this is not true for DTD, CALTECH101, PETS, and CLIPART, where they have far fewer images




Figure 2: **Effect of data curation and labeling.** We compare supervised pre-training on ImageNet-1K to (1) contrastive pre-training on original captions from Flickr with 0.5m samples, and (2) contrastive pre-training on Templated (clean) captions using ImageNet labels with 1.2m samples. Supervised pre-training on ImageNet leads to better transfer accuracy than contrastive pre-training. Improving captions quality from Flickr to Template leads to huge improvements in downstream transfer accuracy, highlighting the importance of captions quality. On a different comparison between ImageNet and LAION distributions, pre-training CLIP on Flickr captions performs better than pre-training on LAION distribution with the same size (0.5m).

per class for fine-tuning on the full dataset. Table 1 compares fine-tune accuracy for different pre-training choices along all downstream datasets. Transfer learning from even the worst pre-training dataset outperforms training from scratch largely. The gap between best and worst-performing pre-training datasets is small. Section 4.9 extends our results to Vision Transformers (Dosovitskiy et al., 2021) instead of ResNet-50.

4.2 Which data distribution is better for transfer learning?

The results presented in Figure 14 demonstrate that pre-training on the Shutterstock and LAION datasets results in superior transfer performance across a range of downstream tasks. A closer look shows the superior performance of Redcaps for PETS. We investigate this further and inspect many pets by looking at random samples from Redcaps at Figure 17. We also look into the most common words in the captions of these pre-training datasets, summarized in Table 7. We observe that "cats" and "dogs" are among the most common words in the Redcaps dataset. Table 7 also shows that "background", "design", "pattern", and "texture" are among the most common words in the captions of Shutterstock, supporting a high correlation to DTD (Describable Textures Dataset). WiT yields the worst performance in most cases because both

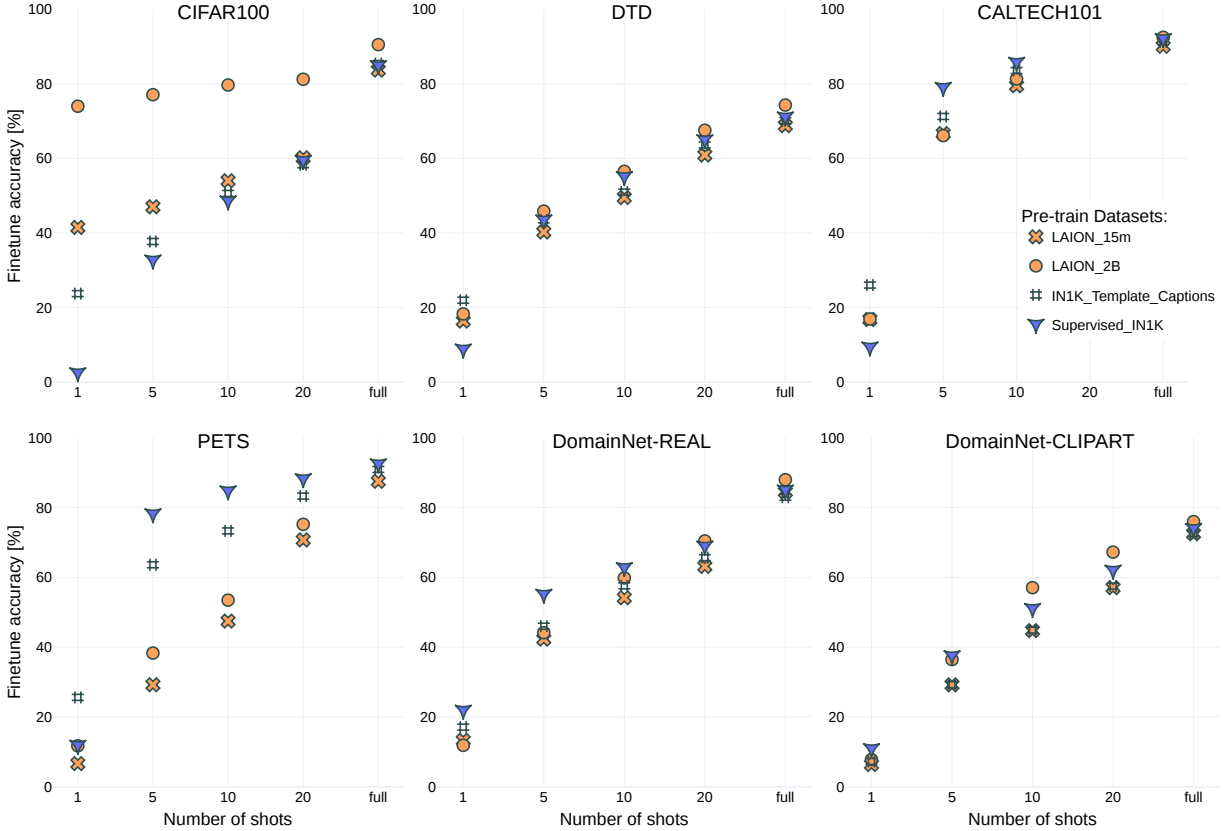


Figure 3: **How much LAION data is worth of ImageNet pre-training?** While Figure 2 shows the superiority of supervised pre-training over contrastive pre-training with the same size, here we increase the size of contrastive pre-training size to see if contrastive pre-training could perform better than supervised pre-training at scale. Including 15x more data from LAION outperform supervised ImageNet pre-training (and template captions) on CIFAR100. However, DTD, REAL, and CLIPART need 2000x more data from LAION to match or outperform ImageNet pre-training. Even including 2000x more data did not help CALTECH101 and PETS, where supervised ImageNet pre-training is still the best choice.

captions and images (Figure 21) are related to topics about people and geography that are far from the studied downstream tasks.

4.3 How much pre-training contributes to downstream performance as opposed to training from scratch?

While transfer learning from a large pre-training dataset outperforms training from scratch for all downstream tasks, the magnitude of the improvement varies for different datasets in Figure 14. We observe a large improvement for PETS, CALTECH-101, and CLIPART. PETS for example has a small number of samples per class for training (30), which makes it hard to train from scratch. It is also scraped from the web (Google search Fei-Fei et al. (2004)), similar to our web-scraped pre-training sources. We also hypothesize that a pre-training shows the best improvement when increases both diversity (how hard pre-train data is to fit) and affinity (how pre-training shifts the decision boundary of the scratch model) (Gontijo-Lopes et al., 2020), meaning it should be semantically close to the classes of target task while enriching the distribution over the samples.




Figure 4: **Effect of the pre-training dataset size.** Increasing the size of the dataset used for pre-training results in better transfer accuracy on downstream tasks. However, the absolute accuracy difference is smaller in the high-shot regime, even when pre-training consists of $100\times$ more data. The benefit of pre-training on LAION-2B is different on target tasks. While there is major gap between LAION-2B and LAION-15m for CIFAR100, the performance gain from scaling up the pre-training dataset on CALTECH101 gets saturated.

4.4 Do well-curated pre-training datasets lead to better transfer?

There has been a significant effort to create computer vision datasets with high-quality labels. On the other hand, many recent datasets are large but noisy. In this section, we are going to investigate: *How much is laborious ImageNet labeling worth?*

To answer this question, we first start by pre-training ResNet-50 on Large Scale Visual Recognition Challenge (ILSVRC) 2012 (Russakovsky et al., 2015), known as ImageNet-1K, using supervised cross-entropy loss and fine-tune on our downstream datasets in Figure 2. To investigate the role of supervision, we then discard ImageNet labels and use CLIP to pre-train on ImageNet. Because the ImageNet dataset has no captions, we include original Flickr captions, which reduces the size of the image and captions to 0.5M samples (Section 4.8 describes the required steps to create ImageNet-Flickr). Figure 2 shows that supervised pre-training on ImageNet outperforms CLIP pre-training on ImageNet with Flickr captions by a large margin in all downstream tasks.

However, such a gap could be attributed to two differences between mentioned pre-trainings: (1) supervised vs. contrastive image-language loss, and (2) the size of training samples for supervised-ImageNet (1.2m) is

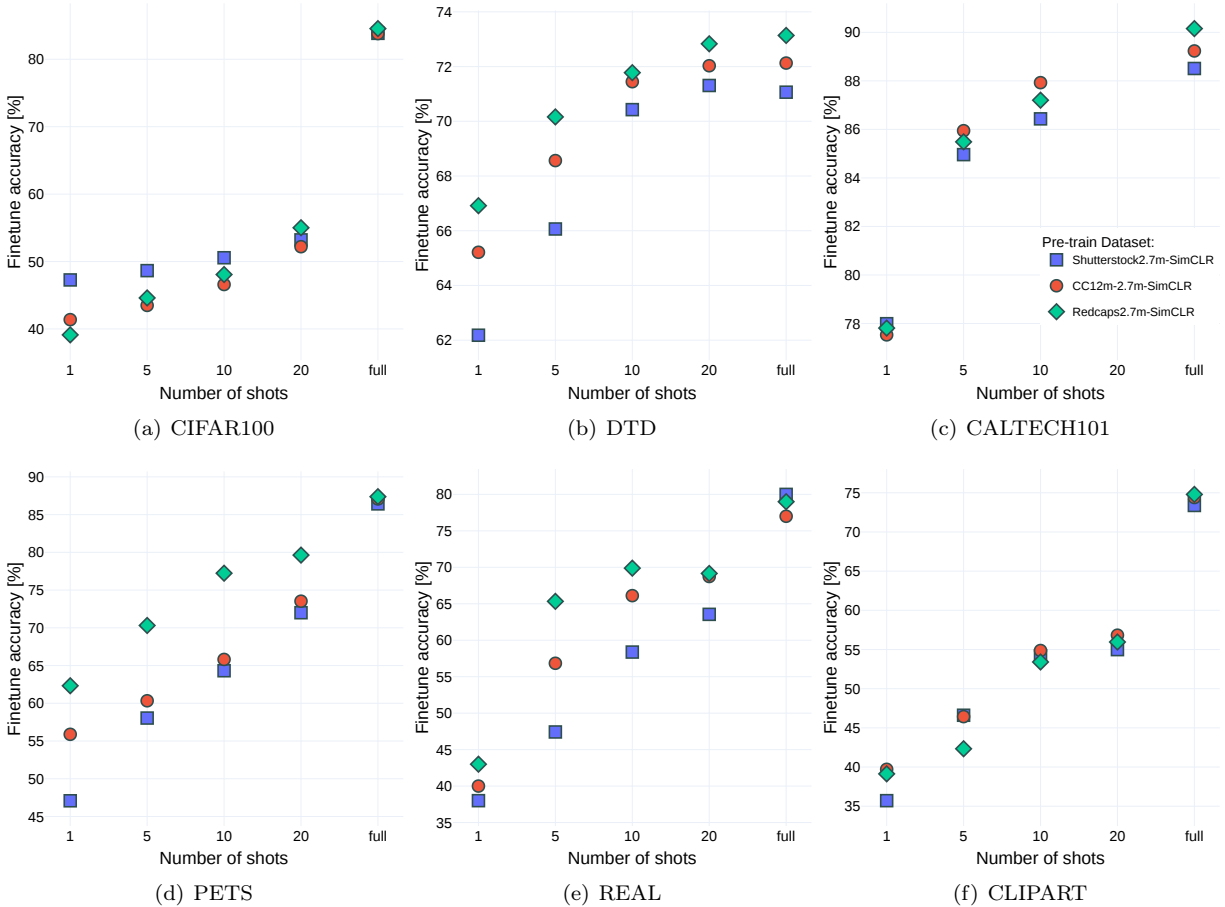


Figure 5: **Effect of the pre-training data distribution when using SimCLR as the pre-training method.** Using different datasets for pre-training leads to a noticeable difference in downstream transfer accuracy. Similarly to the previous results for CLIP pre-training, the absolute difference in downstream transfer accuracy between different pre-training datasets is smaller when many images are available for fine-tuning.

two times larger than CLIP with ImageNet-Flickr captions (0.5m). To remove the second effect we then use all the images from ImageNet, paired with templated clean captions, e.g., “a photo of a *class name*”. This allows us to have a fair comparison between supervised and CLIP pre-training on ImageNet, given the same size. Figure 2 shows that pre-training with clean captions improves the performance of CLIP pre-training by a large margin and outperforms supervised pre-training on CIFAR100. However, supervised pre-training on ImageNet still performs best for the rest of the other datasets.

4.5 How much LAION data is the ImageNet pre-training worth?

Figure 2 compares the ImageNet distribution with LAION. Pre-training CLIP on the ImageNet distribution (with template captions) outperforms LAION-1m by a large margin. Findings from Figure 1 with the same pre-training loss are now extended to different losses in Figure 2, *i.e.*, the gap between the supervised ImageNet (with template captions) pre-training and the contrastive LAION-1m pre-training shrinks as more data for the downstream task are available. Interestingly, pre-training CLIP on LAION-1m is only as good as ImageNet with Flickr captions with half of the data. We also scale LAION pre-training size in Figure 3 to see




Figure 6: **CLIP vs. SimCLR for pre-training.** Overall we observe that SimCLR pre-training leads to a better downstream transfer accuracy than using CLIP pre-training on the same dataset. These differences are more pronounced in the few-shot setting.

if LAION can outperform ImageNet pre-training and downstream performance. Figure 3 shows that including $15\times$ more data from LAION outperforms ImageNet pre-training with template captions only on CIFAR100. However, DTD, REAL, and CLIPART need $2000\times$ more data from LAION to match or outperform ImageNet pre-training. Even including $2000\times$ more data did not help CALTECH101. ImageNet pre-training also outperforms LAION-2B on PETS by a large margin. This is probably because PETS and ImageNet both share many samples of pets like dog breeds.

4.6 How does the downstream performance improve as more data is available for pre-training?

Can we expect that more pre-training data implies a better performance, or can the pre-training effectiveness saturate at some point? We fix the pre-training distribution to YFCC and LAION and compare pre-training on 2.7m samples with 15m samples. We also extended our experiments to see the effect of extreme sample sizes and include ViT-B/32 CLIP model trained on 2b samples from LAION. Figure 4 shows that increasing the size of the dataset used for pre-training results in higher downstream transfer accuracy. However, the magnitude of the improvement varies across different downstream datasets. While increasing the pre-training size of YFCC and LAION improves the CIFAR100 performance by a large margin, this improvement is more modest for the rest of the downstream datasets. Specifically including 2 billion samples from LAION does not help much for CALTECH101 and PETS. Similarly to the findings in Figure 1, in larger sizes, we observe more noticeable differences in downstream performance in the few-shot regime. The difference in the absolute accuracy when more data is available for fine-tuning is usually smaller. However, in contrast to the findings by Abnar et al. (2021), pre-training on the extremely large LAION-2B still manages to boost the downstream performance in the full fine-tuning mode. Figure 4 shows that LAION pre-training outperforms YFCC pre-training on downstream tasks for both 2.7m and 15m subsets of the datasets. Interestingly, pre-training on LAION-2.7m performs similarly to a much larger size of YFCC-15m pre-training, highlighting the efficiency of LAION distributions. Using an extremely large dataset of LAION-2B improves the performance by a significant margin in the few-shot regime for CIFAR100. While differences in absolute accuracy are smaller if more data is available for fine-tuning, LAION-2B pre-training still performs consistently better.

4.7 Effect of perturbing loss

In this section, we replace the pre-training loss from language-image contrastive in CLIP with image-only contrastive loss in SimCLR (Chen et al., 2020). Figure 5 highlights that our observations from Figure 1 are now extended to image-only pre-training, *i.e.*, changing the pre-train dataset leads to differences only in the few-shot downstream performance. Next, we evaluate the distinction between pre-training using CLIP and SimCLR. The results are presented in Figure 6 Overall we find that models pre-trained with SimCLR have better downstream transfer accuracy than models pre-trained with CLIP in the few-shot regime.

Similarly to our observations regarding the effect of the pre-training data distribution (Figures 1, 2, and 5), the absolute accuracy difference is smaller when more data is used for fine-tuning. We note that this is different from what (Santurkar et al., 2022) observed. However, we suspect this difference is because we are fine-tuning all model parameters while they only consider a linear classifier.

The difference in the downstream transfer accuracy for CLIP and SimCLR pre-training varies across different datasets. While SimCLR is only marginally better than CLIP for CIFAR100, the difference is significantly larger for DTD and CALTECH101, especially in the few-shot setting.

4.8 Effect of data curation: ImageNet captioning

We compare CLIP models pre-trained on LAION with CLIP models pre-trained on the following two versions of the curated ImageNet dataset:

- IN1K-Flickr-Captions: This is a subset of the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) 2012 training set, paired with the original image title, description, and tags from Flickr. Therefore, we can use it for CLIP pre-training. To construct this dataset, Fang et al. (2022) start from 14,197,122 image URLs in the ImageNet fall 2011 release, and filter to only include images from Flickr. Next, they restrict the images to the 1,000 classes included in the 2012 ImageNet competition, run the image deduplication routine, and remove text containing profanity. As a result, the dataset of 463,622 images is left along with the newly obtained corresponding text data.
- IN1K-Template-Captions: This dataset includes all data in the ImageNet dataset, paired with templated captions, e.g., “a photo of a classname”. This allows us to use CLIP pre-training but on clean images and text. In terms of ImageNet accuracy, this training scheme is very similar to standard supervised training. However, this is now a controlled experiment as we are always using CLIP pre-training.

4.9 Other architectures

In order to see the effect of architecture on the observed trends, we extend the results to the effect of pre-training distribution in Figure 14 to include Vision Transformers. To do so, we used ViT-B/32 released checkpoints trained on LAION-400m and OpenAI-400m, . Figure 7 shows the effect of data distribution on finetune transfer to CIFAR100, DTD, and CALTECH101 when using ViT instead of ResNet-50. While similar to Figure 14 the difference between the fine-tune performance is minimal, we observe that both models perform also very similarly in the few-shot setting. We hypothesize that this observation could be attributed to the similarity between LAION and OpenAI distributions rather than employing a transformer instead of ResNet-50. A controlled study may include to replicate Figure 14 but with ViT, and we leave that for future work.

4.10 Effect of pre-training data distribution: SimCLR instead of CLIP

In contrast to previous experiments with CLIP where we fine-tuned end-to-end from the zero-shot pre-trained model, in SimCLR finetuning we fine-tune using LP-FT (Kumar et al., 2022) because we are no longer able to start with a zero-shot pre-trained model. When we compare to CLIP, we fine-tune both models with LP-FT to facilitate a fair comparison. LP-FT is the following two-step procedure: for each number of shots k we first freeze the encoder and train a classification head from random initialization using k examples per-class




Figure 7: **Effect of the pre-training data distribution: ViT instead of ResNet-50** While similar to Figure 14 the difference between the fine-tune performance is minimal, we observe that both models perform also very similarly in the few-shot setting. We hypothesize that this observation could be attributed to the similarity between LAION and OpenAI distributions rather than employing transformer instead of ResNet-50.

from the downstream task. In the second step, we initialize the classification head with this linear probe (LP) then unfreeze all weights and finetune (FT) the whole model.

5 Discussion, Limitations, and Future Work

Discussion. As better pre-trained models become available, and more workloads shift from training from scratch to fine-tuning, understanding the transfer learning paradigm becomes increasingly important. Presumably, in the future, a sea of pre-trained models will be available for download from the Internet. Therefore, researchers and practitioners will be faced with the question of where to begin. It will be important to make this choice well, but also to understand to what extent this choice matters. Overall we have observed that different pre-training distributions and methods can lead to differences in downstream transfer accuracy. However, these differences are the largest in the few-shot transfer regime. If many images are used for fine-tuning these differences are mostly diminished. Moreover, while different pre-training decisions lead to similar accuracy in the high-shot regime, they still outperform training from scratch in the setting we consider. We also observed that the pre-training method affects performance on downstream transfer (supervised vs. CLIP vs. SimCLR) and including more pre-training data may compensate for the performance gap between training methods.

Limitations and Future work. There are a number of limitations in our study. For one, we consider only end-to-end fine-tuning, because this method produces the highest accuracy. However, if compute is limited, one may choose to instead use linear probing or other lightweight fine-tuning methods. So far this is not addressed in our study. Another limitation is that we did not do an exhaustive hyperparameter sweep for pre-training. While fine-tuning is cheaper and we are therefore able to do a grid search, for pre-training we are mostly limited to using existing checkpoints. While we think that this reflects a realistic setting, in the future we wish to also better understand the role of hyperparameters.

In addition to the mentioned limitations, future works might include extending experiments to include different samples of ImageNet. One example may include subsets of ImageNet-21K (2.7m in Figure 14 and 15m in Figure 3) and respective comparison to Shutterstock and LAION distributions. Given our observation of the role of data curation, we also hope that our findings stimulate further direction toward creative methods for dataset curation.

6 Acknowledgements

This work is in part supported by the Austrian Marshall Plan Foundation, Google Cloud Research Credit, Rudolf Chaudoire Programm from the Faculty of Electrical and Information Engineering at Graz University of Technology, Graz Center for Machine Learning at Graz University of Technology (GraML), and NSF IIS 1652052, IIS 17303166, DARPA N66001-19-2-4031, DARPA W911NF-15-1-0543 and gifts from Allen Institute for Artificial Intelligence.

References

- Makerere University AI Lab. Cassava leaf disease classification, 2021. <https://www.kaggle.com/competitions/cassava-leaf-disease-classification/overview>. Accessed: 2022-10-20.
- Common crawl. <https://commoncrawl.org/>. Accessed: 2022-09-20.
- Abnar, S., Dehghani, M., Neyshabur, B., and Sedghi, H. Exploring the limits of large scale pre-training. *arXiv preprint arXiv:2110.02095*, 2021.
- Alayrac, J.-B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022.
- Beery, S., Van Horn, G., and Perona, P. Recognition in terra incognita. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 456–473, 2018.
- Bolya, D., Mittapalli, R., and Hoffman, J. Scalable diverse model selection for accessible transfer learning. *Advances in Neural Information Processing Systems*, 34:19301–19312, 2021.
- Changpinyo, S., Sharma, P., Ding, N., and Soricut, R. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3558–3568, 2021.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.
- Cherti, M. and Jitsev, J. Effect of pre-training scale on intra-and inter-domain full and few-shot transfer learning for natural and medical x-ray chest images. *arXiv preprint arXiv:2106.00116*, 2021.
- Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., and Vedaldi, A. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3606–3613, 2014.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Desai, K., Kaul, G., Aysola, Z., and Johnson, J. Redcaps: Web-curated image-text data created by the people, for the people. *arXiv preprint arXiv:2111.11431*, 2021.
- Deshpande, A., Achille, A., Ravichandran, A., Li, H., Zancato, L., Fowlkes, C., Bhotika, R., Soatto, S., and Perona, P. A linearized framework and a new benchmark for model selection for fine-tuning. *arXiv preprint arXiv:2102.00084*, 2021.
- Djolonga, J., Yung, J., Tschannen, M., Romijnders, R., Beyer, L., Kolesnikov, A., Puigcerver, J., Minderer, M., D’Amour, A., Moldovan, D., et al. On robustness and transferability of convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16458–16468, 2021.

- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- Ericsson, L., Gouk, H., and Hospedales, T. M. How well do self-supervised models transfer? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5414–5423, 2021.
- Fang, A., Ilharco, G., Wortsman, M., Wan, Y., Shankar, V., Dave, A., and Schmidt, L. Data determines distributional robustness in contrastive language image pre-training (clip). *arXiv preprint arXiv:2205.01397*, 2022.
- Fei-Fei, L., Fergus, R., and Perona, P. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, pp. 178–178. IEEE, 2004.
- Gontijo-Lopes, R., Smullin, S., Cubuk, E. D., and Dyer, E. Tradeoffs in data augmentation: An empirical study. In *International Conference on Learning Representations*, 2020.
- Goyal, P., Caron, M., Lefaudeaux, B., Xu, M., Wang, P., Pai, V., Singh, M., Liptchinsky, V., Misra, I., Joulin, A., et al. Self-supervised pretraining of visual features in the wild. *arXiv preprint arXiv:2103.01988*, 2021.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Helber, P., Bischke, B., Dengel, A., and Borth, D. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2019.
- Ilharco, G., Wortsman, M., Wightman, R., Gordon, C., Carlini, N., Taori, R., Dave, A., Shankar, V., Namkoong, H., Miller, J., Hajishirzi, H., Farhadi, A., and Schmidt, L. Openclip. July 2021. doi: 10.5281/zenodo.5143773. URL <https://doi.org/10.5281/zenodo.5143773>. If you use this software, please cite it as below.
- Islam, A., Chen, C.-F. R., Panda, R., Karlinsky, L., Radke, R., and Feris, R. A broad study on the transferability of visual representations with contrastive learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8845–8855, 2021.
- Jia, C., Yang, Y., Xia, Y., Chen, Y.-T., Parekh, Z., Pham, H., Le, Q., Sung, Y.-H., Li, Z., and Duerig, T. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pp. 4904–4916. PMLR, 2021.
- Kim, D., Wang, K., Sclaroff, S., and Saenko, K. A broad study of pre-training for domain generalization and adaptation. *arXiv preprint arXiv:2203.11819*, 2022.
- Kornblith, S., Shlens, J., and Le, Q. V. Do better imagenet models transfer better? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2661–2671, 2019.
- Krizhevsky, A., Nair, V., and Hinton, G. Cifar-100 and cifar-10 (canadian institute for advanced research), 2009. URL <http://www.cs.toronto.edu/~kriz/cifar.html>. MIT License.
- Kumar, A., Raghunathan, A., Jones, R., Ma, T., and Liang, P. Fine-tuning can distort pretrained features and underperform out-of-distribution. *arXiv preprint arXiv:2202.10054*, 2022.
- Loshchilov, I. and Hutter, F. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.

- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Miller, J. P., Taori, R., Raghunathan, A., Sagawa, S., Koh, P. W., Shankar, V., Liang, P., Carmon, Y., and Schmidt, L. Accuracy on the line: on the strong correlation between out-of-distribution and in-distribution generalization. In *International Conference on Machine Learning*, pp. 7721–7735. PMLR, 2021.
- Mu, N., Kirillov, A., Wagner, D., and Xie, S. Slip: Self-supervision meets language-image pre-training. *arXiv preprint arXiv:2112.12750*, 2021.
- Neyshabur, B., Sedghi, H., and Zhang, C. What is being transferred in transfer learning? *Advances in neural information processing systems*, 33:512–523, 2020.
- Nguyen, C., Hassner, T., Seeger, M., and Archambeau, C. Leep: A new measure to evaluate transferability of learned representations. In *International Conference on Machine Learning*, pp. 7294–7305. PMLR, 2020.
- Nguyen, T., Ilharco, G., Wortsman, M., Oh, S., and Schmidt, L. Quality not quantity: On the interaction between dataset design and robustness of clip. *arXiv preprint arXiv:2208.05516*, 2022.
- Parkhi, O. M., Vedaldi, A., Zisserman, A., and Jawahar, C. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 3498–3505. IEEE, 2012.
- Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., and Wang, B. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1406–1415, 2019.
- Pham, H., Dai, Z., Ghiasi, G., Liu, H., Yu, A. W., Luong, M.-T., Tan, M., and Le, Q. V. Combined scaling for zero-shot transfer learning. *arXiv preprint arXiv:2111.10050*, 2021.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pp. 8748–8763. PMLR, 2021.
- Raghu, M., Zhang, C., Kleinberg, J., and Bengio, S. Transfusion: Understanding transfer learning for medical imaging. *Advances in neural information processing systems*, 32, 2019.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- Santurkar, S., Dubois, Y., Taori, R., Liang, P., and Hashimoto, T. Is a caption worth a thousand images? a controlled study for representation learning. *arXiv preprint arXiv:2207.07635*, 2022.
- Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., and Komatsuzaki, A. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021.
- Sharma, P., Ding, N., Goodman, S., and Soricut, R. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2556–2565, 2018.
- Srinivasan, K., Raman, K., Chen, J., Bendersky, M., and Najork, M. Wit: Wikipedia-based image text dataset for multimodal multilingual machine learning. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 2443–2449, 2021.
- Sun, C., Shrivastava, A., Singh, S., and Gupta, A. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*, pp. 843–852, 2017.


- Taori, R., Dave, A., Shankar, V., Carlini, N., Recht, B., and Schmidt, L. Measuring robustness to natural distribution shifts in image classification. *Advances in Neural Information Processing Systems*, 33:18583–18599, 2020.
- Thomee, B., Shamma, D. A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., and Li, L.-J. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016.
- Wortsman, M., Ilharco, G., Kim, J. W., Li, M., Kornblith, S., Roelofs, R., Gontijo-Lopes, R., Hajishirzi, H., Farhadi, A., Namkoong, H., and Schmidt, L. Robust fine-tuning of zero-shot models. *arXiv preprint arXiv:2109.01903*, 2021. <https://arxiv.org/abs/2109.01903>.
- You, K., Liu, Y., Wang, J., and Long, M. Logme: Practical assessment of pre-trained models for transfer learning. In *International Conference on Machine Learning*, pp. 12133–12143. PMLR, 2021.
- Zhai, X., Wang, X., Mustafa, B., Steiner, A., Keysers, D., Kolesnikov, A., and Beyer, L. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18123–18133, 2022.



ROME, ITALY - JUNE 13, 2015: Rome hosts a popular Pride celebration - Rome Gay Pride on June 13, 2015. Rome Gay Pride parade takes place on this day, drawing tho...



Abstract Background. Design Template. Modern Pattern. Vector Illustration For Your Design.



Muscular man. Gym characters sport people making exercises bodybuilders posing muscular athletes



Abstract violet fractal composition. Magic explosion star with particles motion illustration.



Conceptual hand writing showing Applied. Business photo text put to practical use as opposed to being theoretical Be applicable.



Baltimore Maryland USA skyline silhouette flat design vector illustration



rocky coast on Atlantic ocean in France, Normandy



Happy rich kanelbulle mascot design carries money bags




Mountain Fuji at top of mountain



Close-up of sliced deli meats with vegetables and baguette on a slate board.



Young woman hitchhiking on countryside road. Traveler woman hitchhiking along lonely road. Pretty young woman tourist hitchhiking. Left alone on the road and lost




Infographics with a pie chart for business or presentation trends



giant centipedes hiding in black ,leather shoes



Blonde d'aquitaine cows at the wash outs of Goeree-Overflakkee at the Haringvliet in the Netherlands



elephant in the savanna of Africa



Green high mountain meadow with rocks closeup as natural background

Figure 8: Random training samples from Shutterstock




lofoten archipelago by <usr>



foggy night in the vancouver forest



bubba is so unbelievably cute when she's sleeping!



duchesse satin wedding guest dress- featuring bonus pockets!



the kids got t-shirts



your present condition!



paused the x-men at just the right time.



homemade flammkuchen for dinner...



in a field of yellow and green



i'm drunk, and this is lucy.



eerie section of trail on a long-forgotten country backroad. - long path, catskills park ny



my handsome new neighbour



dressing up for the family photo



shot from our airbnb porch view on oia on santorini in greece



completed a small remodel of the half bath. first timer.



such a pretty girl

Figure 9: Random training samples from Redcaps



Juniper Berries Eastern red cedar (Juniperus virginiana) laden with berries at the High Line



Rendlesham Forest Suffolk Spider



Energy Saving 20W CFL bulb equivalent to 100W incandescent bulb. It's like magic.



PISM's analysts On 5 October 2012, in the presence of PISM Council members and Directors, the PISM staff inaugurated the autumn season at Warecka Street. Our gu...



Parque Mayer (Lisboa/ Portugal) Obrigado por todas as visitas, comentários e dicas ;-). Thanks for all your visits, comments and advice ;-).



Roof Repairs Roof Repairs, Lester Public Library, Two Rivers, Wisconsin - www.greatlakesroofing.net/



Alaska Trip-Glacier Bay, Sitka 1976 Glacier Bay 062 My blog here Musings from the Silent Generation Glenn



Nash, North Dakota Nash, North Dakota. From everydot.com/.



Point Mugu State Park On the way to Santa Barbara



Outdoor practice Heikki Karinen teaches how to do makeshift bandages



dancing monk note the audience expressions!



Boats in Porvoo Plus more of those cute red storage houses.



Lowland Paca This Lowland Paca, Cuniculus paca, was photographed in Panama, as part of a research project utilizing motion-activated camera-traps. You are invite...



Cuff Point The old haunt, from the corner of Hackney Road



QLD Police Traffic Branch Commodore SS with customer!



CHELSEA FOOTBALL CLUB Chelsea Magazine - Issue 63, November 2009

Figure 10: Random training samples from YFCC



Busy Slaying Vampires Mens T-Shirt



Webcam site Stripchat




Yellow sandals for women pointy and low heeled Beatnik Françoise Mustard



OOZOO TIMEPIECES bordeaux croco leather strap



3 Bedrooms Terraced House for sale in Eastbourne Road, Walton, Liverpool, Merseyside, L9



Paris Marriott Rive Gauche Hotel & Conference Center photo 27



Minimum Wage Barbie



Diamond Art Deco 18 Karat White Gold Dangle Drop Earrings



Video editing with laptop. Professional editor adding special effects or color grading footage for commercial film or movie.



I'm a Rugby Referee - Men's Premium Hoodie



Marika Airbrush Yoga Leggings



Julianne Hough: 'Footloose' Premiere with Kenny Wormald!



Collins, Jackie Lovers and Gamblers



All you need is Oils SVG



Girls Dora the Explorer Costume - HalloweenCostumes4U.com - Kids Costumes



SCDMS SHIP TO SHORE COVERAGE

Figure 11: Random training samples from LAION



<PERSON> `` The wolf and the lamb shall feed together, and the lion shall eat straw like the bullock: and dust shall be the serpent's meat. They shall not hurt no...



Islamic vector geometric ornaments based on traditional arabic art. Oriental seamless pattern. Muslim mosaic. Turkish, Arabian tile on a white background. Mosque ...



Biker girl in a leather jacket on a black and red color motorcycle



Light Touch Wall digital marketing activation at the Canberra Centre.



Today's wedding dress inspiration brings us fabulous bridal gowns from creative designer <PERSON>. The Divine Affection latest bridal collection of <PERSON> wedd...



Illustration of hand holding the id card. Vector illustration flat design.



Easy Cabbage Rolls that are <PERSON>, <PERSON> and have no rice! <PERSON> budget friendly comfort food recipe adapted from my Russian grandmother!



<PERSON>: U. <PERSON> in United States Army. First <PERSON> appointed to that position. First, &, so far, only <PERSON> to serve on Joint Chiefs of Staff. Black H...



Wedding rings on a bouquet of roses stock photos




<PERSON> tattoo, the American number 23 from Akron, United States



All Balls Swinging Arm Bearing Kit for Yamaha XT225 | XT250 Serow 1993 to 2007



Search the hidden word, the simple educational kid game. stock illustration




Different types of photo frames with circles and squares on the wall - background template stock illustration



The Russian army entering Prussia, 1914 : News Photo



The art of good drinking



Modern Bathroom Makeovers 20 Design Ideas For a Small Bathroom Remodel. Modern Bathroom Designs On A Budget Minimalist Small Bathrooms, Modern Small Bathrooms, Mo...

Figure 12: Random training samples from Conceptual Captions



Japanese chamber pot from the Edo period



Ratel



15th race in 1982



Dedication plaque at Oregon Dunes Overlook, Oregon Dunes National Recreation Area



Robinson in March 2018



85-15 Wareham Place, Donald Trump's childhood home



Hugo van der Goes, Saint Luke Drawing the Virgin, c. 1470-80. National Museum of Ancient Art, Lisbon



Museum of Arts & Design at 2 Columbus Circle, nearly completed in July 2008. A piece by David Dunlap's in the NY Times reveals that the appearance of the letter ...



Paul McAuley at Worldcon 2005 in Glasgow



Maui Veterans Highway shown just above Kealia Pond National Wildlife Refuge as it enters Kihei



Construction of restrooms and locker rooms with east side stands and pavilion



McLaren 600LT



Buzz Aldrin received praise for his performance.



James Barber House



A Fiat G.91PAN, in service in the Freccce Tricolori from 1963 to 1982



Alexa Stirling, c. 1919

Figure 13: Random training samples from WIT

Appendix

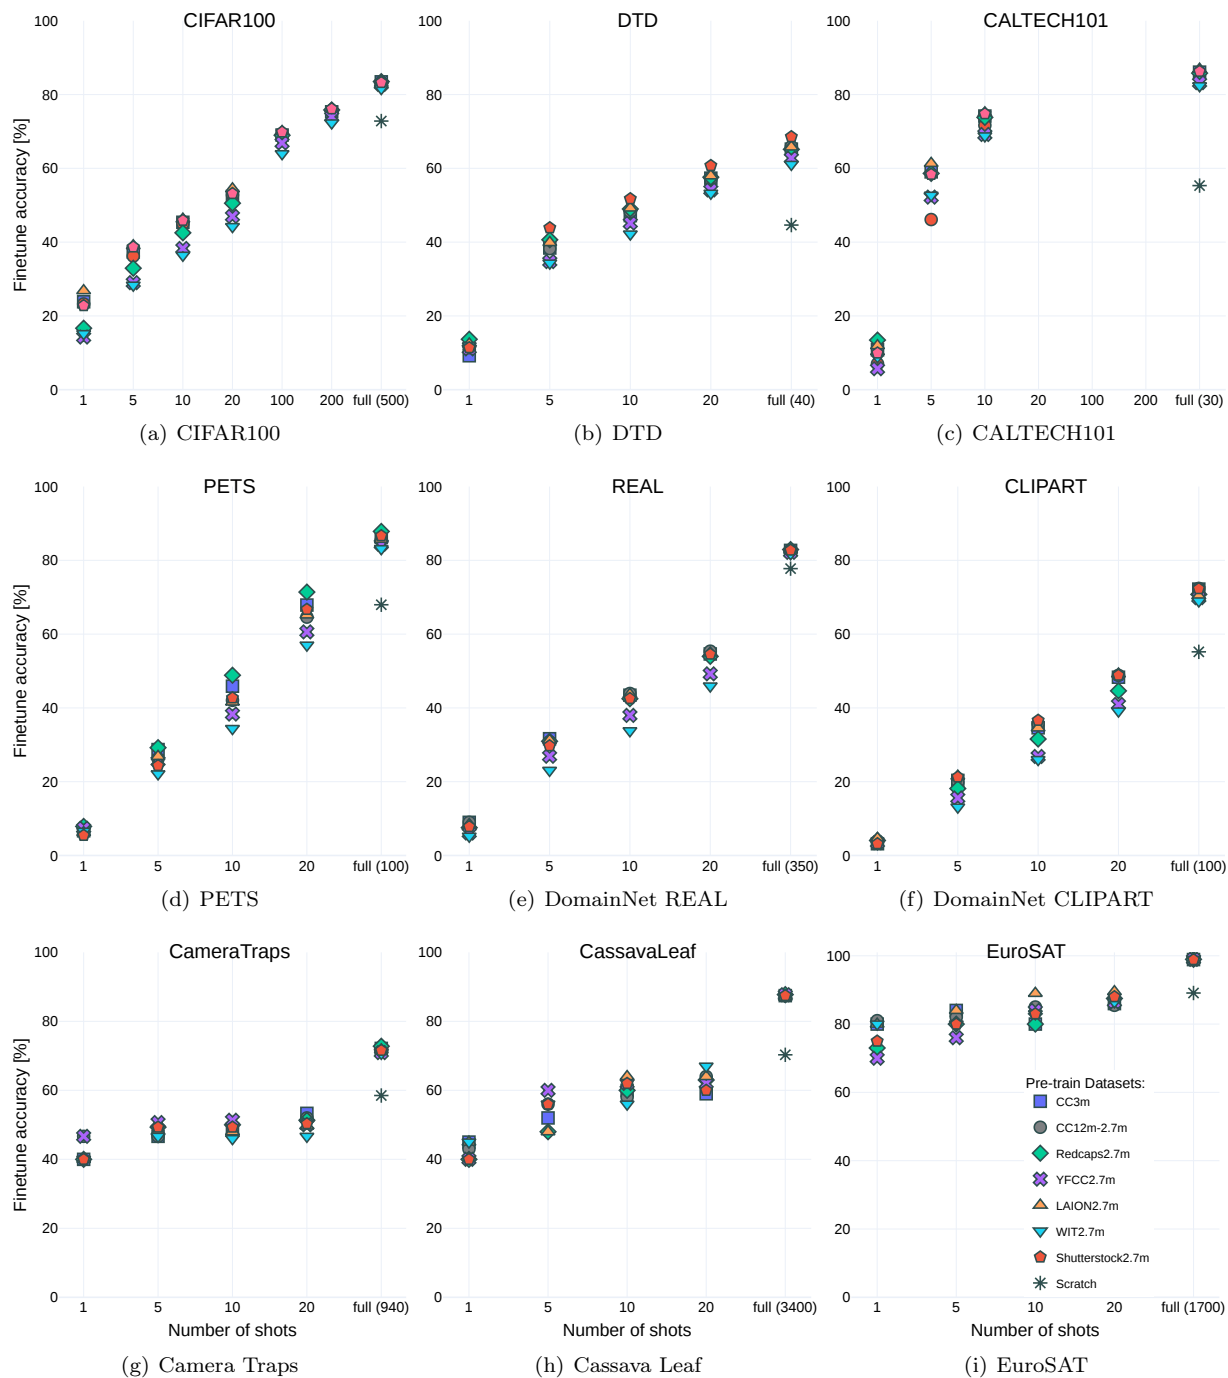


Figure 14: **Effect of the pre-training data distribution.** While Figure 1 shows the aggregated results all downstream datasets, here we include the performance for each pair of (pretraining, downstream) datasets in detail. In the low-shot setting, different pre-training datasets lead to noticeable differences in downstream performance. If many samples are available for fine-tuning, the difference in absolute accuracy between models pre-trained on different sources largely evaporates.

A Effect of the pre-training data distribution

Figure 14 shows a detailed for aggregated results shown in Figure 1. In the low-shot setting, different pre-training datasets lead to noticeable differences in downstream performance. If many samples are available for fine-tuning, the difference in absolute accuracy between models pre-trained on different sources largely evaporates.

Figure 14 compares different data sources for pre-training. While Shutterstock shows superior performance on the first six datasets (except for PETS), the best pre-training distribution changes between Camera Traps, Cassava Leaf, and EuroSAT. Changing the pre-training dataset leads to noticeable differences in the downstream low-shot performance of nine datasets.

B Training Details

B.1 CLIP training

Our CLIP models are trained from scratch on each of the pre-training datasets unless otherwise mentioned and follow the training code from the OpenCLIP GitHub repository (Ilharco et al., 2021). CLIP models are trained using AdamW optimizer (Loshchilov & Hutter, 2017) with default PyTorch parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, batch size 1024, and weight decay of 0.1. For learning rate, we start with a learning rate of 10^{-3} and apply a cosine-annealing learning rate schedule (Loshchilov & Hutter, 2016) with 5,000 steps warm-up. We use the same data augmentations as in (Radford et al., 2021).

B.2 SimCLR training

Our SimCLR implementation closely follows the training code from the SLIP (Mu et al., 2021). SimCLR models are also trained for 16 epochs from scratch using AdamW optimizer (Loshchilov & Hutter, 2017) with $\beta_1 = 0.9$, $\beta_2 = 0.98$, $\epsilon = 10^{-8}$, batch size 1024, and weight decay of 0.1. we start with a learning rate of 10^{-3} and apply a cosine-annealing learning rate schedule (Loshchilov & Hutter, 2016) with 2 epochs of warm-up. The hidden dimension of SimCLR MLP projection head is set to 4,094 and the output embedding dimension of MLP projection head is set to 256.

B.3 Finetuning details

Each pretrained model is finetuned on the specific downstream task for 128 epochs while the learning rate is mostly from 0.0001, 0.0003, 0.001, 0.003 as starting and applying a cosine-annealing learning rate schedule (Loshchilov & Hutter, 2016) with 500 steps warm-up and batch size of 128. For each fine-tuning, we choose the best-performing result on the test set among the performed grid search. We use the implementation from the WiSE-FT GitHub repository for fine-tuning, where we have only one model and $\alpha = 1$ (Wortsman et al., 2021). For a list of all 4000 experiments, including their hyperparameters and performance see https://github.com/AnonymousMLSubmission/DataDistributionTransfer/blob/main/Hyperparameters_results.csv

C Effect of pre-training data distribution: SimCLR instead of CLIP

In contrast to previous experiments with CLIP where we fine-tuned end-to-end from the zero-shot pre-trained model, in SimCLR finetuning we fine-tune using LP-FT (Kumar et al., 2022) because we are no longer able to start with a zero-shot pre-trained model. When we compare to CLIP, we fine-tune both models with LP-FT to facilitate a fair comparison. LP-FT is the following two-step procedure: for each number of shots k we first freeze the encoder and train a classification head from random initialization using k examples per-class from the downstream task. In the second step, we initialize the classification head with this linear probe (LP) then unfreeze all weights and finetune (FT) the whole model.

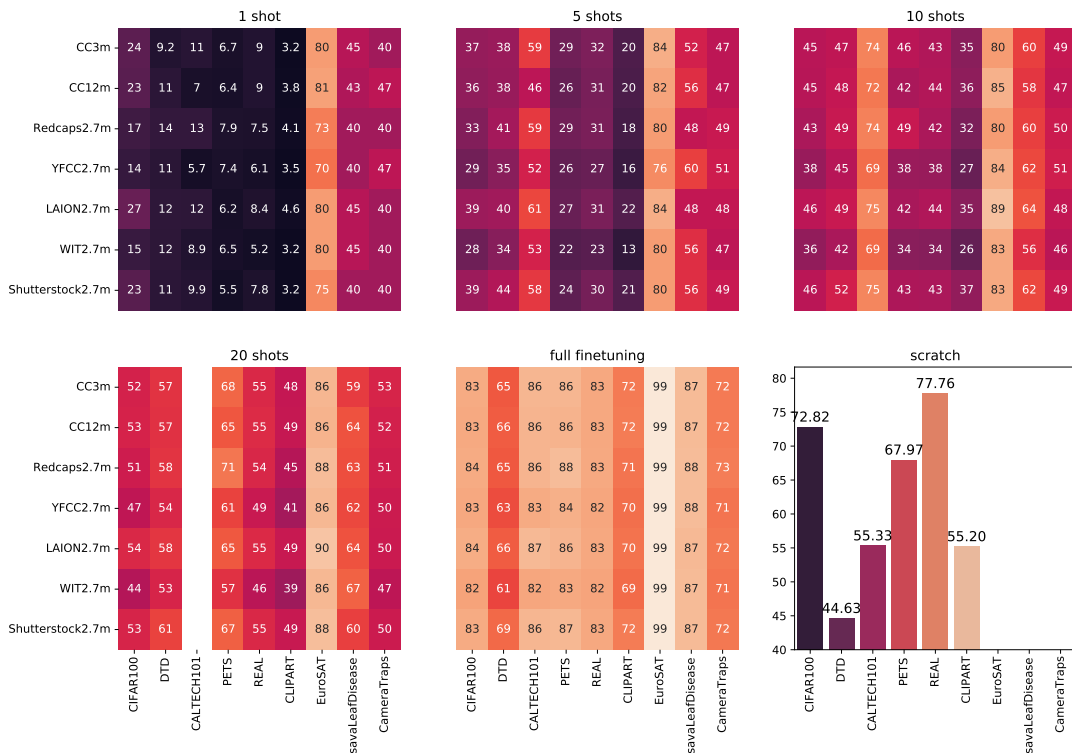


Figure 15: **Effect of pre-training data distribution: a better view.** We change the presentation of Figure 14 for a better view of exact performance numbers on different data distributions and datasets.

D Extended related works

Transfer learning is widely used in deep learning research and practice and has become a cornerstone in both computer vision and natural language processing. Through the years, there have been many questions on why transfer helps and how to choose a good pre-trained model to transfer from. [Neyshabur et al. \(2020\)](#) separated the effect of feature reuse from that of learning low-level pre-training data statistics. [Raghu et al. \(2019\)](#) investigate the similarity of the pre-training and downstream datasets by looking into medical datasets and found that transfer learning from ImageNet pre-trained models shows little benefit in performance. [Ericsson et al. \(2021\)](#) studied the downstream performance of self-supervised models and found that the best self-supervised models of that time could outperform supervised pre-training as an upstream source of knowledge transfer and that the performance of self-supervised models on ImageNet is indicative of downstream performance on natural image classification tasks. Similarly, [Islam et al. \(2021\)](#) found that contrastively trained models consistently outperform standard cross-entropy models in transfer learning. [Goyal et al. \(2021\)](#) showed that self-supervised models outperform supervised models on ImageNet, even when trained on random and uncurated images from the web. Moreover, they showed that these models are also good at few shot learning by achieving 77.9% top-1 accuracy using only 10% on ImageNet.

Building on contrastive techniques, [Radford et al. \(2021\)](#) introduced CLIP which learns a joint embedding space for both images and their descriptive captions, making it possible to effectively leverage a large-scale dataset from the Internet. Flamingo ([Alayrac et al., 2022](#)), a visual language model, is another successful example in the line of multimodal models and enables visual question answering and image captioning. CLIP and similar models like ALIGN ([Jia et al., 2021](#)), BASIC ([Pham et al., 2021](#)), and LiT ([Zhai et al.,](#)

2022) demonstrated unprecedented robustness to challenging data distribution shifts. This accomplishment raised questions on the probable sources of such robustness—whether this robustness is caused by language supervision, the pre-training data distribution, size, or contrastive loss functions.

Fang et al. (2022) investigated this question and found that the diverse training distribution is the main cause of the robustness properties of CLIP. Nguyen et al. (2022) explored the role of the pre-training dataset for CLIP with a testbed of six pre-training sources, finding that no single pre-training dataset consistently performs best. In recent work, Santurkar et al. (2022) carefully investigated the effect of language supervision in CLIP-like models, finding it an important factor if the pre-training dataset is large and the captions are descriptive enough. Unlike their work, we consider end-to-end fine-tuning which result in higher accuracy. Cherti & Jitsev (2021) study the effect of scaling the pre-training model and data for both in-domain and out-of-domain transfer. They conduct supervised pre-training while varying pre-training model size and data source (ImageNet-1k/21k or large medical chest XRay datasets), and transfer pre-trained models to different natural or medical targets. They find that, when performing transfer to large X-Ray targets, pre-training on natural ImageNet-21k is as good or better than pre-trained medical X-Ray data. Djolonga et al. (2021) also investigate the impact of the pre-training data size and model scale, finding that increasing both the training set and model sizes significantly improve the distributional shift robustness.

E Datasets

E.1 Downstream tasks

We have used 9 different downstream datasets. Table 5 describes the first six datasets in Figure 14. While these six datasets are internet-crawled datasets and are more common in transfer learning in computer vision benchmarks, we include three new downstream datasets that are domain-specific, *i.e.*, the dataset is created after a specific challenge is defined in a specific domain.

- EuroSAT (Helber et al., 2019): The task is to classify land use and land cover based on Sentinel-2 satellite images. The dataset covers 13 spectral bands and consists of 10 classes within a total of 27,019 labeled and geo-referenced images. we create an 80%-20% random class-balanced split with the provided dataset.
- Cassava Leaf Disease Classification (Cas): The dataset contains 21,397 images from the Kaggle competition, to give farmers access to methods for diagnosing plant diseases. The images are labeled as healthy or as one of four different diseases. we split the dataset with 80%-20% random class-balanced ratio for train and test, respectively.
- Caltech Camera Traps-20 (Beery et al., 2018): CCT-20 contains 57,864 images in 15 classes, taken from camera traps deployed to monitor animal populations. Classes are either single species *e.g.*, "Coyote") or groups of species, *e.g.*, "Bird"). CCT-20 is a subset of the iWildCam Challenge 2018, whose yearly editions have been hosted on Kaggle. Here we study the subset of CCT-20 that come from the same locations ⁴, including 14,071 and 16,395 images for train and test respectively.

E.2 Pre-training datasets

Our study covers 7 pre-training datasets as follow:

- YFCC: Our experiments mostly include YFCC-2.7M, a random subset of YFCC-15M. The 15M subset of the YFCC-100M dataset (Thomee et al., 2016) was filtered to only include images with English titles or descriptions. The dataset contains 14,829,396 images with natural language captions associated with each image. The images and captions are collected from Flickr.

⁴"Cis" in the main dataset refers to images from locations seen during training, and "trans" refers to new locations not seen during training

Table 5: Details on the downstream datasets used in the experiments.

Downstream Task	Description
CIFAR100	The task consists in classifying natural images (100 classes, with 500 training images each). Some examples include apples, bottles, dinosaurs, and bicycles. The image size is 32x32.
DTD	The task consists in classifying images of textural patterns (47 classes, with 120 training images each). Some of the textures are banded, bubbly, meshed, lined, or porous. The image size ranges between 300x300 and 640x640 pixels.
CALTECH-101	The task consists in classifying images of objects (9144 images in 101 classes plus a background clutter class), including animals, airplanes, chairs, or scissors. The image size varies, but it typically ranges from 200-300 pixels per edge.
PETS	The task consists in classifying images of cat and dog breeds (7000 images in 37 classes). Images dimensions are typically 200 pixels or larger
REAL	The task is a subset of larger DomainNet from six distinct domains, including photos (real), painting, clipart, quickdraw, infograph, and sketch. Total size of 172,000
CLIPART	The task is a subset of larger DomainNet from six distinct domains, including photos (real), painting, clipart, quickdraw, infograph, and sketch. Total size of 172,000

- LAION (Schuhmann et al., 2021): The images and corresponding alt-texts come from web pages collected by Common Crawl (Com) between 2014 and 2021. We randomly select a subset of 2.7M and 15M samples for our experiments.
- Redcaps (Desai et al., 2021): Redcaps contains 11,882,403 examples from 350 manually curated subreddit collected between 2008 and 2020. The subreddits are selected to contain a large number of image posts that are mostly photographs and not images of people.
- Shutterstock: 11,800,000 images and captions from the Shutterstock website.
- Conceptual Captions-3m (Sharma et al., 2018): The raw descriptions in Conceptual Captions are harvested from the alt-text HTML attribute associated with web images. This dataset contains 2,799,553 samples, denoted as CC_2.7m in the plots.
- Conceptual Captions-12m (Changpinyo et al., 2021): A dataset with 12 million image-text pairs. It is larger than CC_2.7m and covers a much more diverse set of visual concepts. We randomly select 2.7M samples from this dataset, denoted as CC_12_2.7m.
- WIT (Srinivasan et al., 2021): Image-text pairs come from Wikipedia pages. We use reference description as the source of text data and obtain 5,038,295 examples in total after filtering to include only the English language.

Table 6 shows their main source and total size. We also show some examples of image-caption pairs randomly selected from Shutterstock in Figure 16, Redcaps in Figure 17, YFCC-15m in Figure 18, LAION-15m in Figure 19, Conceptual Captions in Figure 20, and WIT in Figure 21. Table 7 also shows the most common words in captions of these pre-training datasets.

Looking at Redcaps samples in Figure 17 and also the top 20 captions shows many samples of animals. This is showing why Redcaps perform better on PETS. Samples from WIT in Figure 21 and also its top 20 words mostly featuring geographical locations, which is rare in our downstream task, hence performing worst compared to other pre-training distributions. Shutterstock top 20 words also include words like "pattern", "texture", "and design" which are close to DTD classes, hence showing superior performance in this downstream task.

Table 6: Details on pre-training datasets


Dataset	Source	Total size
YFCC	Flickr	14,826,000
LAION	Common Crawl	15,504,742
CC-12M	Unspecified web pages	9,594,338
RedCaps	Reddit	11,882,403
WIT	Wikipedia	5,038,295
ShutterStock	ShutterStock	11,800,000
IN1K-Captions	ImageNet	463,622

Table 7: Most common words in captions of pre-training distributions

Pre-training dataset	Top 20 words in 1M sample of captions
Shutterstock	background , vector, illustration, design , icon, pattern , texture , style, woman, concept, hand, color, flower, view, template, line, business, logo, card, symbol
Redcaps	day, today, year, time, cat, plant, friend, anyone, picture, baby, guy, week, dog, home, morning, night, month, way, boy, work
YFCC-15m	photo, day, park, street, city, picture, view, time, world, year, house, state, center, part, garden, shot, image, building, road, museum
LAION-15m	photo, stock, image, black, woman, design, set, vector, white, print, home, men, blue, dress, art, card, sale, gold, bag, cover
CC-12m	illustration, stock, art, design, photo, image, background, room, vector, house, home, woman, wedding, style, photography, royalty, car, fashion, girl, world
CC-3m	background, actor, artist, player, illustration, view, woman, man, football, team, tree, premiere, city, vector, day, girl, beach, game, hand, people
WIT	view, church, station, map, house, building, hall, museum, city, location, street, park, river, state, john, county, town, center, bridge, world



ROME, ITALY - JUNE 13, 2015: Rome hosts a popular Pride celebration - Rome Gay Pride on June 13, 2015. Rome Gay Pride parade takes place on this day, drawing tho...



Abstract Background. Design Template. Modern Pattern. Vector Illustration For Your Design.



Muscular man. Gym characters sport people making exercises bodybuilders posing muscular athletes



Abstract violet fractal composition. Magic explosion star with particles motion illustration.



Conceptual hand writing showing Applied. Business photo text put to practical use as opposed to being theoretical Be applicable.



Baltimore Maryland USA skyline silhouette flat design vector illustration



rocky coast on Atlantic ocean in France, Normandy



Happy rich kanelbulle mascot design carries money bags




Mountain Fuji at top of mountain



Close-up of sliced deli meats with vegetables and baguette on a slate board.



Young woman hitchhiking on countryside road. Traveler woman hitchhiking along lonely road. Pretty young woman tourist hitchhiking. Left alone on the road and lost




Infographics with a pie chart for business or presentation trends



giant centipedes hiding in black ,leather shoes



Blonde d'aquitaine cows at the wash outs of Goeree-Overflakkee at the Haringvliet in the Netherlands



elephant in the savanna of Africa



Green high mountain meadow with rocks closeup as natural background

Figure 16: Random training samples from Shutterstock




lofoten archipelago by <usr>



foggy night in the vancouver forest



bubba is so unbelievably cute when she's sleeping!



duchesse satin wedding guest dress- featuring bonus pockets!



the kids got t-shirts



your present condition!



paused the x-men at just the right time.



homemade flammkuchen for dinner...



in a field of yellow and green



i'm drunk, and this is lucy.



eerie section of trail on a long-forgotten country backroad. - long path, catskills park ny



my handsome new neighbour



dressing up for the family photo



shot from our airbnb porch view on oia on santorini in greece



completed a small remodel of the half bath. first timer.



such a pretty girl

Figure 17: Random training samples from Redcaps



Juniper Berries Eastern red cedar (Juniperus virginiana) laden with berries at the High Line



Rendlesham Forest Suffolk Spider



Energy Saving 20W CFL bulb equivalent to 100W incandescent bulb. It's like magic.



PISM's analysts On 5 October 2012, in the presence of PISM Council members and Directors, the PISM staff inaugurated the autumn season at Warecka Street. Our gu...



Parque Mayer (Lisboa/ Portugal) Obrigado por todas as visitas, comentários e dicas ;-). Thanks for all your visits, comments and advice ;-).



Roof Repairs Roof Repairs, Lester Public Library, Two Rivers, Wisconsin - www.greatlakesroofing.net/



Alaska Trip-Glacier Bay, Sitka 1976 Glacier Bay 062 My blog here Musings from the Silent Generation Glenn



Nash, North Dakota Nash, North Dakota. From everydot.com/.



Point Mugu State Park On the way to Santa Barbara



Outdoor practice Heikki Karinen teaches how to do makeshift bandages



dancing monk note the audience expressions!



Boats in Porvoo Plus more of those cute red storage houses.



Lowland Paca This Lowland Paca, Cuniculus paca, was photographed in Panama, as part of a research project utilizing motion-activated camera-traps. You are invite...



Cuff Point The old haunt, from the corner of Hackney Road



QLD Police Traffic Branch Commodore SS with customer!



CHELSEA FOOTBALL CLUB Chelsea Magazine - Issue 63, November 2009

Figure 18: Random training samples from YFCC



Busy Slaying Vampires Mens T-Shirt



Webcam site Stripchat




Yellow sandals for women pointy and low heeled Beatnik Françoise Mustard



OOZOO TIMEPIECES bordeaux croco leather strap



3 Bedrooms Terraced House for sale in Eastbourne Road, Walton, Liverpool, Merseyside, L9



Paris Marriott Rive Gauche Hotel & Conference Center photo 27



Minimum Wage Barbie



Diamond Art Deco 18 Karat White Gold Dangle Drop Earrings



Video editing with laptop. Professional editor adding special effects or color grading footage for commercial film or movie.



I'm a Rugby Referee - Men's Premium Hoodie



Marika Airbrush Yoga Leggings



Julianne Hough: 'Footloose' Premiere with Kenny Wormald!



Collins, Jackie Lovers and Gamblers



All you need is Oils SVG



Girls Dora the Explorer Costume - HalloweenCostumes4U.com - Kids Costumes




SCDMS SHIP TO SHORE COVERAGE

Figure 19: Random training samples from LAION



<PERSON> `` The wolf and the lamb shall feed together, and the lion shall eat straw like the bullock: and dust shall be the serpent's meat. They shall not hurt no...



Islamic vector geometric ornaments based on traditional arabic art. Oriental seamless pattern. Muslim mosaic. Turkish, Arabian tile on a white background. Mosque ...



Biker girl in a leather jacket on a black and red color motorcycle



Light Touch Wall digital marketing activation at the Canberra Centre.



Today's wedding dress inspiration brings us fabulous bridal gowns from creative designer <PERSON>. The Divine Affection latest bridal collection of <PERSON> wedd...



Illustration of hand holding the id card. Vector illustration flat design.



Easy Cabbage Rolls that are <PERSON>, <PERSON> and have no rice! <PERSON> budget friendly comfort food recipe adapted from my Russian grandmother!



<PERSON>: U. <PERSON> in United States Army. First <PERSON> appointed to that position. First, &, so far, only <PERSON> to serve on Joint Chiefs of Staff. Black H...



Wedding rings on a bouquet of roses stock photos




<PERSON> tattoo, the American number 23 from Akron, United States



All Balls Swinging Arm Bearing Kit for Yamaha XT225 | XT250 Serow 1993 to 2007



Search the hidden word, the simple educational kid game. stock illustration




Different types of photo frames with circles and squares on the wall - background template stock illustration



The Russian army entering Prussia, 1914 : News Photo



The art of good drinking



Modern Bathroom Makeovers 20 Design Ideas For a Small Bathroom Remodel. Modern Bathroom Designs On A Budget Minimalist Small Bathrooms, Modern Small Bathrooms, Mo...

Figure 20: Random training samples from Conceptual Captions



Japanese chamber pot from the Edo period



Ratel



15th race in 1982



Dedication plaque at Oregon Dunes Overlook, Oregon Dunes National Recreation Area



Robinson in March 2018



85-15 Wareham Place, Donald Trump's childhood home



Hugo van der Goes, Saint Luke Drawing the Virgin, c. 1470-80. National Museum of Ancient Art, Lisbon



Museum of Arts & Design at 2 Columbus Circle, nearly completed in July 2008. A piece by David Dunlap's in the NY Times reveals that the appearance of the letter ...



Paul McAuley at Worldcon 2005 in Glasgow



Maui Veterans Highway shown just above Kealia Pond National Wildlife Refuge as it enters Kihei



Construction of restrooms and locker rooms with east side stands and pavilion



McLaren 600LT



Buzz Aldrin received praise for his performance.



James Barber House



A Fiat G.91PAN, in service in the Freccce Tricolori from 1963 to 1982



Alexa Stirling, c. 1919

Figure 21: Random training samples from WIT