

Large-scale Dataset Pruning with Dynamic Uncertainty

Muyang He^{1,2*} Shuo Yang³ Tiejun Huang^{1,2} Bo Zhao^{1†}

¹Beijing Academy of Artificial Intelligence

²Peking University ³University of Technology Sydney

isaache@pku.edu.cn, zhaobo@baai.ac.cn

Abstract

The state of the art of many learning tasks, e.g., image classification, is advanced by collecting larger datasets and then training larger models on them. As the outcome, the increasing computational cost is becoming unaffordable. In this paper, we investigate how to prune the large-scale datasets, and thus produce an informative subset for training sophisticated deep models with negligible performance drop. We propose a simple yet effective dataset pruning method by exploring both the prediction uncertainty and training dynamics. We study dataset pruning by measuring the variation of predictions during the whole training process on large-scale datasets, i.e., ImageNet-1K and ImageNet-21K, and advanced models, i.e., Swin Transformer and ConvNeXt. Extensive experimental results indicate that our method outperforms the state of the art and achieves 25% lossless pruning ratio on both ImageNet-1K and ImageNet-21K. The code and pruned datasets are available at <https://github.com/BAAI-DCAI/Dataset-Pruning>.

1. Introduction

The emerging large models, e.g., vision transformers [10, 12, 21] in computer vision, can significantly outperform the traditional neural networks, e.g., ResNet [16], when trained on large-scale datasets, e.g., ImageNet-21K [9] and JFT-300M [35]. However, storing large datasets and training on them are expensive and even unaffordable, which prevents individuals or institutes with limited resources from exceeding the state of the art. How to improve the data-efficiency of deep learning and achieve good performance with less training cost is a long-standing problem.

It is known that large-scale datasets have many redundant and less-informative samples which contribute little to model training. Dataset pruning (or coreset selection)

[5, 6, 13–15, 27, 31–33, 37, 41, 43, 49] aims to remove those less-informative training samples and remain the informative ones of original dataset, such that models trained on the remaining subset can achieve comparable performance. Nevertheless, how to find the informative subset is a challenging problem, especially for large-scale datasets, since sample selection is a combinatorial optimization problem.

Limitations of previous work. Existing dataset pruning works select important samples based on geometry/distribution [5, 31, 41, 49], prediction uncertainty [6, 14, 32], training error/gradient [27, 37], optimization [2, 18, 19, 43], etc. However, these methods have several major limitations: 1) Confusion between hard samples and mislabeled samples. Both of them lead to more prediction errors and larger loss. Exploiting the uncertainty or prediction error solely may obtain inferior subsets. 2) Ignoring training dynamics. Existing methods generally utilize a pretrained model to select samples. The geometry of data distribution thus does not reflect that throughout the training process. 3) Difficulty in scalability and generalization. Most of pruning methods focus on small-scale datasets, e.g., CIFAR-10/100 [20] which contains only 50K training images with 32×32 resolution, and traditional convolutional models, e.g., ResNet. They have difficulty in scaling up to large datasets with millions of samples due to the high computational complexity and memory cost.

Our method. In this paper, we design a simple yet effective sample selection metric dubbed Dynamic-Uncertainty (*Dyn-Unc*) for dataset pruning. We extract the prediction uncertainty by measuring the variation of predictions in a sliding window, and then introduce the training dynamics into it by averaging the variation throughout the whole training process. To address the first limitation, *Dyn-Unc* favors the uncertain samples rather than easy-to-learn ones or hard-to-learn ones during model training. *Dyn-Unc* can better separate informative samples from noise or mislabeled samples, as the model will gradually produce less uncertainty on mislabeled samples. For the second limitation, *Dyn-Unc* introduces the training dynamics and learns the property of data throughout the training process instead of

*Intern at Beijing Academy of Artificial Intelligence

†Corresponding author




Figure 1. Comparison to state-of-the-art dataset pruning methods on ImageNet-1K.

a certain stage. Since our method requires only one-run model training on the whole dataset before pruning it, it is scalable to large datasets.

As shown in Fig. 1 and Tab. 3, our *Dyn-Unc* significantly surpasses the previous state-of-the-art methods and achieve 25% lossless pruning ratio on ImageNet-1K and ImageNet-21K. Besides, experimental results (Tab. 4, Tab. 5) demonstrate that the coreset selected by *Dyn-Unc* using one architecture can generalize well to other unseen architectures.

2. Related work

2.1. Dataset Pruning

Dataset pruning is mainly rooted in coreset selection, which focuses on selecting a small but highly informative subset from a large dataset for training models. Dataset pruning methods typically calculate a scalar score for each training example based on predefined criteria. These can be grouped into geometry-based, uncertainty-based, error-based, bilevel-optimization-based, etc. For instance, (1) *Geometry*-based methods employ geometric measures in the feature space to select samples that preserve original geometry. A representative method is *Herdin* [5], which incrementally selects a sample that enables the new subset center to be closest to the whole-dataset center. *K-Center* [31, 40] method minimizes the largest distance from any data point to its nearest center in the selected subset, while *Moderate* [41] targets at data points with distances close to the median. (2) *Uncertainty*-based methods identify the most challenging samples, which are typically defined as instances that the model exhibits the least confidence in [6, 14, 32], or those that lie near the decision boundary where prediction variability is high [11, 24]. For example, [4] adjust the weights assigned to training samples by taking into account the variance of their predictive distribution.

(3) *Error*-based methods focus on the samples that contribute most to the loss. *GraNd* and *EL2N* [27] methods use the gradient norm and the L2-distance between the normalized error of predicted probabilities and one-hot labels to identify important samples. Another popular method, *Forgetting* [37], considers retaining the most forgettable samples. These are defined as those ones flipping most frequently from being correctly classified to being misclassified throughout training process. (4) There also exist some methods based on *bilevel optimization* [2, 18, 19], which optimize the subset selection (weights) that leads to best-performing models trained on the subset. However, these methods have difficulty in scaling up to large datasets due to the complex bilevel optimization.

While the above methods exhibit decent performance on small datasets and traditional (convolutional) models, we have empirically discovered that those studies based on small datasets have difficulty in generalizing to larger datasets and modern large models (e.g., Vision Transformers). Recognizing this limitation, our work takes an unprecedented step towards pruning highly extensive datasets, specifically ImageNet-1K and ImageNet-21K, and evaluating them by using state-of-the-art large models like Swin Transformer and ConvNeXt. This advancement marks a crucial milestone in the field, addressing the need for effective data pruning methods that can handle very large datasets and cutting-edge models.

2.2. Dataset Distillation

Another research topic that focuses on reducing the amount of training data is dataset distillation [39, 44, 48]. Instead of selecting a subset, these studies learn to *synthesize* a highly compact yet informative dataset that could be utilized for model training from scratch. The synthesis process involves aligning various factors such as performance [25, 34, 39],

gradient [45, 48], distribution [46, 47], feature map [38], and training trajectories [3, 8], between the model trained on the original dataset and the one trained on the synthetic data. Despite their promising prospects, these techniques, in their current state, cannot scale to large-scale datasets and models due to the computational load associated with the pixel-level optimization for image synthesis [8]. In addition, existing dataset distillation methods show marginal performance improvements over subset selection methods in the experiments of synthesizing a large number of training samples [47], since optimizing increasingly more variables (pixels) is challenging.

3. Method

3.1. Notation

Given a large-scale labeled dataset \mathcal{T} containing n training samples, denoted by $\mathcal{T} = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)\}$, where $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$ is the data point and $y \in \{0, \dots, C-1\}$ is the label belonging to C classes. The goal of dataset pruning is to remove the less-informative training samples and then output a subset $\mathcal{S} \subset \mathcal{T}$ with $|\mathcal{S}| < |\mathcal{T}|$, namely *coreset*, which can be used to train models and achieve comparable generalization performance to those trained on the whole dataset:

$$|\mathbb{E}_{\mathbf{x} \sim P_{\mathcal{D}}}[\ell(\phi_{\theta\mathcal{T}}(\mathbf{x}), y)] - \mathbb{E}_{\mathbf{x} \sim P_{\mathcal{D}}}[\ell(\phi_{\theta\mathcal{S}}(\mathbf{x}), y)]| < \epsilon. \quad (1)$$

$P_{\mathcal{D}}$ is the real data distribution. $\phi_{\theta\mathcal{T}}(\cdot)$ and $\phi_{\theta\mathcal{S}}(\cdot)$ are models parameterized with θ and trained on \mathcal{T} and \mathcal{S} respectively. ϵ is a small positive number. The pruning ratio is calculated as $r = 1 - \frac{|\mathcal{S}|}{|\mathcal{T}|}$.

3.2. Dataset Pruning

It is a common practice to classify training samples into three categories [4, 36]: easy, uncertain and hard samples. Previous works have shown that selecting easy, uncertain or hard samples may work well in different learning tasks. For example, *Herdling* [5] method tends to select easy samples that are close to class center. *GraNd* and *EL2N* [27] pick hard samples with large training gradients or errors. Different from them, [4, 36] measure the uncertainty of training samples and keep those with large prediction uncertainty. Thus, the best selection strategy is determined by multiple factors, including dataset size, data distribution, model architecture and training strategy.

Prediction Uncertainty. We design a simple yet effective large-scale dataset pruning method that selects samples by considering both prediction uncertainty and training dynamics. Larger dataset probably contains more easy samples that are replaceable for model training and hard/noisy samples that cannot be learned due to inconsistency between images and labels [29]. In other words, the model

trained on the rest can generalize well to those easy samples, meanwhile overfitting noisy samples does not improve model’s generalization performance. Hence, we choose to select samples with large prediction uncertainty and remove the rest.

Given a model θ^k trained on the whole dataset at k_{th} epoch, its prediction on the target label is denoted as $\mathbb{P}(y|\mathbf{x}, \theta^k)$. The uncertainty is defined as the standard deviation of the predictions in successive J training epochs with models $\{\theta^k, \theta^{k+1}, \dots, \theta^{k+J-1}\}$:

$$U_k(\mathbf{x}) = \sqrt{\frac{\sum_{j=0}^{J-1} [\mathbb{P}(y|\mathbf{x}, \theta^{k+j}) - \bar{\mathbb{P}}]^2}{J-1}}, \quad (2)$$

where $\bar{\mathbb{P}} = \frac{\sum_{j=0}^{J-1} \mathbb{P}(y|\mathbf{x}, \theta^{k+j})}{J}$.

J is the range for calculating the deviation.

Training Dynamics. Various uncertainty metrics have been designed in previous works, which are typically measured in a certain epoch, for example, in the last training epoch. However, the uncertainty metric may vary for models at different training epochs. Hence, we enhance the uncertainty metric with training dynamics. Specifically, we measure uncertainty in a sliding window with the length J , and then average the uncertainty throughout the whole training process:

$$U(\mathbf{x}) = \frac{\sum_{k=0}^{K-J-1} U_k(\mathbf{x})}{K-J}, \quad (3)$$

where K is the training epochs.

The training algorithm is depicted in Algorithm 1. Given a dataset \mathcal{T} , we train the model ϕ_{θ} to compute the dynamic uncertainty of each sample. In each training iteration, the prediction and loss are computed for every sample. Note that the computation can be in parallel with GPU. The loss averaged on a batch is used to update model parameters. Meanwhile, we compute the uncertainty of each sample based on its predictions over the past J epochs using Eq. (2). After K training epochs, the dynamic uncertainty is calculated using Eq. (3). Then, we sort all samples in the descending order of dynamic uncertainty $U(\cdot)$ and output the front $(1-r) \times |\mathcal{T}|$ samples as the pruned dataset \mathcal{S} , where r is the pruning ratio.

3.3. Analysis of Pruned Data

Following [36], we analyze the training data in terms of the prediction accuracy and variance. We train a Swin Transformer on ImageNet-1K and record the predictions of each training sample throughout all training epochs. Note that the prediction is a scalar (probability) that corresponds to the true label. Then, the mean and standard deviation of predictions over all training epochs are calculated for each

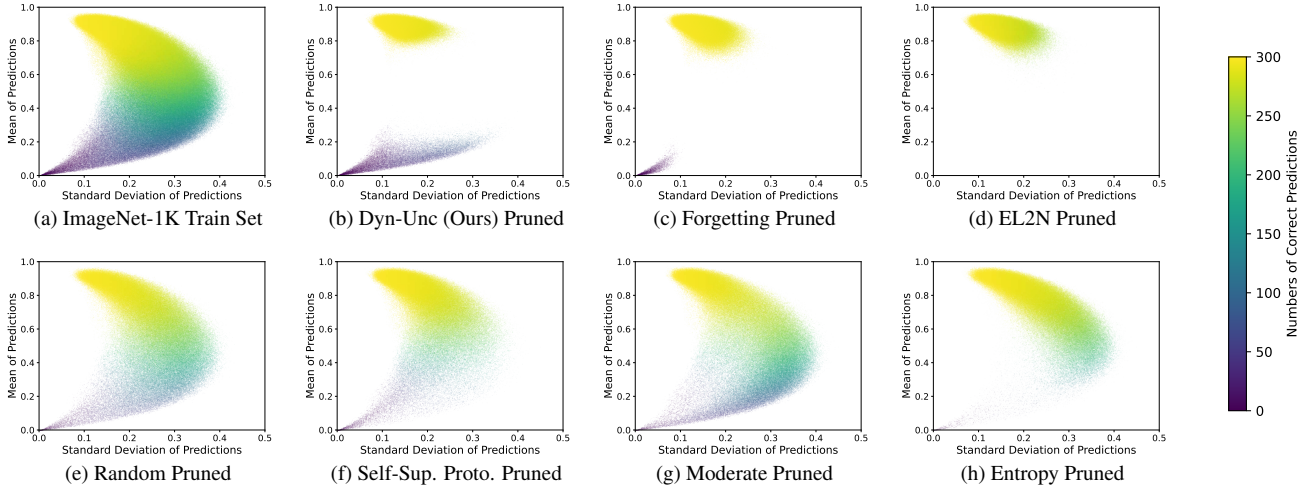


Figure 2. Visualization of data distribution and pruned data. We train a Swin Transformer on ImageNet-1K and record the predictions. The mean and standard deviation of predictions over all epochs are calculated for each sample, i.e., each point in the figure. The color represents the number of correct predictions in all epochs. (a) The distribution of the whole train set of ImageNet-1K. (b) - (h) The distribution of samples pruned by every method at the pruning ratio of 25%.

Algorithm 1: *Dyn-Unc.*

Input: Training set \mathcal{T} , pruning ratio r .

- 1 **Required:** Deep model ϕ_θ parameterized with θ , data augmentation \mathcal{A} , training epochs K , length of uncertainty window J , learning rate η .
- 2 **for** $k = 0, \dots, K - 1$ **do**
- 3 ▷ In parallel on GPUs.
- 4 Sample a batch $\mathcal{B} \sim \mathcal{T}$.
- 5 **for** $(x_i, y_i) \in \mathcal{B}$ **do**
- 6 Compute prediction $\mathbb{P}(y_i|x_i, \theta)$ and loss $\ell(\phi_\theta(\mathcal{A}(x_i)), y_i)$
- 7 **if** $k \geq J$ **then**
- 8 Compute uncertainty $U_{k-J}(x_i)$ using Eq. (2)
- 9 Update $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}$, where $\mathcal{L} = \frac{\sum \ell(\phi_\theta(\mathcal{A}(x_i)), y_i)}{|\mathcal{B}|}$
- 10 **for** $(x_i, y_i) \in \mathcal{T}$ **do**
- 11 Compute dynamic uncertainty $U(x_i)$ using Eq. (3)
- 12 Sort \mathcal{T} in the descending order of $U(\cdot)$
- 13 $\mathcal{S} \leftarrow$ front $(1 - r) \times |\mathcal{T}|$ samples in the sorted \mathcal{T}

Output: Pruned dataset \mathcal{S}

training sample, which compose the y-axis and x-axis respectively. The color reflects the number of correct predictions throughout the training. As shown in Fig. 2, the training set can be roughly split into three groups, namely 1) easy samples that have high mean accuracy and low deviation, 2) hard samples that have low mean accuracy and

low deviation, 3) uncertain samples, namely the rest ones. Previous works about sample selection [4, 6, 36], especially active learning, suggest uncertainty sampling for minimizing the uncertainty of model prediction. Most of these metrics are statistic, namely only one deviation/variance is calculated on a set of predictions/variables. Instead, the proposed metric is an average over a set of deviations throughout training process.

We visualize the data (25%) pruned by five state-of-the-art methods and ours in Fig. 2. Our method remove a large ratio of easy samples with high mean accuracy and low deviation and many hard samples with both low mean accuracy and low deviation. Note that our method is different from the naive pruning with statistic metric, for example using y-axis, namely, the mean accuracy. *Forgetting* method removes more easy samples while fewer hard samples, since some hard samples may cause continuous overturning, i.e., large forgetfulness, during model training. However, such samples may be less informative for model training and meanwhile slow convergence. Error-based method *EL2N* remove easy samples only. *Moderate* and *Self-Sup. Proto.* both remove many easy samples. For both *Forgetting* and ours, the uncertain samples are rarely removed, since the two methods both consider the training dynamics. Different from *Forgetting* and ours, *Self-Sup. Proto.*, *Moderate* and *Entropy* remove many uncertain samples, as they merely take the feature distribution or final prediction into consideration.

Fig. 3 demonstrates some hard samples that are kept or removed by our method. Obviously, these hard samples removed by our method are likely mislabeled. For example,

the “cowboy boot” is labeled as “buckle” but only a small region of the image contains a black buckle. The hard samples kept by our method are still recognizable though they contain some variances. For example, the “black swans” are with some “white swans”, and the “broom” image contains a man with a “broom”.

4. Experiments

4.1. Datasets and Settings

Datasets and Architectures. We evaluate our method on ImageNet-1K and ImageNet-21K [9] which have 1.3M and 14M training samples respectively. We implement main experiments with Swin Transformer [21], and also test our method on ConvNeXt [22] and ResNet [16]. We set the input size to be 224×224 .

Competitors. We compare to multiple representative and state-of-the-art dataset pruning methods, namely *Random*, *Forgetting* [37], *Entropy* [6], *EL2N* [27], *Self-Supervised Prototypes* [33] and *Moderate* [41]. Tab. 1 summarizes the characteristics of these methods. Note that there exist some other pruning methods. We do not compare with them due to their difficulty in scaling up to large datasets and models [43] or different experimental settings [28].

Implementation details. All experiments are conducted with NVIDIA A100 GPUs and PyTorch [26]. To accelerate data loading and pre-processing, we use NVIDIA Data Loading Library (DALI). For our method, we use $J = 10$ as the window’s length. All experiments are evaluated on the validation set of ImageNet-1K and ImageNet-Real [1]. ImageNet-Real reassesses the labels of ImageNet-1K-val and largely remedies the errors in the original labels, reinforcing a better benchmark. Top-1 accuracy is reported on both validation sets.

In the main experiments, we compare all seven methods with ImageNet-1K and Swin-T under the pruning ratios of 25%, 30%, 40% and 50%. For *EL2N*, the score is averaged over ten independent runs at epoch 30. For *Self-Sup. Proto.*, a Swin-B pre-trained and fine-tuned on ImageNet-1K with SimMIM [42] serves as the feature extractor. We train a Swin-T for implementing other dataset pruning methods. We employ an AdamW [23] optimizer and train models for 300 epochs with a cosine decay learning rate scheduler and 20 epochs of linear warm-up. The batch size is 1024. The initial learning rate is 0.001, and the weight decay is 0.05. We apply AutoAugment [7] implemented in NVIDIA DALI, label smoothing, random path drop and gradient clipping.

We also implement experiments on ImageNet-21K to indicate the good generalization ability and scalability of our method. We first pre-train Swin-T on ImageNet-21K (Fall 2011 release) for 90 epochs with 5 epochs of warm-up. Then we fine-tune the model on ImageNet-1K for 30 epochs

with 5 epochs of warm-up using a cosine decay learning rate scheduler, a batch size of 1024, an initial learning rate of 4×10^{-5} and a weight decay of 10^{-8} .

4.2. Pruning ImageNet-1K

Tab. 2 displays the performances of seven methods under different pruning rate. Obviously, our method *Dyn-Unc* achieves the best performances in all settings. Specifically, when pruning 25% training samples, our method obtains 79.54% and 84.88% top-1 accuracies on ImageNet-1K-val and ImageNet-1K-Real respectively, which are comparable to the upper-bound performances without pruning, namely 79.58% and 84.93%. These results indicate that our method can achieve 25% lossless dataset pruning ratio.

We observe consistent rankings of methods whether on the original validation set or Real. The performance gaps between our method and others are remarkable. The gap increases for more challenging settings in which more training data are removed. For instance, our method outperforms *Forgetting* by 0.84%, 0.87%, 0.94% and 3.32% when pruning 25%, 30%, 40% and 50% samples, considering that the standard deviation is around 0.2%.

Forgetting is always the runner-up when pruning less than 40% training data. This phenomenon emphasizes the priority of prediction uncertainty and training dynamics based methods. However, *Forgetting* suffers from a large performance drop from 40% to 50% sample pruning. We compare the pruned samples of ours and *Forgetting* in Fig. 4. The possible reason is that *Forgetting* throws more easy samples, which makes the model training unstable.

Generally, feature distribution based methods *Self-Sup. Proto.* and *Moderate* do not work well. *Random* baseline exceeds all other methods except *Moderate* and our *Dyn-Unc* on both two validation sets when pruning 50% training data. The average ranking shows that *Entropy* performs the worst in the ImageNet-1K pruning task. The possible reason may be that this method only considers the output of a trained model. We also illustrate these results in Fig. 1 for easier visual comparison.

Besides, we also test our method on ConvNeXt. With 25% training data pruned, *Dyn-Unc* achieves 78.77% and 84.13% accuracies on ImageNet-1K-val and ImageNet-1K-Real respectively and the results of no pruning are 78.79% and 84.05%, suggesting that our method can generalize to different models. We further study the cross-architecture generalization ability of our method in Sec. 4.4.

4.3. Pruning ImageNet-21K

In this paper, we for the first time test dataset pruning algorithms on ImageNet-21K which consists of 14 million images. The expensive computational cost prevents previous works from studying this dataset. As shown in Tab. 3, we prune ImageNet-21K with different ratios and then test on



Figure 3. Illustrations of hard samples kept and removed by our method at the pruning ratio of 25%. The removed ones are likely mislabeled or confusing, while the kept ones are certainly recognizable.

Table 1. Comparison to other dataset pruning methods.

Method	Prediction Uncertainty	Training Dynamics	Feature Distribution	Training Error	Label Supervised	Class Balance
Random Forgetting	✓	✓			✓	
Entropy	✓				✓	
EL2N				✓	✓	
Self-Sup.Proto. Moderate			✓		✓	✓
Dyn-Unc (Ours)	✓	✓	✓		✓	

ImageNet-1K-val and ImageNet-1K-Real. When pruning 25% training data, the Swin-T trained on our pruned dataset can achieve 82.14% and 87.31% accuracies on ImageNet-1K-val and ImageNet-1K-Real respectively. The former has 0.13% performance drop while the latter has 0.07% improvement. Considering the deviation, the results indicate that our method can achieve 25% lossless pruning ratio on ImageNet-21K.

4.4. Cross-architecture Generalization

Few works test the cross-architecture generalization performance of pruned dataset, although it is important for real-world applications in which the new architectures will be trained. We conduct experiments to verify the pruned dataset can generalize well to those unseen architectures that are inaccessible during dataset pruning. The representative architectures, Swin-T, ConvNeXt-T and ResNet-50, are taken into consideration here. Dataset pruned by one architecture is used to train others. We do experiments at the pruning ratio of 25%. We train ConvNeXt-T with the batch size of 2048, initial learning rate of 0.008. The rest

hyper-parameters are the same as training Swin-T. For training ResNet-50, we employ the SGD optimizer for 90 epochs with step decay, batch size of 2048, initial learning rate of 0.8 and weight decay of 10^{-4} .

The results in Tab. 4 show that the dataset pruned by our method has good generalization ability to unseen architectures. For example, pruning ImageNet-1K with Swin-T and then training ConvNeXt-T and ResNet-50 causes only 0.13% and 0.4% performance drop on ImageNet-1K-val compared to training the two architectures on the whole dataset respectively. Notably, pruning with Swin-T or ResNet-50 leads to performance improvement of training ConvNeXt-T over the whole dataset training when trained models are validated on ImageNet-1K-Real. It is reasonable because dataset pruning with proper ratio can remove noisy and less-informative data, thus improve dataset quality.

We further compare to *Forgetting* and *Moderate* in cross-architecture experiment of pruning ImageNet-1K with Swin-T and then testing with other architectures. As shown in Tab. 5, our method performs slightly worse than the

Table 2. Experimental results on ImageNet-1K. Our method achieves the best performances in all settings. The results indicate that our method can achieve 25% lossless dataset pruning ratio.

ImageNet-1K-val Acc. (%)									
Method\Pruning Ratio	25%	rank	30%	rank	40%	rank	50%	rank	avg rank
Random	77.82	5	77.18	5	75.93	4	74.54	3	4.25
Forgetting	78.70	2	78.27	2	77.55	2	74.32	4	2.5
Entropy	77.11	7	76.62	7	74.87	7	72.65	6	6.75
EL2N	78.51	3	78.15	3	75.87	5	71.10	7	4.5
Self-Sup. Proto.	78.24	4	77.23	4	75.69	6	73.48	5	4.75
Moderate	77.74	6	77.06	6	75.94	3	74.98	2	4.25
Dyn-Unc (ours)	79.54 ± 0.13	1	79.14 ± 0.07	1	78.49 ± 0.22	1	77.64 ± 0.17	1	1
No Pruning	79.58 ± 0.15								

ImageNet-1K-ReaL Acc. (%)									
Method\Pruning Ratio	25%	rank	30%	rank	40%	rank	50%	rank	avg rank
Random	83.55	5	82.96	5	81.90	4	80.82	3	4.25
Forgetting	84.20	2	83.75	2	82.94	2	79.62	4	2.5
Entropy	82.77	7	82.45	7	80.85	7	78.96	6	6.75
EL2N	83.84	3	83.40	3	81.12	6	75.93	7	4.75
Self-Sup. Proto.	83.71	4	82.74	6	81.33	5	79.28	5	5
Moderate	83.42	6	83.00	4	81.95	3	81.05	2	3.75
Dyn-Unc (ours)	84.88 ± 0.15	1	84.62 ± 0.05	1	84.13 ± 0.26	1	83.22 ± 0.19	1	1
No Pruning	84.93 ± 0.11								




Figure 4. Comparison of samples pruned by ours and *Forgetting* at the ratio of 50%. *Forgetting* method throws too many easy samples.

Table 3. Experimental results on ImageNet-21K.

Pruning Ratio	Accuracy (%)	
	IN-1K-val	IN-1K-ReaL
0%	82.27	87.24
25%	82.14	87.31
30%	81.96	87.16
40%	81.68	87.15
50%	81.26	86.84

upper-bound of no pruning, but remarkably exceeds other methods. Especially, our method surpasses *Forgetting* and

Moderate by 0.95% and 1.68% when testing with ResNet-50 on ImageNet-1K-val. The promising results indicate that our method has good generalization ability on unseen architectures and the pruned dataset can be used in real-world applications in which the architectures of downstream tasks are unknown.

4.5. Out-of-distribution Detection

To further verify the robustness and reliability of our method, we evaluate our method on ImageNet-O [17], an out-of-distribution detection dataset designed for evaluating models trained on ImageNet. ImageNet-O consists of images from classes that are not found in ImageNet-1K but in

Table 4. Cross-arch. generalization performance of Dyn-Unc.

Accuracy (%)		
Pruning → Test Model	IN-1K-val	IN-1K-ReaL
ConvNeXt → Swin	78.82	84.05
ResNet → Swin	79.03	84.78
Swin No Pruning	79.58	84.93
Swin → ConvNeXt	78.66	84.15
ResNet → ConvNeXt	78.68	84.26
ConvNeXt No Pruning	78.79	84.05
Swin → ResNet	75.51	82.17
ConvNeXt → ResNet	75.39	82.10
ResNet No Pruning	75.91	82.64

Table 5. Comparison to others w.r.t. cross-arch. generalization.

Swin → ConvNeXt Acc. (%)		
Method\Val. Set	IN-1K-val	IN-1K-ReaL
Forgetting	78.50	83.84
Moderate	77.25	82.91
Dyn-Unc (ours)	78.66	84.15
No Pruning	78.79	84.05
Swin → ResNet Acc. (%)		
Method\Val. Set	IN-1K-val	IN-1K-ReaL
Forgetting	74.56	81.52
Moderate	73.83	80.91
Dyn-Unc (ours)	75.51	82.17
No Pruning	75.91	82.64

ImageNet-21K. ImageNet-O collects images that are likely to be wrongly classified with high confidence by models trained on ImageNet-1K. Hence, this dataset can be used to test the robustness and reliability of models. We use AUPR (area under the precision-recall curve) metric [17, 30], and higher AUPR means better performing OOD detector. For ImageNet-O, random chance level for the AUPR is approximately 16.67%, and the maximum is 100% [17].

Based on models trained in the aforementioned ImageNet-1K experiments, we report the OOD detection results in Tab. 6. Interestingly, our dataset pruning method can improve the OOD detection performance slightly, from 21.98% (no pruning) to 22.61% (pruning 30% samples). The possible reason is that dataset pruning prevents models from overfitting many easy and noisy samples, and thus improves model’s generalization ability.

4.6. Ablation Study of Observation Number

[36] conducts experiments on NLP tasks and shows that training on ambiguous data can achieve quite good in-distribution and out-of-distribution performance. We apply

Table 6. Out-of-distribution detection performance of Dyn-Unc.

Swin-T	
Pruning Ratio	ImageNet-O AUPR (%)
0%	21.98 ± 0.14
25%	22.51 ± 0.38
30%	22.61 ± 0.27
40%	21.97 ± 0.44
50%	21.85 ± 0.40

Table 7. Ablation study on observation number at 25% pruning ratio.

Accuracy (%)		
Method (observation numbers)	IN-1K-val	IN-1K-ReaL
[36] (only 1)	75.05	81.60
Dyn-Unc (1×, Ours)	79.54	84.88
Multi. Init. (3×)	79.19	84.75
Multi. Arch. (3×)	79.34	84.60

it to image classification task. Different from our method, [36] takes first five or six epochs into account and observe only one variance of these epochs. We compare to this method in Tab. 7. The significant performance gap implies that measuring one variance is not enough and training dynamics is important for sample selection.

Our *Dyn Unc* observes $K - J$ variances throughout the K -epoch training with a sliding window of length J . We try to figure out whether more observations from multiple model training processes with different initializations and architectures (Swin-T, ConvNeXt and ResNet) can produce better metric and therefore better performance. In Tab. 7, the results show that using more observations does not improve the performance.

5. Conclusion

In this paper, we push the study of dataset pruning to large datasets, i.e., ImageNet-1K/21K and advanced models, i.e., Swin Transformer and ConvNeXt. A simple yet effective dataset pruning method is proposed based on the prediction uncertainty and training dynamics. The extensive experiments verify that our method achieves the best results in all settings comparing to the state of the art. Notably, our method achieves 25% lossless pruning ratio on both ImageNet-1K and ImageNet-21K. The cross-architecture generalization and out-of-distribution detection experiments show promising results that pave the way for real-world applications.

Acknowledgement. This work is funded by National Science and Technology Major Project (2021ZD0111102) and NSFC-62306046.

References

- [1] Lucas Beyer, Olivier J Hénaff, Alexander Kolesnikov, Xiao-hua Zhai, and Aaron van den Oord. Are we done with imagenet? *arXiv preprint arXiv:2006.07159*, 2020. [5](#)
- [2] Zalán Borsos, Mojmír Mutný, and Andreas Krause. Coresets via bilevel optimization for continual learning and streaming. *Advances in Neural Information Processing Systems*, 33:14879–14890, 2020. [1](#), [2](#)
- [3] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A. Efros, and Jun-Yan Zhu. Dataset distillation by matching training trajectories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. [3](#)
- [4] Haw-Shiuan Chang, Erik Learned-Miller, and Andrew McCallum. Active bias: Training more accurate neural networks by emphasizing high variance samples. *Advances in Neural Information Processing Systems*, 30, 2017. [2](#), [3](#), [4](#)
- [5] Yutian Chen, Max Welling, and Alex Smola. Super-samples from kernel herding. *The Twenty-Sixth Conference Annual Conference on Uncertainty in Artificial Intelligence*, 2010. [1](#), [2](#), [3](#)
- [6] Cody Coleman, Christopher Yeh, Stephen Mussmann, Baharan Mirzasoleiman, Peter Bailis, Percy Liang, Jure Leskovec, and Matei Zaharia. Selection via proxy: Efficient data selection for deep learning. In *International Conference on Learning Representations*, 2020. [1](#), [2](#), [4](#), [5](#)
- [7] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. [5](#)
- [8] Justin Cui, Ruochen Wang, Si Si, and Cho-Jui Hsieh. Scaling up dataset distillation to imagenet-1k with constant memory. In *International Conference on Machine Learning*, pages 6565–6590. PMLR, 2023. [3](#)
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [1](#), [5](#)
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [1](#)
- [11] Melanie Ducoffe and Frederic Precioso. Adversarial active learning for deep networks: a margin based approach. *arXiv preprint arXiv:1802.09841*, 2018. [2](#)
- [12] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19358–19369, 2023. [1](#)
- [13] Dan Feldman. Introduction to core-sets: an updated survey. *arXiv preprint arXiv:2011.09384*, 2020. [1](#)
- [14] Yariv Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International conference on machine learning*, pages 1183–1192. PMLR, 2017. [1](#), [2](#)
- [15] Chengcheng Guo, Bo Zhao, and Yanbing Bai. Deepcore: A comprehensive library for coreset selection in deep learning. In *Database and Expert Systems Applications: 33rd International Conference, DEXA 2022, Vienna, Austria, August 22–24, 2022, Proceedings, Part I*, pages 181–195. Springer, 2022. [1](#)
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [1](#), [5](#)
- [17] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021. [7](#), [8](#)
- [18] Krishnateja Killamsetty, Durga Sivasubramanian, Ganesh Ramakrishnan, and Rishabh Iyer. Glisten: Generalization based data subset selection for efficient and robust learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8110–8118, 2021. [1](#), [2](#)
- [19] Krishnateja Killamsetty, Xujiang Zhao, Feng Chen, and Rishabh Iyer. Retrieve: Coreset selection for efficient and robust semi-supervised learning. *Advances in Neural Information Processing Systems*, 34:14488–14501, 2021. [1](#), [2](#)
- [20] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Cite-seer, 2009. [1](#)
- [21] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. [1](#), [5](#)
- [22] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. [5](#)
- [23] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. [5](#)
- [24] Katerina Margatina, Giorgos Vernikos, Loïc Barrault, and Nikolaos Aletras. Active learning by acquiring contrastive examples. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 650–663, 2021. [2](#)
- [25] Timothy Nguyen, Zhouong Chen, and Jaehoon Lee. Dataset meta-learning from kernel-ridge regression. In *International Conference on Learning Representations*, 2021. [2](#)
- [26] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, 2019. [5](#)

- [27] Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. Deep learning on a data diet: Finding important examples early in training. *Advances in Neural Information Processing Systems*, 34, 2021. 1, 2, 3, 5
- [28] Ziheng Qin, Kai Wang, Zangwei Zheng, Jianyang Gu, Xiangyu Peng, Xu Zhao Pan, Daquan Zhou, Lei Shang, Baigui Sun, Xuansong Xie, and Yang You. Infobatch: Lossless training speed up by unbiased dynamic data pruning. In *The Twelfth International Conference on Learning Representations*, 2024. 5
- [29] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. 3
- [30] Takaya Saito and Marc Rehmsmeier. The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets. *PloS one*, 10(3): e0118432, 2015. 8
- [31] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018. 1, 2
- [32] Yanyao Shen, Hyokun Yun, Zachary C. Lipton, Yakov Kroprod, and Animashree Anandkumar. Deep active learning for named entity recognition. In *International Conference on Learning Representations*, 2018. 1, 2
- [33] Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari Morcos. Beyond neural scaling laws: beating power law scaling via data pruning. *Advances in Neural Information Processing Systems*, 35:19523–19536, 2022. 1, 5
- [34] Felipe Petroski Such, Aditya Rawal, Joel Lehman, Kenneth O Stanley, and Jeff Clune. Generative teaching networks: Accelerating neural architecture search by learning to generate synthetic training data. *International Conference on Machine Learning (ICML)*, 2020. 2
- [35] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*, pages 843–852, 2017. 1
- [36] Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A Smith, and Yejin Choi. Dataset cartography: Mapping and diagnosing datasets with training dynamics. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293, 2020. 3, 4, 8
- [37] Mariya Toneva, Alessandro Sordani, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J. Gordon. An empirical study of example forgetting during deep neural network learning. In *International Conference on Learning Representations*, 2019. 1, 2, 5
- [38] Kai Wang, Bo Zhao, Xiangyu Peng, Zheng Zhu, Shuo Yang, Shuo Wang, Guan Huang, Hakan Bilen, Xinchao Wang, and Yang You. Cafe: Learning to condense dataset by aligning features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12196–12205, 2022. 3
- [39] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A Efros. Dataset distillation. *arXiv preprint arXiv:1811.10959*, 2018. 2
- [40] Gert W Wolf. Facility location: concepts, models, algorithms and case studies. series: Contributions to management science, 2011. 2
- [41] Xiaobo Xia, Jiale Liu, Jun Yu, Xu Shen, Bo Han, and Tongliang Liu. Moderate coreset: A universal method of data selection for real-world data-efficient deep learning. In *The Eleventh International Conference on Learning Representations*, 2023. 1, 2, 5
- [42] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022. 5
- [43] Shuo Yang, Zeke Xie, Hanyu Peng, Min Xu, Mingming Sun, and Ping Li. Dataset pruning: Reducing training data by examining generalization influence. In *The Eleventh International Conference on Learning Representations*, 2023. 1, 5
- [44] Bo Zhao. Data-efficient neural network training with dataset condensation. 2023. 2
- [45] Bo Zhao and Hakan Bilen. Dataset condensation with differentiable siamese augmentation. In *International Conference on Machine Learning*, pages 12674–12685. PMLR, 2021. 3
- [46] Bo Zhao and Hakan Bilen. Synthesizing informative training samples with GAN. In *NeurIPS 2022 Workshop on Synthetic Data for Empowering ML Research*, 2022. 3
- [47] Bo Zhao and Hakan Bilen. Dataset condensation with distribution matching. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6514–6523, 2023. 3
- [48] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. Dataset condensation with gradient matching. In *International Conference on Learning Representations*, 2021. 2, 3
- [49] Haizhong Zheng, Rui Liu, Fan Lai, and Atul Prakash. Coverage-centric coreset selection for high pruning rates. In *The Eleventh International Conference on Learning Representations*, 2023. 1