

How Much More Data Do I Need?

Estimating Requirements for Downstream Tasks

Rafid Mahmood¹ James Lucas^{1,2,3} David Acuna^{1,2,3} Daiqing Li¹ Jonah Philion^{1,2,3}
 Jose M. Alvarez¹ Zhiding Yu¹ Sanja Fidler^{1,2,3} Marc T. Law¹
¹NVIDIA ²University of Toronto ³Vector Institute

{rmahmood, jalucas, dacunamarrer, daiqingl, jphilion, josea, zhidingy, sfidler, marcl}@nvidia.com

Abstract

Given a small training data set and a learning algorithm, how much more data is necessary to reach a target validation or test performance? This question is of critical importance in applications such as autonomous driving or medical imaging where collecting data is expensive and time-consuming. Overestimating or underestimating data requirements incurs substantial costs that could be avoided with an adequate budget. Prior work on neural scaling laws suggest that the power-law function can fit the validation performance curve and extrapolate it to larger data set sizes. We find that this does not immediately translate to the more difficult downstream task of estimating the required data set size to meet a target performance. In this work, we consider a broad class of computer vision tasks and systematically investigate a family of functions that generalize the power-law function to allow for better estimation of data requirements. Finally, we show that incorporating a tuned correction factor and collecting over multiple rounds significantly improves the performance of the data estimators. Using our guidelines, practitioners can accurately estimate data requirements of machine learning systems to gain savings in both development time and data acquisition costs.

1. Introduction

Before deploying a deep learning model, designers may mandate that the model meet a baseline performance, such as a target metric over a held out validation or test set. For example, an object detector may require a minimum mean average precision before being deployed in a safety-critical application. One of the most effective ways of meeting the target performance is by collecting more training data for a given model. However, how much more data is needed?

Overestimating data requirements can incur costs from unnecessary collection, cleaning, and annotation. For example, annotating segmentation data sets may require 15 to 40 seconds per object [2], meaning annotating a driving




Figure 1. Extrapolating accuracy on ImageNet [7] as a function of data set size from 10% of the data set (125,000 images; dotted) and 50% (600,000 images; dashed) using four regression functions. The vertical dashed lines show how much data is needed to meet a target 67% validation accuracy according to each dashed curve. All the dashed curves can accurately extrapolate performance as they are given a sufficient amount of images. Although the functions have an error of 1-6% from the ground truth (67% at 900,000 images), they mis-estimate the data requirement by 120,000 to 310,000 images.

data set of 100,000 images with on average 10 cars per image can take between 170 and 460 days-equivalent of time. On the other hand, underestimating means having to collect more data at a later stage, incurring future costs and workflow delays. For instance in autonomous vehicle applications, each period of data collection requires managing a fleet of drivers to record driving videos. Thus, accurately estimating how much data is needed for a given task can reduce both costs and delays in the deep learning workflow.

There is a growing body of literature on estimating the sample complexity of machine learning models [4, 11, 12]. Recently proposed *neural scaling laws* suggest that generalization scales with the data set size according to a power

law [3, 14, 15, 20, 26]. Rosenfeld *et al.* [26] propose fitting a power law function using the performance statistics from a small data set to extrapolate the performance for larger data sets; while not a focus of their paper, they suggest this can be used to estimate the data requirements. However, the power law function is not the only possible choice. We propose in this paper to use it with similar functions that can be more accurate in practice. Figure 1 illustrates the data collection process in image classification with the ImageNet data set [7] for the power law function and several effective alternatives. When using small data sets to extrapolate, the fitted functions may diverge in different ways from the ground truth performance curve. More importantly, even a small error in extrapolating accuracy can lead to large errors in over or under-estimating the data requirements, which may present huge operational costs.

In this paper, we ask: given a small training data set and a model not yet meeting target performance in some metric, what is the least amount of data we should collect to meet the target? Generalizing the estimation of data requirements from power laws, we investigate several alternate regression functions and show that all of them are well-suited towards estimating model performance. Moreover, each function is almost always either overly optimistic (*i.e.* under-estimating the data requirement) or pessimistic (*i.e.* over-estimating), meaning that there is no unique best regression function for all situations, but using all of the different functions, we can approximately bound the true data requirement. Through a simulation of the data collection workflow, we show that incrementally collecting data over multiple rounds is critical to meeting the requirement without significantly exceeding it. Finally, we introduce a simple correction factor to help these functions meet data requirement more often; this factor can be learned by simulating on prior tasks. We explore classification, detection, and segmentation tasks with different data sets, models, and metrics to show that our results hold in every setting considered.

Altogether, our empirical findings and proposed enhancements yield easy-to-implement guidelines for data collection in real-world applications: practitioners should allocate for up to five rounds of data collection and use the correction factor introduced in this paper to augment an optimistic regression function (*e.g.* Power Law, Logarithmic, Algebraic Root) in order to accurately estimate data requirements and ultimately collect only a relatively small amount more than the minimum data required to meet the desired performance. We believe that this approach can improve workflows and yield large cost savings in the future.

2. Related work

Neural scaling laws. Prior work has estimated model performance as a function of data set size [1, 4, 11, 26, 30, 35]. The recent literature on neural scaling laws argues that

model performance (usually defined as decreasing loss over a validation set) scales with data set size according to a power law function, *i.e.* $V \propto \theta_1 n^{\theta_2}$ where n is the data set size. Hestness *et al.* [14] empirically validate that power laws accurately model negative validation set loss and top-1 error over different image classification, language, and audio tasks. Bahri *et al.* [3] prove that for over-parametrized networks, under Lipschitz continuity of the loss function, model, and data distribution, the out-of-sample loss scales in $\mathcal{O}(n^{-1/\theta})$. Rosenfeld *et al.* [26] fit power law functions using small data subsets. Finally, Hoiem *et al.* [15] use power laws to construct learning curves and investigate modeling questions. One key difference between these studies and our own is that we focus on estimating target data requirements given an approximate relationship between data size and model performance; such as a power law function. More broadly, this area also relates to the study of learning curves in classical machine learning [10, 19, 31]. Our work differentiates from this literature through a detailed simulation that investigates the operational costs of poorly estimating a learning curve.

Active learning. In this work, we consider collecting data over multiple rounds. This is related to active learning [6], where a model selects which data to use during multiple rounds of training. The focus of active learning is to intelligently select this data given a fixed collection budget [23, 27–29, 33], sometimes with a focus on the performance on rare categories [25]. However, the goal of this work is to predict the optimal collection budget itself. This paper focuses on random sampling, but includes experiments with active learning in the Appendix to demonstrate that our insights on estimating the data requirement hold independent of the sampling strategy.

Statistical learning theory. Loosely speaking, statistical learning theory seeks to relate model performance and data set size. Accurate theoretical characterizations of this relationship could be used to infer the target data requirements, but these results are typically only tight asymptotically; if at all. More recent work has explored empirically estimating this theoretical relationship [17, 18]. Bisla *et al.* [4] build models of generalization for deep neural networks under assumptions on the training and test behaviour that are validated empirically. Bisla *et al.* highlight the utility in being able to estimate data requirements from such a model, but do not explore this empirically as we do in this work.

3. Main problem

In this section, we mathematically define the data collection problem and the general solution method. The goal of this problem is to estimate the data set size that returns a desired performance in a limited number of rounds. We first model performance as a function of data set size and then solve for the data given an input performance.

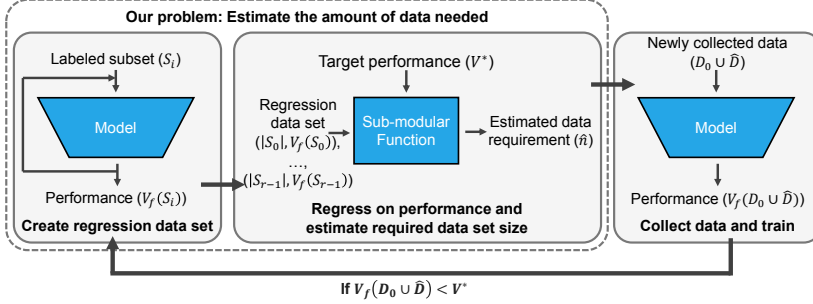


Figure 2. The iterative data collection loop.

3.1. The data collection problem

Let $z \sim p(z)$ be data drawn from a distribution p . For instance, $z := (x, y)$ may correspond to images x and labels y . Consider a prediction problem for which we currently have an initial training data set $\mathcal{D}_0 := \{z_i\}_{i=1}^{n_0}$ of n_0 points and a model f . Let $V_f(\mathcal{D})$ be a score function of the model after it is trained on a set \mathcal{D} . Our goal is to obtain a pre-determined target score $V^* > V_f(\mathcal{D}_0)$.

To achieve our goal, we sample \hat{n} additional points to create $\hat{\mathcal{D}} := \{\hat{z}_i\}_{i=1}^{\hat{n}} \sim p(z)$ and then evaluate $V_f(\mathcal{D}_0 \cup \hat{\mathcal{D}})$. If we do not meet the target, we must determine a larger \hat{n} and augment $\hat{\mathcal{D}}$ with more data. Because each data point incurs a cost from collecting, cleaning, and labeling, we ideally want the fewest number of points \hat{n} that achieve the target. Furthermore, because initiating a round of data collection is itself expensive and time-consuming, we are only permitted a maximum of T rounds; failing to meet the requirement within T rounds means failing to solve the problem. This problem is summarized in the following iterative sequence. Initialize $\hat{\mathcal{D}} = \emptyset$. Then in each round, repeat:

- (1) Estimate the amount of additional data \hat{n} needed.
- (2) Sample points until $|\hat{\mathcal{D}}| = \hat{n}$ and then evaluate the score. If $V_f(\mathcal{D}_0 \cup \hat{\mathcal{D}}) \geq V^*$, then terminate. Otherwise, repeat for another round up to T rounds.

The objective of the data collection problem is to select the minimum \hat{n} such that $V_f(\mathcal{D}_0 \cup \hat{\mathcal{D}}) \geq V^*$ within T rounds. This paper focuses on the first step of the loop: accurately estimating the \hat{n} required to meet V^* .

3.2. Regressing performance using data set size

Figure 2 illustrates our data collection pipeline to estimate \hat{n} , motivated by the following empirical observation.

Observation from [11, 26]. Let $\mathcal{D}_0 \subset \mathcal{D}_1 \subset \dots$ be a growing sequence of data sets and let $n_i = |\mathcal{D}_i|$ for each i in the sequence. Then, the piecewise linear function

$$v(n) := \begin{cases} \frac{V_f(\mathcal{D}_0)}{n_0} n, & n \leq n_0 \\ \frac{V_f(\mathcal{D}_i) - V_f(\mathcal{D}_{i-1})}{n_i - n_{i-1}} (n - n_{i-1}) + V_f(\mathcal{D}_{i-1}), & n_{i-1} \leq n \leq n_i \end{cases}$$

is concave and monotonically increasing.

Recall that $V_f(\mathcal{D}_i)$ is the model score after it is trained on \mathcal{D}_i . We refer to $v(n)$ as the model score function over

Regression Function	$\hat{v}(n; \theta)$
Power Law	$\theta_1 n^{\theta_2} + \theta_3$
Arctan	$\frac{200}{\pi} \arctan\left(\theta_1 \frac{\pi}{2} n + \theta_2\right) + \theta_3$
Logarithmic	$\theta_1 \log(n + \theta_2) + \theta_3$
Algebraic Root	$\frac{100n}{(1 + \theta_1 n ^{\theta_2})^{1/\theta_2}} + \theta_3$

Table 1. Four concave monotonic increasing regression functions explored in this paper. The set of learnable parameters is $\theta := \{\theta_1, \theta_2, \theta_3\}$.

Algorithm 1 The data collection problem

- 1: **Input:** Initial data set \mathcal{D}_0 , Score function $V_f(\mathcal{D})$, Target score V^* , Maximum rounds T , Regression model $\hat{v}(n; \theta)$, Initial regression set size r
- 2: Set $n_0 \leftarrow |\mathcal{D}_0|$, $\hat{\mathcal{D}} = \emptyset$
- 3: CREATE REGRESSION DATA SET
- 4: Sample subsets $\mathcal{S}_0 \subset \mathcal{S}_1 \subset \dots \subset \mathcal{S}_{r-1} = \mathcal{D}_0$
- 5: Evaluate $V_f(\mathcal{S}_i)$ and create $\mathcal{R} \leftarrow \{(|\mathcal{S}_i|, V_f(\mathcal{S}_i))\}_{i=0}^{r-1}$
- 6: PERFORM DATA COLLECTION
- 7: **repeat**
- 8: Fit $\theta^* \leftarrow \arg \min_{\theta} \sum_{(n,v) \in \mathcal{R}} (v - \hat{v}(n; \theta))^2$
- 9: Minimize \hat{n} subject to $\hat{v}(n_0 + \hat{n}; \theta^*) \geq V^*$
- 10: Sample points from $p(z)$ until $|\hat{\mathcal{D}}| = \hat{n}$
- 11: Train model and evaluate score $V_f(\mathcal{D}_0 \cup \hat{\mathcal{D}})$
- 12: Update $\mathcal{R} \leftarrow \mathcal{R} \cup \{(n_0 + \hat{n}, V_f(\mathcal{D}_0 \cup \hat{\mathcal{D}}))\}$
- 13: **until** $V_f(\mathcal{D}_0 \cup \hat{\mathcal{D}}) \geq V^*$ or T rounds have passed
- 14: **Output:** Final collected data set $\mathcal{D}_0 \cup \hat{\mathcal{D}}$

the training data set size. The observation implies that intuitively, as we collect more data, the marginal value of each additional data point should decrease (e.g. Figure 1). Furthermore, we can model $v(n)$ by regression using concave, monotonically increasing functions. Within the data collection loop, we first estimate \hat{n} by using the available data, \mathcal{D}_0 and $\hat{\mathcal{D}}$, and the corresponding scores by fitting a regression model $\hat{v}(n; \theta)$ of $v(n)$, where θ is the set of regression parameters. We consider four functions that satisfy the Observation (see Table 1) from the learning curve literature [31]. While we could use more complicated models, we find these simpler structured functions with a small number of parameters are easier to fit to smaller data sets of learning statistics. Using the fitted regression function, we solve for the smallest \hat{n} such that $\hat{v}(n_0 + \hat{n}; \theta) \geq V^*$.

Algorithm 1 summarizes the main steps. We first create a regression data set by selecting r subsets $\mathcal{S}_0 \subset \mathcal{S}_1 \subset \dots \subset \mathcal{S}_{r-1} = \mathcal{D}_0$ and computing their scores; this yields a set of r pairs $\mathcal{R} := \{(|\mathcal{S}_i|, V_f(\mathcal{S}_i))\}_{i=0}^{r-1}$. Then, in the data collection loop, we select a function $\hat{v}(n; \theta)$ from Table 1 and fit the set of parameters θ via least squares minimization. Finally, we minimize \hat{n} subject to $\hat{v}(\hat{n}; \theta^*) \geq V^*$, and then collect \hat{n} new points. In subsequent rounds of data collection as we obtain $\hat{\mathcal{D}}$, we augment \mathcal{R} with pairs $(|\mathcal{D}_0| + |\hat{\mathcal{D}}|, V_f(\mathcal{D}_0 \cup \hat{\mathcal{D}}))$, and then re-fit $\hat{v}(n; \theta)$.

Data set	Task	Score	Full data set size
CIFAR10 [21]	Classification	Accuracy	50,000
CIFAR100 [21]	Classification	Accuracy	50,000
ImageNet [7]	Classification	Accuracy	1,281,167
VOC [8,9]	2-D Object Detection	Mean AP	16,551
nuScenes [5]	3-D Object Detection	Mean AP	28,130
BDD100K [34]	Semantic Segmentation	Mean IoU	7,000
nuScenes [5]	BEV Segmentation	Mean IoU	28,130

Table 2. Data sets, tasks, and score functions considered.

The existing literature shows that power laws can estimate model accuracy using data set size, but the practical application of estimating the required data set size to meet a target score presents three major challenges. We highlight them below using the ImageNet data set in Figure 1.

All of the functions in Table 1 fit the model score. With enough data, all of the regression functions in Table 1 can accurately fit $v(n)$. When fit using $|\mathcal{D}_0| = 600,000$ images ($\approx 50\%$ of the data set), Figure 1 shows that each link function (dashed curves) achieves at most 6% error from the ground truth accuracy when extrapolating. Although power laws are theoretically motivated [3, 16], is there empirical justification for using them over other functions?

Extrapolating accuracy with small data sets is hard. With limited data, all of the regression functions extrapolate $v(n)$ poorly. Figure 1 shows how each curve (dotted curves) diverges significantly from the ground truth when fitting with $|\mathcal{D}_0| = 125,000$ images ($\approx 10\%$ of the data set). Further, some curves provide better fit than power laws. This small data regime was observed in [14, 26] who proposed jointly regressing on data set and model size; while this improves extrapolating performance, it also requires a $2\times$ larger \mathcal{R} obtained by sampling subsets and modifying different models. This can grow computationally expensive and time-consuming; as a result, we focus on simple estimators using a small number of training statistics, i.e., $r \leq 10$. **Small accuracy errors yield large data errors.** Suppose we must build a model meeting 67% test set accuracy on ImageNet, which requires 900,000 data points. Even though the functions fit using 600,000 images achieve error $|67\% - \hat{v}(900,000; \theta)|$ between 1 to 6%, they misestimate the data requirement between 120,000 to 310,000 images—collecting up to 34% less data than actually required. Since the tolerance for extrapolation errors is low, we must determine best practices for estimating data needs.

4. Empirical findings

We investigate the three challenges using regression and simulation over different data sets and tasks. We first summarize our experimental setup before analyzing the results.

4.1. Data and methods

We assess the data collection problem on image classification, object detection, and semantic segmentation

tasks summarized in Table 2. In classification, we train ResNets [13] on the CIFAR10 [21], CIFAR100 [21], and ImageNet [7] data sets, where we determine the amount of data needed to meet a target validation set accuracy. We train SSD300 [22] for 2-D object detection using the PASCAL VOC data sets [8,9], where we determine the amount of data needed to meet a target mean average precision (AP). For 3-D object detection, we train the FCOS3D network architecture [32] on different subsets of the nuScenes training set. We report mean average precision (mAP) following the nuScenes 3D detection evaluation protocol [5]. Samples are obtained randomly across different scenes. We explore semantic segmentation using BDD100K [34], which is a large-scale driving dataset collected over 50K drives with various geographic, environmental, and weather conditions. For multi-view Bird’s-Eye-View (BEV) segmentation, we train the “Lift Splat” architecture [24] on the nuScenes data set [5]. Here, we report mean intersection-over-union (IoU). For each task, we fix the architecture of the model and learning algorithm including data sampling. Details are in the supplementary content.

For each data set and task, we have an initial dataset \mathcal{D}_0 (e.g. $n_0 = 10\%$ of the training data set). In our analyses, we report n_0 in terms of the relative size of \mathcal{D}_0 w.r.t. the full training data set. We first create a regression data set \mathcal{R} according to Algorithm 1 by sampling r subsets that grow linearly in size (i.e. each $|\mathcal{S}_i| = |\mathcal{D}_0|(i+1)/r$ for all $i \in \{0, \dots, r-1\}$). To ensure that this regression procedure is inexpensive, we use a small $r \leq 10$. Then to evaluate our regression functions on extrapolating performance and estimating data requirements, we sample larger subsets $\mathcal{D}_1 \subset \mathcal{D}_2 \subset \dots$ of growing size (e.g. 10%, 20%, 30%, ..., 100% of the full training data set). For each subset, we train our model and evaluate the score $V_f(\mathcal{D}_i)$. Using these sets, we construct the piecewise linear score function $v(n)$, which we use as a ground truth.

We perform two types of experiments. In the first, preliminary, analysis we fit each regression function from Table 1 using \mathcal{R} and then evaluate their error with respect to predicting $V_f(\mathcal{D}_i)$ for all $|\mathcal{D}_i| > |\mathcal{D}_0|$. This analysis reveals how well each of the regression functions can extrapolate the model’s score for larger data sets. Our second, main, analysis is a simulation of the data collection problem in Algorithm 1 where we initialize with $n_0 = 10\%$ of the full training data set ($n_0 = 20\%$ for VOC) and estimate how much data is needed to obtain different target values V^* within $T = 1, 3, 5$ rounds. Here, we repeat the same steps described in the Data Collection stage of Algorithm 1, except with one difference. In our simulations, rather than sampling more data and evaluating $V_f(\mathcal{D}_0 \cup \hat{\mathcal{D}})$ in each round (e.g. lines 11–13), we evaluate $v(n_0 + \hat{n})$ to obtain the model score. This simulation approximates the true data collection problem, while simplifying experimen-

	Data set	n_0	r	Power Law	Arctan	Logarithmic	Algebraic Root
Classification	CIFAR10	10%	5	39.02 ± 20.3	7.98 ± 7.1	32.28 ± 13.1	33.63 ± 22.1
	CIFAR10	20%	10	15.26 ± 1.3	1.0 ± 0.6	11.53 ± 1.5	4.97 ± 1.6
	CIFAR10	50%	17	6.0 ± 0.5	0.38 ± 0.3	4.4 ± 0.5	0.76 ± 0.4
	CIFAR100	10%	5	34.38 ± 35.1	13.3 ± 5.3	17.25 ± 21.8	26.29 ± 16.8
	CIFAR100	20%	10	29.52 ± 3.9	4.71 ± 2.0	19.87 ± 2.5	40.33 ± 1.5
	CIFAR100	50%	17	5.49 ± 0.2	0.69 ± 0.2	5.42 ± 0.2	3.65 ± 0.3
	ImageNet	10%	4	23.89 ± 7.4	3.19 ± 2.1	17.2 ± 3.2	60.1 ± 1.1
	ImageNet	20%	8	10.12 ± 0.4	2.38 ± 0.5	9.46 ± 0.6	7.61 ± 1.0
	ImageNet	50%	15	5.06 ± 0.1	0.74 ± 0.2	3.81 ± 0.2	1.64 ± 0.2
Detection	VOC	20%	4	4.66 ± 3.1	2.98 ± 1.6	3.23 ± 2.1	3.28 ± 1.8
	VOC	30%	6	3.16 ± 1.6	2.31 ± 1.2	2.55 ± 1.3	2.83 ± 1.3
	VOC	50%	10	1.15 ± 0.5	0.79 ± 0.5	1.08 ± 0.4	1.13 ± 0.5
	nuScenes	10%	2	6.57 ± 0.5	13.43 ± 0.3	0.79 ± 0.2	4.53 ± 0.4
	nuScenes	20%	4	2.10 ± 2.1	1.65 ± 1.0	1.73 ± 1.3	2.32 ± 1.6
	nuScenes	50%	6	0.69 ± 0.2	0.71 ± 0.1	0.51 ± 0.2	0.36 ± 0.2
Segmentation	BDD100K	10%	5	9.85 ± 7.9	8.12 ± 7.6	9.18 ± 8.9	5.82 ± 2.3
	BDD100K	20%	10	2.98 ± 1.2	0.76 ± 0.3	1.60 ± 0.9	2.76 ± 1.2
	BDD100K	50%	17	1.30 ± 0.5	0.95 ± 0.3	0.82 ± 0.2	1.10 ± 0.5
	nuScenes	10%	5	2.78 ± 0.0	2.30 ± 0.7	2.03 ± 0.9	1.47 ± 0.6
	nuScenes	20%	10	0.61 ± 0.2	3.34 ± 0.6	0.91 ± 0.7	2.31 ± 1.0
	nuScenes	50%	17	0.38 ± 0.3	2.40 ± 0.1	0.28 ± 0.2	1.77 ± 1.7

Table 3. Mean±standard deviation of multiple runs evaluating the RMSE on extrapolating performance in each task when trained on small subsets of the data. We report n_0 in terms of the percentage of the true data set. The lowest error for each setting is bolded. We provide regression plots and alternate error metrics in the supplementary content. Given 50% of the data, every function achieves low regression error, whereas, with 10% of the data all of the functions have significant error in their estimation. Furthermore, the alternative functions typically outperform the Power Law across different values of n_0 and over different tasks.

tation since we do not have to repeatedly re-train our model.

4.2. Analysis

Regression. Table 3 summarizes the Root Mean Squared Error (RMSE) of each regression function when extrapolating the score for larger data sets. In each data set and task, we perform three runs with different random seeds, showing how well we can extrapolate with small, medium, and large subsets of the data. In the supplement, we provide regression plots for $v(n)$ versus $\hat{v}(n; \theta^*)$ and a table summarizing regression error in terms of relative error ratio.

We validate that the first two challenges mentioned in Section 3.2 hold for every task that we consider. Given a sufficient amount of initial data \mathcal{D}_0 to fit a regression model (*i.e.* when n_0 is equal to 50% of the full data set size), every link function achieves a low RMSE (whose range is the interval $[0, 100]$). Moreover, there is always at least one regression function that achieves an RMSE less than 1. When n_0 is equal to 10% of the full data set size, most of the link functions yield high RMSE, suggesting that the functions are susceptible to diverging from the true $v(n)$ when fitted on a small data set. Finally, for most data sets, our alternative regression functions consistently yield low RMSE. In particular, the Arctan function is the best for all of the classification data sets, and often cuts the RMSE from the Power Law by half. These results show that extrapolating model performance from small data sets is difficult, but furthermore, other regression functions instead of the Power Law may obtain more accurate regressions of the score.

Simulation. We simulate data collection for each of the different regression functions by sweeping a range of targets V^* when n_0 and T are given. Figure 3 reports the ratio

of the final data collected by each function versus the minimum data required according to the ground truth score, *i.e.* $(n_0 + \hat{n}) / (n_0 + n^*)$ where n^* is the smallest value satisfying $v(n_0 + n^*) = V^*$. The value of n^* is easy to find since $v(n)$ is a piecewise linear monotonically increasing function.

In evaluating how each regression function collects data, there are two scenarios to consider. If the ratio is less than one, the function is described as an optimistic predictor of the score that under-estimates how much data will be needed. A ratio less than one means that using this regression function, we will not collect enough data to meet V^* within T rounds, thereby failing to solve the problem. On the other hand if the ratio is greater than one, the function is a pessimistic predictor that over-estimates how much data will be needed. An ideal data collection policy will achieve the smallest ratio greater than one. Our experiments show that in general, the Arctan function is the most pessimistic and often achieves the largest ratios by a large margin.

We first validate the third challenge from Section 3.2 and remark that low regression error does not necessarily translate to better data collection. On CIFAR100, ImageNet, and VOC, using Arctan may lead to collecting up to five times more data than is actually needed; with BEV segmentation on nuScenes, it may lead to over 10 times more. Recall from Figure 1 that on ImageNet, we require approximately 900,000 images to reach a target $V^* = 67\%$. Using Arctan when initialized with $n_0 = 10\%$ of the data would result in collecting approximately 4.5 million images in the first round alone, whereas all of the other regression functions achieve a ratio approximately equal to 1. Although Table 3 showed that Arctan achieved the lowest RMSE (3.19) of all functions in regression, using it to estimate data re-

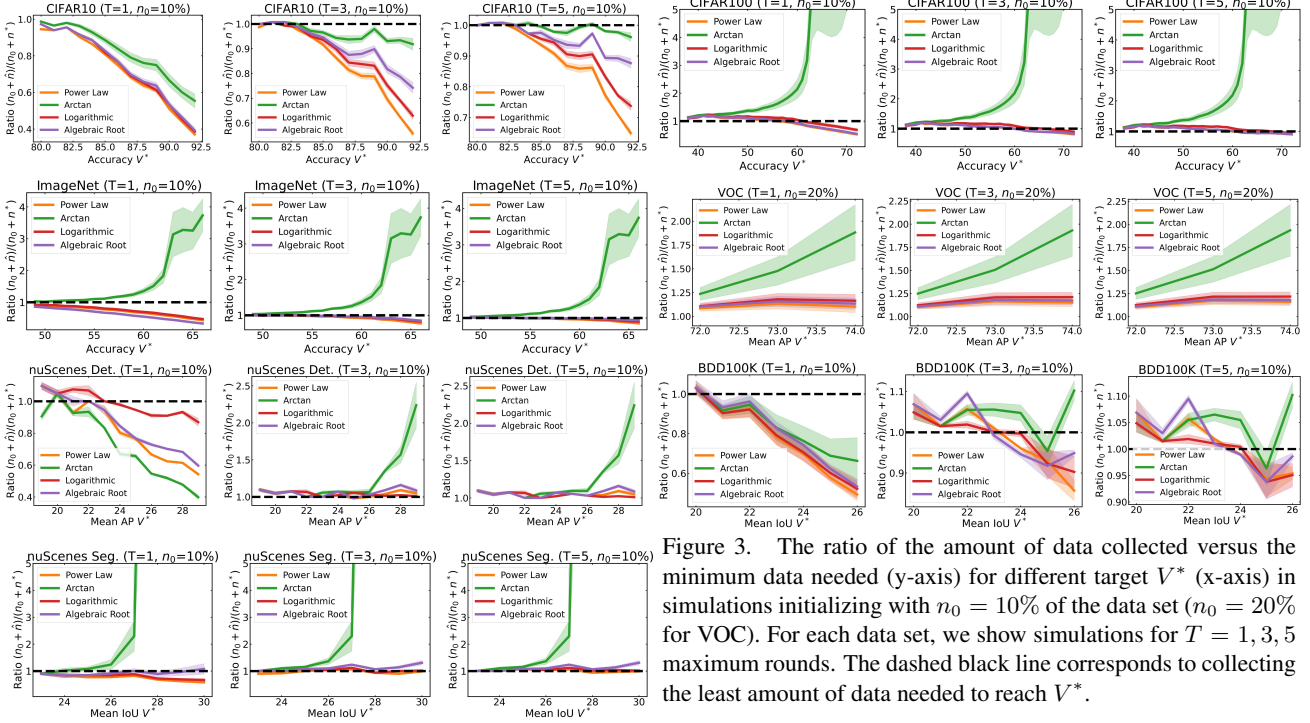


Figure 3. The ratio of the amount of data collected versus the minimum data needed (y-axis) for different target V^* (x-axis) in simulations initializing with $n_0 = 10\%$ of the data set ($n_0 = 20\%$ for VOC). For each data set, we show simulations for $T = 1, 3, 5$ maximum rounds. The dashed black line corresponds to collecting the least amount of data needed to reach V^* .

quirements would lead to an unnecessarily expensive data collection procedure. This reveals that simply analyzing regression error is insufficient when determining good data collection policies, necessitating our simulation approach.

For most regression functions, collecting enough data requires multiple rounds. When $T = 1$, the Power Law, Logarithmic, and Algebraic Root functions under-estimate the data requirement for all data sets and tasks except for VOC. However when $T = 5$, for every data set except for CIFAR10, all of the functions yield ratios greater than 0.9 over the entire range of V^* . That is, we can consistently reach at least 90% of the data needed with any of the functions.

Ultimately, even with $T = 5$, these estimators can still under-estimate the requirement when V^* is large (e.g. on ImageNet, the Power Law, Logarithmic, and Algebraic Root functions achieve ratios less than 1 for $V^* \geq 62\%$). From an operational perspective, although these methods do not incur large costs, they also fail to solve the problem. In the next section, we show simple techniques to correct these estimators and better guide data collection.

Ablations. In the supplement, we perform ablations that evaluate regressions and simulations on different model depths and widths for CIFAR100. We also consider alternate score functions such as collecting enough data to meet a target performance on a specific class using nuScenes. Finally, we explore estimating requirements when using active learning rather than random sampling for CIFAR100. Our results indicate the same trends, further supporting the challenges towards estimating the data requirement.

5. Towards better estimates of data

We previously showed that some optimistic estimators fail to collect enough data to meet V^* whereas other pessimistic estimators lead to collecting far more data than required. Here, we first introduce a correction factor, which is a bias term that addresses the problem of under-estimating data requirements. We then show how analyzing both the optimistic and pessimistic regression functions considered in this paper can lead to a collection of estimates that often bound the true data requirement.

5.1. A correction factor to help meet the target

From Algorithm 1, in each round of data collection, we minimize \hat{n} subject to $\hat{v}(n_0 + \hat{n}; \theta^*) \geq V^*$. Ideally, we would want to minimize the true data requirement, i.e. solving for n^* satisfying $v(n_0 + n^*) = V^*$. However, our simulations show that most of the regression functions are optimistic and under-estimate how much data is needed. Intuitively, a simple way to correct for collecting less than the data needed to meet V^* is to impose a correction factor $\tau \geq 0$ and instead estimate the data required to meet a “corrected” higher target $V^* + \tau$. As a result, we fix a constant τ and modify Algorithm 1 so that in each round, we now minimize \hat{n} subject to $\hat{v}(n_0 + \hat{n}; \theta^*) \geq V^* + \tau$.

In order to determine how large this correction factor should be, we treat it as a hyper-parameter to fit. For instance, suppose that we have the full CIFAR10 data set and we want to construct a T -round collection policy for future data sets. We first simulate data collection with $\tau = 0$ for CIFAR10 with each regression function to obtain the plots

	Data set	n_0	T	Power Law		Arctan		Logarithmic		Algebraic Root	
				Without	With	Without	With	Without	With	Without	With
Classification	CIFAR100	10%	1	0.53	0.91	1.13	1.36	0.68	<u>1.54</u>	0.54	0.82
	CIFAR100	10%	3	0.81	<u>1.09</u>	1.13	1.19	0.9	1.08	0.83	0.94
	CIFAR100	10%	5	0.9	<u>1.03</u>	1.13	1.19	0.94	<u>1.11</u>	0.91	1.01
Classification	ImageNet	10%	1	0.43	<u>1.16</u>	1.02	1.35	0.47	<u>1.28</u>	0.33	0.5
	ImageNet	10%	3	0.77	<u>1.10</u>	1.03	1.08	0.83	<u>1.06</u>	0.85	1.03
	ImageNet	10%	5	0.85	<u>1.07</u>	1.03	1.08	0.9	<u>1.06</u>	0.94	1.03
Detection	VOC	20%	1	1.08	6.42	1.24	5.05	1.11	7.4	1.1	6.03
	VOC	20%	3	1.1	2.75	1.25	1.64	1.12	2.23	1.11	1.54
	VOC	20%	5	1.1	2.03	1.25	1.64	1.13	2.23	1.11	1.54
Detection	nuScenes	10%	1	0.56	<u>2.9</u>	0.39	0.51	0.83	<u>32.45</u>	0.61	<u>2.9</u>
	nuScenes	10%	3	0.94	<u>1.05</u>	1.0	1.0	1.0	1.68	0.94	<u>1.07</u>
	nuScenes	10%	5	1.0	1.09	1.0	1.0	1.0	1.68	1.0	1.07
Segmentation	BDD100K	10%	1	0.49	<u>2.45</u>	0.66	1.79	0.52	<u>5.2</u>	0.53	<u>2.17</u>
	BDD100K	10%	3	0.86	<u>1.76</u>	0.95	1.2	0.9	<u>1.58</u>	0.92	<u>1.19</u>
	BDD100K	10%	5	0.94	<u>1.48</u>	0.96	1.2	0.94	<u>1.58</u>	0.94	<u>1.19</u>
Segmentation	nuScenes	10%	1	0.58	<u>24.58</u>	0.9	4.63	0.67	<u>27.46</u>	0.83	<u>51.12</u>
	nuScenes	10%	3	0.9	<u>1.42</u>	1.0	1.25	0.94	<u>1.31</u>	1.0	1.25
	nuScenes	10%	5	0.94	<u>1.07</u>	1.0	1.25	1.0	<u>1.31</u>	1.0	1.25




Figure 4. For $T = 5$, the ratio of the amount of data collected versus the minimum data needed to meet different target V^* when using regression functions with correction factors fitted using CIFAR10.

in Figure 3. We then increase τ until the entire ratio curve for that function is above 1. In other words, we solve for the smallest τ such that the data collection policy will collect just enough data to meet all target values V^* for CIFAR10 (for a given fixed T and function). We then use this fitted τ as a correction factor for future data sets.

By combining the correction factor with multiple rounds of data collection, we can consistently collect just above the minimum data requirement. Table 4 compares the effect of using τ for each of the regression functions on the minimum ratio over all V^* for each data set. We use the CIFAR10 data set to fit τ for each setting of T and regression function. Without correction, the Power Law, Logarithmic, and Algebraic Root functions achieve ratios less than 1 for every data set except VOC. Using τ , these functions almost always achieve ratios between 1 to 2. Furthermore for each data set, these three regression functions achieve their respective lowest ratios (above 1) when $T = 5$. Figure 4 further plots simulations using τ over all V^* for each data set with $T = 5$. Here, the Power Law, Logarithmic, and Algebraic Root functions achieve ratios between 1.03 to 2.5 for all V^* with every data set. Furthermore, there is no consistently best regression function for all data sets. For instance, the Algebraic Root function dominates over VOC, but the Power Law is particularly effective on nuScenes BEV segmentation when V^* is large. However, recall that Arctan naturally over-estimates the data requirement, so this func-

Table 4. The minimum ratio $\frac{n_0 + \hat{n}}{n_0 + n^*}$ for each regression function without (baseline) and with using τ when estimating data requirements. The best ratio (*i.e.* smallest value greater than 1) for each data set is bolded. Instances where using τ for a given regression function increased the ratio from below to above 1 are underlined. Power Law, Logarithmic, and Algebraic Root improve for nearly every setting to obtain ratios above 1. Furthermore, these functions achieve their best performance when $T = 5$.

tion does not benefit from correction. We conclude that correcting any of the three optimistic estimators, Power Law, Logarithmic, or Algebraic Root, and collecting data over five rounds is enough to approximately minimize the total data collected while still meeting the desired target.

5.2. Empirical bounds on the data requirement

If the correction factor is poorly fit or the number of rounds is constrained to be small, we may still under- or over-estimate the data requirement. From Table 4, on nuScenes segmentation with $T = 1$, the Power Law without τ leads to estimating 58% of the required data, whereas using τ leads to estimating 28 times more data than needed.

In some applications, modelers may also desire rule-of-thumb estimates of the amount of data that they should immediately collect. We now consider the problem where we have a single $T = 1$ round remaining with n_0 data points; in a single round or the final of multiple rounds, we must meet the data collection target. As a result here, we seek to obtain worst and best-case estimates (*i.e.* upper and lower bounds) on how much more data is needed. All of the different regression functions yield an ensemble of predictions. Then, the largest prediction is the worst-case estimate and the smallest prediction is the best-case estimate.

For each data set, we set $T = 1$ and sweep over n_0 and V^* to estimate the data requirement with eight regression functions. Figure 5 (top row) shows for each n_0 , the

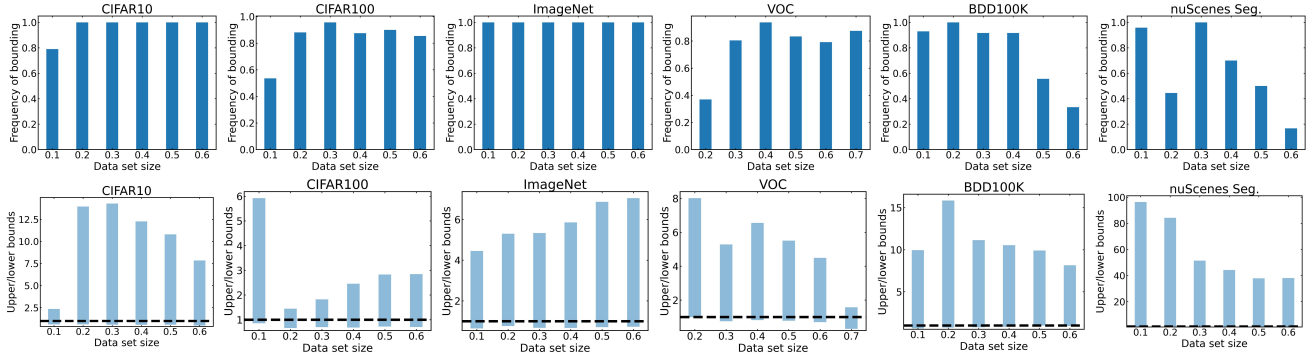


Figure 5. (Top row) For $T = 1$ and varying n_0 , the frequency of instances where the largest and smallest \hat{n} estimated by the different functions upper and lower bound the true n^* . (Bottom row) The largest and smallest ratios $(n_0 + \hat{n})/(n_0 + n^*)$ estimated by the different functions. The dashed black line corresponds to ratio 1.

frequency of instances of V^* in which the most optimistic and pessimistic regression function bound the true data requirement. The bottom row further shows the average value of these upper and lower bounds. For image classification, our estimators bound the true requirement over 80% of the time. This trend also holds for VOC wherein we bound the requirement over 80% of the time for $n_0 \geq 30\%$ of the data set. Since BDD100K and nuScenes BEV segmentation are more challenging data sets, our probability of bounding the data requirement can at times decrease. Because training the 3-D object detector on nuScenes is computationally far more expensive than the other tasks, we omit their plots and only report values for $n_0 = 10\%, 20\%, 50\%$. Here, the range of estimators bound the true data requirement 88%, 91%, and 83% of the time with ratios in the interval $[0.56, 31.1]$, $[0.76, 40.8]$, $[0.56, 26.9]$, respectively. Nonetheless, the results show that if we are given a single round with a large initial data set, we will be able to accurately estimate upper and lower bounds on the data requirement. Moreover, even if we are given multiple rounds to collect data, on the final round, we should be able to obtain upper and lower intervals for the requirement. In practical applications, these bounds can guide modelers to make optimistic or pessimistic choices, for example if the real-world deadline for training and deploying a model is strict.

6. Discussion

In this work, we propose an effective solution to the problem of estimating how much data must be collected to meet a target performance. While the problem of predicting a model’s performance has received growing research interest as a springboard for various design decisions, we find that estimating performance does not capture the downstream problem of estimating data requirements. Even small errors in predicting performance can yield large errors in data collection, meaning that the error permissible from a good data estimator is far smaller than intuition suggests. Furthermore, errors are divided into under- or over-

estimation, where each poses different challenges to data collection. To better analyze data collection strategies, we formulate an iterative data collection simulation. Our experiments draw several high-level insights:

- Different techniques estimate either far more data or far less data than needed. Using multiple rounds of data collection with techniques that under-estimate can lead to collecting up to 90% of the true amount of data needed.
- By simulating on previous tasks, we can identify which approaches under-estimate data requirements and learn a correction factor to address this deficiency. Using a correction factor and collecting for up to five rounds allows us to collect at most one to two times the minimum amount of data needed for any desired performance.
- With only one round of data collection remaining, we can use all of the different regression functions to obtain an interval that often bounds the true data requirement. These bounds can guide modelers to collect data more or less aggressively with respect to practical requirements.

Limitations. The data collection problem and the simulation proposed in this work approximate real collection practices. Our simulation relies on a pre-constructed ground truth $v(n)$ rather than sampling points, training a model, and evaluating $V_f(\mathcal{D})$. The latter is computationally too expensive to perform for the range of settings explored in this paper. The quality of our simulation depends on the number of subsets used to construct $v(n)$. More subsets means $v(n)$ better approximates $V_f(\mathcal{D})$ and from inspection (see the supplementary content), all of our $v(n)$ appear to be visually smooth curves. Moreover in our data collection problem, we assume that the model f and sampling strategy $p(z)$ are constant. In practice, designers may update f in between rounds; this may be incorporated in a more complete model of the deep learning workflow. In addition, secondary metrics can be used to optimize $p(z)$. For example, if a classifier is particularly poor for a single class in a given round, modelers may seek to obtain more samples of that specific class in the next round. We leave these more sophisticated problem settings to future research.

References

- [1] Samira Abnar, Mostafa Dehghani, Behnam Neyshabur, and Hanie Sedghi. Exploring the limits of large scale pre-training. In *International Conference on Learning Representations*, 2022. 2
- [2] David Acuna, Huan Ling, Amlan Kar, and Sanja Fidler. Efficient interactive annotation of segmentation datasets with polygon-rnn++. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1
- [3] Yasaman Bahri, Ethan Dyer, Jared Kaplan, Jaehoon Lee, and Utkarsh Sharma. Explaining neural scaling laws. *arXiv preprint arXiv:2102.06701*, 2021. 2, 4
- [4] Devansh Bisla, Apoorva Nandini Saridena, and Anna Chormanska. A theoretical-empirical approach to estimating sample complexity of dnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3270–3280, 2021. 1, 2
- [5] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11621–11631, 2020. 4
- [6] David A Cohn, Zoubin Ghahramani, and Michael I Jordan. Active learning with statistical models. *Journal of Artificial Intelligence Research*, 4:129–145, 1996. 2
- [7] 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, 2, 4
- [8] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>. 4
- [9] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>. 4
- [10] Rosa L Figueroa, Qing Zeng-Treitler, Sasikiran Kandula, and Long H Ngo. Predicting sample size required for classification performance. *BMC Medical Informatics and Decision Making*, 12(1):1–10, 2012. 2
- [11] Lewis J Frey and Douglas H Fisher. Modeling decision tree performance with the power law. In *Seventh International Workshop on Artificial Intelligence and Statistics*. PMLR, 1999. 1, 2, 3
- [12] Baohua Gu, Feifang Hu, and Huan Liu. Modelling classification performance for large data sets. In *International Conference on Web-Age Information Management*, pages 317–328. Springer, 2001. 1
- [13] 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. 4
- [14] Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md Patwary, Mostofa Ali, Yang Yang, and Yanqi Zhou. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409*, 2017. 2, 4
- [15] Derek Hoiem, Tanmay Gupta, Zhizhong Li, and Michal Shlapentokh-Rothman. Learning curves for analysis of deep networks. In *International Conference on Machine Learning*, pages 4287–4296. PMLR, 2021. 2
- [16] Marcus Hutter. Learning curve theory. *arXiv preprint arXiv:2102.04074*, 2021. 4
- [17] Yiding Jiang, Pierre Foret, Scott Yak, Daniel M Roy, Hossein Mobahi, Gintare Karolina Dziugaite, Samy Bengio, Suriya Gunasekar, Isabelle Guyon, and Behnam Neyshabur. Neurips 2020 competition: Predicting generalization in deep learning. *arXiv preprint arXiv:2012.07976*, 2020. 2
- [18] Yiding Jiang, Parth Natekar, Manik Sharma, Sumukh K Aithal, Dhruva Kashyap, Natarajan Subramanyam, Carlos Lassance, Daniel M Roy, Gintare Karolina Dziugaite, Suriya Gunasekar, et al. Methods and analysis of the first competition in predicting generalization of deep learning. In *NeurIPS 2020 Competition and Demonstration Track*, pages 170–190. PMLR, 2021. 2
- [19] S Jones, S Carley, and M Harrison. An introduction to power and sample size estimation. *Emergency Medicine Journal: EMJ*, 20(5):453, 2003. 2
- [20] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020. 2
- [21] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009. 4
- [22] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Proceedings of the European Conference on Computer Vision*, pages 21–37. Springer, 2016. 4
- [23] Rafid Mahmood, Sanja Fidler, and Marc T. Law. Low-budget active learning via wasserstein distance: An integer programming approach. In *International Conference on Learning Representations*, 2022. 2
- [24] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *Proceedings of the European Conference on Computer Vision*, 2020. 4
- [25] Fait Poms, Vishnu Sarukkai, Ravi Teja Mullapudi, Nimit S Sohoni, William R Mark, Deva Ramanan, and Kayvon Fatahalian. Low-shot validation: Active importance sampling for estimating classifier performance on rare categories. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10705–10714, 2021. 2
- [26] Jonathan S Rosenfeld, Amir Rosenfeld, Yonatan Belinkov, and Nir Shavit. A constructive prediction of the generalization error across scales. In *International Conference on Learning Representations*, 2020. 2, 3, 4

- [27] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018. 2
- [28] Burr Settles. Active learning literature survey. 2009. 2
- [29] Samarth Sinha, Sayna Ebrahimi, and Trevor Darrell. Variational adversarial active learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5972–5981, 2019. 2
- [30] 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. 2
- [31] Tom Viering and Marco Loog. The shape of learning curves: a review. *arXiv preprint arXiv:2103.10948*, 2021. 2, 3
- [32] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 913–922, October 2021. 4
- [33] Donggeun Yoo and In So Kweon. Learning loss for active learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 93–102, 2019. 2
- [34] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 2636–2645, 2020. 4
- [35] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2