
On the Predictability of Pruning Across Scales

Jonathan Rosenfeld¹ Jonathan Frankle¹ Michael Carbin¹ Nir Shavit¹

Abstract

We show that the error of iteratively magnitude-pruned networks empirically follows a scaling law with interpretable coefficients that depend on the architecture and task. We functionally approximate the error of the pruned networks, showing it is predictable in terms of an invariant tying width, depth, and pruning level, such that networks of vastly different pruned densities are interchangeable. We demonstrate the accuracy of this approximation over orders of magnitude in depth, width, dataset size, and density. We show that the functional form holds (generalizes) for large scale data (e.g., ImageNet) and architectures (e.g., ResNets). As neural networks become ever larger and costlier to train, our findings suggest a framework for reasoning conceptually and analytically about a standard method for unstructured pruning.

1. Introduction

For decades, neural network *pruning*—eliminating unwanted parts of a network—has been a popular approach for reducing network sizes or the computational demands of inference (LeCun et al., 1990; Reed, 1993; Han et al., 2015). In practice, pruning can reduce the parameter-counts of contemporary models by 2x (Gordon et al., 2020) to 5x (Renda et al., 2020) with no increase in error. More than 80 pruning techniques have been published in the past decade (Blalock et al., 2020), but, despite this enormous volume of research, there remains little guidance on important aspects of pruning. Consider a seemingly simple question one might ask when using a particular pruning technique:

Given a family of neural networks (e.g., ResNets on ImageNet of various widths and depths), which family member should we prune (and by how much) to obtain the network with the smallest parameter-count such that error does not exceed some threshold ϵ_k ?

¹MIT CSAIL. Correspondence to: Jonathan Rosenfeld <jonsr@csail.mit.edu>.

As a first try, we could attempt to answer this question using brute force: we could prune every member of a network family (i.e., perform grid search over widths, depth, and pruned densities) and select the smallest pruned network that satisfies our constraint on error. However, depending on the technique, pruning one network (let alone grid searching) could take days or weeks on expensive hardware.

If we want a more efficient alternative, we will need to make assumptions about pruned networks: namely, that there is some *structure* to the way that their error behaves. For example, that pruning a particular network changes the error in a predictable way. Or that changing its width or depth changes the error when pruning it in a predictable way. We could then train a smaller number of networks, characterize this structure, and estimate the answer to our question.

We have reason to believe that such structure does exist for certain pruning methods, since there are already techniques that take advantage of it implicitly. For example, Cai et al. (2019) create a single neural network architecture that can be scaled down to many different sizes; to choose which sub-network to deploy, Cai et al. train an auxiliary, black-box neural network to predict subnetwork performance. Although this black-box approach implies the existence of structure for this pruning method, it does not reveal this structure explicitly or make it possible to reason analytically in a fashion that could answer our research question.

For other aspects of deep learning beyond pruning, such structure has been observed—and, further, codified explicitly—yielding insights and predictions in the form of scaling laws. Tan & Le (2019) design the *EfficientNet* family by developing a heuristic for predicting efficient tradeoffs between depth, width, and resolution. Hestness et al. (2017) observe a power-law relationship between dataset size and the error of vision and NLP models. Rosenfeld et al. (2020) use a power law to predict the error of all variations of architecture families and dataset sizes, jointly, for computer vision and natural language processing settings. Kaplan et al. (2020) develop a similar power law for language models that incorporates the computational cost of training.

Inspired by this work, we address our research question about pruning by finding a scaling law to predict the error of pruned networks. We focus on a pruning method called *iterative magnitude pruning (IMP)*, where weights with the

lowest magnitudes are pruned in an unstructured fashion interspersed with re-training to recover accuracy. This method is a standard way to prune (Han et al., 2015) that gets state-of-the-art tradeoffs between error and unstructured density (Gale et al., 2019; Renda et al., 2020). To the best of our knowledge, this is the first explicit scaling law that holds for pruned networks, let alone entire network families.

To formulate such a predictive scaling law, we consider the dependence of generalization error on the pruning-induced *density* for networks of different depths and widths trained on different dataset sizes. We begin by developing a functional form that accurately estimates the generalization error of a specific model as it is pruned (Section 3). We then account for other architectural degrees of freedom, expanding the functional form for pruning into a scaling law that jointly considers density alongside width, depth, and dataset size (Section 4). The basis for this joint scaling law is an *invariant* we uncover that describes ways that we can interchange depth, width, and pruning without affecting error. The result is a scaling law that accurately predicts the error of pruned networks across scales. And, now that we have established this functional form, fitting it requires only a small amount of data (Section 5), making it efficient to use on new architectures and datasets (Appendix E). Finally, we use this scaling law to answer our motivating question (Section 7). In summary, our contributions are as follows:

- We develop a scaling law that accurately estimates the error when pruning a single network with IMP.
- We observe and characterize an *invariant* that allows error-preserving interchangeability among depth, width, and pruning density.
- Using this invariant, we extend our single-network scaling law into a joint scaling law that predicts the error of all members of a network family at all dataset sizes and all pruning densities.
- In doing so, we demonstrate that there is structure to the behavior of the error of iteratively magnitude-pruned networks that we can capture explicitly with a simple functional form and interpretable parameters.
- Our scaling law enables a framework for reasoning analytically about IMP, allowing us to answer our motivating question and similar questions about pruning.

2. Experimental Setup

Pruning. We study *iterative magnitude pruning* (IMP) (Janowsky, 1989; Han et al., 2015). IMP prunes by removing a fraction—typically 20%, as we do here—of individual weights with the lowest magnitudes in an unstructured fashion at the end of training.¹ We choose these weights globally

¹We do not prune biases or BatchNorm, so pruning 20% of weights prunes fewer than 20% of parameters.

throughout the network, i.e., without regard to specific layers. We use per-weight magnitude pruning because it is generic, well-studied (Han et al., 2015), and produces state-of-the-art tradeoffs between density and error (Gale et al., 2019; Blalock et al., 2020; Renda et al., 2020).

Pruning weights typically increases the error of the trained network, so it is standard practice to further train after pruning to reduce error. For IMP, we use a practice called *weight rewinding* (Frankle et al., 2020; Renda et al., 2020), in which the values of unpruned weights are *rewound* to their values earlier in training (in our case, epoch 10) and the training process is repeated from there to completion. To achieve density levels below 80%, this process is repeated *iteratively*—pruning by 20%, rewinding, and retraining—until a desired density level is reached. For a formal statement of this pruning algorithm, see Appendix A.

Datasets. In the main body of the paper, we study the image classification tasks CIFAR-10 and ImageNet. Our scaling law predicts the error when training with the entire dataset and smaller *subsamples*. We include subsampling because it provides a cost-effective way to collect some of the data for fitting our functional form. To subsample a dataset to a size of n , we randomly select n of the training examples without regard to individual classes such that in expectation we preserve the original dataset distribution (we always retain the entire test set). When performing iterative pruning, we maintain the same subsample for all pruning iterations. We consider other datasets in Appendix E.

Networks. In the main body of the paper, we study ResNets for CIFAR-10 and ImageNet.² We develop a scaling law that predicts the error (when pruned) of an entire *family* of networks with varying widths and—in the case of the CIFAR-10 ResNets—depths. To vary width, we multiply the number of channels in each layer by a *width scaling factor*. To vary depth of the CIFAR-10 ResNets, we vary the number of residual blocks. We refer to a network by its depth l (the number of layers in the network, not counting skip connections) and its width scaling factor w . We consider other networks in Appendix E.

Notation and terminology. Throughout the paper, we use the following notation and terminology:

- $\mathbb{D}_N = \{\mathbf{x}_i, y_i\}_{i=1}^N$ is a labeled training set with N examples. A *subsample* of size n is a subset of \mathbb{D}_N with n examples selected uniformly at random.
- l and w are, respectively, the depth (i.e., the number of layers, excluding skip connections) and the width scaling factor of a particular network.
- Networks that vary by width and depth are a *family*.

²See Appendix B for full details on architectures and hyperparameters. Note that CIFAR-10 and ImageNet ResNets have different architectures (He et al., 2016).

Network Family	Densities (d)	Depths (l)	Width Scalings (w)	Subsample Sizes (n)
CIFAR-10 ResNet	$0.8^i, i \subseteq \{0, \dots, 40\}$ (see below)	$l \subseteq \{8, 14, 20, 26, 50, 98\}$ (see below)	$2^i, i \subseteq \{-4, \dots, 2\}$ (see below)	$\frac{N}{i}, i \in \{1, 2, 4, 8, 16, 32, 64\}$
ImageNet ResNet	$0.8^i, i \subseteq \{0, \dots, 30\}$ (see below)	50	$2^i, i \subseteq \{-4, \dots, 0\}$ (see below)	$\frac{N}{i}, i \in \{1, 2, 4\}$

Table 1. The ranges of settings we consider in our experiments in the main body of the paper. We consider all densities $d \in \{0.8^i \mid i \in \mathbb{N}_0\}$ where (1) the network is not disconnected and (2) the network does better than random chance; this value varies based on the configuration (l, w, n) and the random seed. Where neither of these conditions apply, we cap i at 40 (CIFAR-10 ResNets) or 30 (ImageNet ResNets). We consider all configurations of (l, w, n) before which increasing depth or width of the unpruned network increases test error; configurations beyond this point are overly large for a given dataset subsample. By this criterion, we use 152 configurations of (l, w, n) for the CIFAR-10 ResNets and 15 for the ImageNet ResNets. Taking into account all feasible densities, we use a total of 4,301 CIFAR-10 ResNet points and 274 ImageNet ResNet points. Note that we use these configurations to find and evaluate the functional form. Once we do so, far fewer configurations are needed to fit the functional form for each setting (see Section 5).

- d is the *density* of a pruned network (i.e., the fraction of weights that have not been pruned).
- $\epsilon(d, l, w, n)$ is the test error of a network with the specified density, depth, width scaling, and dataset size.
- $\epsilon_{np}(l, w, n) = \epsilon(1, l, w, n)$ is the test error of the unpruned network with the specified depth, width scaling, and dataset size. When clear from context, we omit (w, l, n) and write ϵ_{np} .
- $\hat{\epsilon}(\epsilon_{np}, d \mid l, w, n)$ is an estimate of the error of a pruned model for a scaling law that has been fitted to a specific network with the specified depth, width scaling, and dataset size (Section 3).
- $\hat{\epsilon}(\epsilon_{np}, d, l, w, n)$ is an estimate of the error of a pruned model with the specified depth, width scaling, and dataset size for a scaling law that has been fitted to a network family (Section 4).

Dimensions. In developing our scaling laws, we vary four dimensions: dataset subsample size (n) and network degrees of freedom density (d), network depth (l), and width scaling factor (w). In the main body of the paper, we consider the ranges of these values as specified in Table 1. See the caption in Table 1 for full details. We train three replicates of each CIFAR-10 configuration with different seeds.

3. Modeling the Error of a Pruned Network

Our goal in this section is to develop a functional form that models the error of a member of a network family as it is pruned (using IMP) based on its unpruned error $\epsilon_{np}(w, l, n)$. In other words, we wish to find a function $\hat{\epsilon}(\epsilon_{np}, d \mid l, w, n)$ that predicts the error at each density d for a network with a specific depth l , width scaling factor w , and dataset size n .

Intuition. Since IMP prunes 20% at a time, it produces pruned networks at intermediate densities $d_k = 0.8^k$ in the process of pruning to density $d_K = 0.8^K$. In Figure 1 (left), we plot the error of these pruned networks for CIFAR-10 ResNets with depth $l = 20$ and different widths w . All of these curves follow a similar pattern:³

³The same patterns occur for l and n for CIFAR-10 and w and n for ImageNet (see Appendix C). We focus on width for CIFAR-10 here for illustration.

Observation 1: Low-error plateau. The densest networks (right part of curves) have similar error to the unpruned network: $\epsilon_{np}(w)$. We call this the *low-error plateau*.

Observation 2: Power-law region. When pruned further, error increases linearly on the logarithmic axes of the figure. Linear behavior on a logarithmic scale is the functional form of a *power law*, where error relates to density via exponent γ and coefficient c : $\epsilon(d, w) \approx cd^{-\gamma}$. In particular, γ is the slope of the line on the logarithmic axes.

Observation 3: High-error plateau. When pruned further, error again flattens; we call this the *high-error plateau* and call the error of the plateau ϵ^\dagger .

Figure 1 (center) labels these regions for CIFAR-10 ResNet-20 ($w = 1, n = 1$) and shows an approximation of these regions that is piece-wise linear on logarithmic axes. These observations are our starting point for developing a functional form that estimates error when pruning.

Functional form. Our next task is to find a functional form that captures these observations about the relationship between density and error. In prior work, Rosenfeld et al. (2020) observe that the relationship between width and error shares the same general shape: it has a region of lower error, a power-law region, and region of higher error. However, this relationship is different enough from the one we observe (see Appendix G) to merit an entirely new functional form.

To develop this functional form, we note that the three regions of the curves in Figure 1 (the low-error plateau, the power-law region, and the high-error plateau) can be described by three power laws: two plateaus with exponent zero and an intermediate region with exponent γ . A functional family that arises frequently in systems that exhibit different power-law regions is the *rational family*. The particular family member we consider is as follows:⁴

$$\hat{\epsilon}(\epsilon_{np}, d \mid l, w, n) = \epsilon_{np} \left\| \frac{d - j p \left(\frac{\epsilon^\dagger}{\epsilon_{np}} \right)^{\frac{1}{\gamma}}}{d - j p} \right\|^\gamma \text{ where } j = \sqrt{-1} \quad (1)$$

⁴The expression $\left\| \frac{d - j a}{d - j b} \right\|^\gamma = \left(\frac{d^2 + a^2}{d^2 + b^2} \right)^{\frac{\gamma}{2}}$ meaning Eq. 1 can be rewritten as $\epsilon_{np} \left[(d^2 + p^2 (\epsilon^\dagger / \epsilon_{np})^{2/\gamma}) / (d^2 + p^2) \right]^{\gamma/2}$




Figure 1. Relationship between density and error when pruning CIFAR-10 ResNets; w varies, $l = 20$, $n = N$ (left). Low-error plateau, power-law region, and high-error plateau when $l = 20$, $w = 1$, $n = N$ (center). Visualizing Eq. 1 and the roles of free parameters (right).




Figure 2. Estimated (blue dots) and actual error (solid lines) when pruning CIFAR-10 ResNets; w varies, $l = 20$, $n = N$ (left). Estimated versus actual error for the same networks (center). Estimated versus actual error for all CIFAR-10 ResNet configurations (right).

This function’s shape is controlled by ϵ_{np} , ϵ^\dagger , γ , and p (visualized in Figure 1, right). ϵ_{np} and ϵ^\dagger are the values of the low and high-error plateaus. γ is the slope of the power-law region on logarithmic axes. p controls the density where the high-error plateau transitions to the power-law region.

Fitting. To fit $\hat{\epsilon}(\epsilon_{np}, d | l, w, n)$ to actual data $\epsilon(d, l, w, n)$, we estimate values for the free parameters ϵ^\dagger , γ , and p by minimizing the relative error $\delta \triangleq \frac{\hat{\epsilon}(\epsilon_{np}, d | l, w, n) - \epsilon(d, l, w, n)}{\epsilon(d, l, w, n)}$ using least squares regression. The fit is performed separately for each configuration (l, w, n) for all densities, resulting in per-configuration estimates of $\hat{\epsilon}^\dagger$, $\hat{\gamma}$, and \hat{p} .

Evaluating fit. For a qualitative view,⁵ we plot the actual error⁶ $\epsilon(d, l, w, n)$ and the estimated error $\hat{\epsilon}(\epsilon_{np}, d | l, w, n)$ as a function density for CIFAR-10 ResNets of varying widths (Figure 2, left). Our estimated error appears to closely follow the actual error. The most noticeable deviations occur at large densities, where the error decreases slightly on the low-error plateau whereas we treat it as flat (see Section 6).

Quantitatively, we measure the extent to which estimated error departs from the actual error using the mean μ and standard deviation σ of the relative deviation δ . Figure 2

(center) compares the estimated and actual errors for the networks in Figure 2 (left); Figure 2 (right) shows the same comparison for all configurations of l , w , and n on CIFAR-10 and the more than 4,000 pruned ResNets that result. The relative deviation on all configurations has mean $\mu < 2\%$ and standard deviation $\sigma < 4\%$; this means that, if the actual error is 10%, the estimated error is $9.8 \pm 0.4\%$ ($\hat{\epsilon} = (1 - \delta)\epsilon$).

4. Jointly Modeling Error For All Dimensions

In Section 3, we found a functional form $\hat{\epsilon}(\epsilon_{np}, d | l, w, n)$ (Eq. 1) that accurately predicts the error when pruning a *specific* member of a network family (with depth l and width w) trained with a dataset of size n . The parameters governing Equation 1 (ϵ^\dagger , p , and γ) varied between and depended on the specific configuration of l , w , n . However, we are interested in a single *joint* scaling law $\hat{\epsilon}(\epsilon_{np}, d, l, w, n)$ that, given the unpruned network error $\epsilon_{np}(l, w, n)$, accurately predicts error across *all* dimensions: all members of a network family that vary in depth and width, all densities, and all dataset sizes. Importantly, the parameters of this scaling law must be constants as a function of all dimensions. In this section, we develop this joint scaling law.

Intuition: the error-preserving invariant. Our desired scaling law $\hat{\epsilon}(\epsilon_{np}, d, l, w, n)$ will be a four-dimensional function of d , w , l , and n . To develop an intuition for

⁵Error is a 4-dimensional function, so we can only qualitatively examine 2D projections. All such projections are in Appendix D.

⁶We compute the error as the mean across three replicates with different random seeds and dataset subsamples.




Figure 3. Projections of $\epsilon(d, l, w, n)$ onto two-dimensional planes for the CIFAR-10 ResNets, showing contours of constant error. For low enough densities, the contours have linear slopes on the logarithmic axes—depicted by a reference black-dotted line. The density/depth plane (left). The density/width plane (right).

the relationship between these inputs, we study the interdependence between density and depth or width by examining two-dimensional projections of the actual error $\epsilon(d, l, w, n)$ in Figure 3. These plots display contours of constant error as density and depth or width vary.

Consider the projection onto the plane of density and depth (Figure 3, left). The constant-error contours are linear except for in the densest networks, meaning each contour traces a power-law relationship between d and l . In other words, we can describe all combinations of densities and widths that produce error ϵ_v using $l^\phi d = v$, where v is a constant at which network error is ϵ_v and ϕ is the slope of the contour on the logarithmic axes. The contours of density and width also have this pattern (Figure 3, right), meaning we can describe a similar relationship $w^\psi d = v'$. Finally, we can combine these observations about depth and width into the expression $l^\phi w^\psi d = v''$.

We refer to the expression $l^\phi w^\psi d$ as the *error-preserving invariant*, and we denote it m^* . This invariant captures the observation that there exist many interchangeable combinations of depth, width, and density that achieve the same error and tells us which combinations do so. For example, networks of vastly different densities reach the same error if we vary l and w according to the invariant.

Functional form. The invariant allows us to convert the functional form $\hat{\epsilon}(\epsilon_{np}, d | l, w, n)$ for a specific l, w , and n from Section 3 into a joint functional form $\hat{\epsilon}(\epsilon_{np}, d, l, w, n)$ for all l, w , and n . Rewriting the definition of the invariant, $d = \frac{m^*}{l^\phi w^\psi}$. We can substitute this for d in the functional form from Section 3. Finally, by rewriting p as $\frac{p'}{l^\phi w^\psi}$ and canceling, we arrive at the following expression:

$$\begin{aligned} \hat{\epsilon}(\epsilon_{np}, d | l, w, n) &= \epsilon_{np} \left\| \frac{m^* - jp' \left(\frac{\epsilon_{np}^\dagger}{\epsilon_{np}} \right)^{\frac{1}{\gamma}}}{m^* - jp'} \right\|^\gamma \\ \epsilon_{np} \left\| \frac{l^\phi w^\psi d - jp' \left(\frac{\epsilon_{np}^\dagger}{\epsilon_{np}} \right)^{\frac{1}{\gamma}}}{l^\phi w^\psi d - jp'} \right\|^\gamma &= \hat{\epsilon}(\epsilon_{np}, d, l, w, n) \end{aligned} \quad (2)$$

which is the joint functional form $\hat{\epsilon}(\epsilon_{np}, d, l, w, n)$ of all four dimensions d, l, w , and n . Critically, for this to be a useful joint form, the free parameters ϵ^\dagger, p' , and γ must be constants shared across all possible values of d, l, w , and n . We will assume this is the case and directly quantify how well this assumption holds in the evaluation section below.

For qualitative intuition as to why this is a reasonable assumption, consider the relationship between m^* and the test error of pruned networks as we vary depth, width, and dataset size (Figure 4). Across all projections, the annotated ϵ^\dagger (error of the high-error plateau), γ (slope of the power-law region) and p' (value of m^* where the high-error plateau transitions to the power-law region) appear the same.

The preceding discussion addresses how we handle l, w , and d in our joint scaling law. We address dataset size n in Eq. 2 implicitly through the way that it affects ϵ_{np} , and we validate that this is a reasonable choice through the evaluation below. We retain the explicit form $\hat{\epsilon}(\dots, n)$ to stress that the lack of explicit dependency on n is non-trivial and was not known prior to our work.

Fitting. To fit $\hat{\epsilon}(\epsilon_{np}, d, l, w, n)$ to the actual data $\epsilon(d, l, w, n)$, we estimate values for the free parameters $\epsilon^\dagger, \gamma, p', \phi$ and ψ by minimizing the relative error $\delta \triangleq \frac{\hat{\epsilon}(\epsilon_{np}, d, l, w, n) - \epsilon(d, l, w, n)}{\epsilon(d, l, w, n)}$ using least squares regression. The fit is performed jointly over all configurations of d, l, w , and n , resulting in joint estimates of $\hat{\epsilon}^\dagger, \hat{\gamma}, \hat{p}, \hat{\phi}$, and $\hat{\psi}$. One can also perform a partial fit for a subset of dimensions (e.g., just d, l , and n) by omitting ϕ and/or ψ (see Appendix D).

Evaluating fit. In Figure 5, we plot the actual error $\epsilon(d, l, w, n)$ and the estimated error $\hat{\epsilon}(\epsilon_{np}, d, l, w, n)$ for the CIFAR-10 ResNets (all widths, depths and dataset sizes) and ImageNet ResNets (all widths and dataset sizes for depth 50). As in Section 3, our estimated error appears to closely follow the actual error. Deviations arise mainly at high densities where error decreases below ϵ_{np} and low densities approaching high error saturation.

We again quantify the fit of the estimated error using the mean μ and standard deviation σ of the relative deviation δ . The relative deviation on the joint scaling laws for the CIFAR-10 and ImageNet networks has a mean $\mu < 2\%$ and standard deviation of $\sigma < 6\%$.

To contextualize these results, Figure 5 (right) quantifies the variation in error we see over multiple replicates of the CIFAR-10 experiments due to using different random seeds. It plots the minimum, maximum, and mean errors across the three replicates we ran.⁷ The variation across trials has a standard deviation of $\sigma = 3.4\%$, sizeable relative to the estimation error of $\sigma = 5.8\%$ for the joint scaling law. This

⁷We only ran a single replicate of the ImageNet experiments due to the significant cost of collecting data.




Figure 4. Relationship between m^* and error when pruning CIFAR-10 ResNets and varying w (left, $l = 20, n = N$), l (center, $w = 1, n = N$), n (right, $l = 20, w = 1$). We annotate $\gamma, \epsilon^\uparrow$, and p' ; they qualitatively appear to take on similar values in all cases, an observation that we use to inform the design of our joint scaling law.




Figure 5. Estimated versus mean actual error for all configurations (d, w, n) for ImageNet (left) and (d, w, l, n) for CIFAR-10 (center). The variation in error when running the same experiment on CIFAR-10 three times with different random seeds (right).

indicates that a significant portion of our error may stem from measurement noise.

The functional form has just five parameters and obtains an accurate fit on 4,301 points on CIFAR-10 and 274 points on ImageNet, suggesting it is a good approximation. In Appendix E, we show that it achieves a similarly good fit for additional architectures and datasets. In Section 5, we show that, although we use a large number of points to develop and evaluate our functional form here, it is possible to get a good fit with far fewer points and the fit has low sensitivity to the choice of points.

5. Sensitivity of Fit to Number of Points

In Section 4, we showed that our scaling law was accurate when we fit it on all of the available data. Now that we possess the functional form and know that it can accurately model the behavior of IMP, we study the amount of data necessary to obtain a stable,⁸ accurate fit. This question is especially relevant when the functional form is applied to new settings—new networks, datasets, or pruning algorithms—and we must collect new data to do so. The functional form

⁸Stability is defined as a small change in output relative to a change in input. The requirement here is that a change in choice of points leads to a small expected change in estimation accuracy.

has only five parameters, suggesting that few experiments will be necessary to obtain an accurate fit.

Experiments. To evaluate the effect of the number of points on the stability and accuracy of the fit, we randomly sample varying numbers of points, fit the scaling law to those points, and evaluate the quality of the fit over all points. We sample these points in two ways.

Experiment 1. Randomly sample T networks (w, l, n, d) . This experiment evaluates the stability and accuracy of the fit when naively varying the number of points.

Experiment 2. Randomly sample T network configurations (w, l, n) and include all densities d for each configuration. This experiment captures the specific use case of IMP, where obtaining data at density d requires obtaining all densities $d' > d$. As such, we anticipate that data will be obtained by iteratively pruning a small number of configurations (w, l, n) to low density.

Results. We perform each experiment for many different values of T on the CIFAR-10 ResNets pruned with IMP. We repeat the experiment at each value of T 30 times with different samples of points and report the mean and standard deviation of μ and σ for the fit. Experiments 1 and 2 respectively appear in Figure 6 left and right. The shaded




Figure 6. The effect of the number of points used to fit our scaling law (on the CIFAR-10 ResNets) on μ and σ . Left: experiment 1 from Section 5 (random points w, l, d, n). Right: experiment 2 from Section 5 (random configurations w, l, n and all densities).

areas represent one standard deviation from the mean in each direction. On Experiment 1, when just 40 networks (w, l, d, n) are available, the standard deviation on both μ and σ is just one percentage point. On Experiment 2, when just 15 random configurations of (w, l, n) are available at all densities, we similarly achieve standard deviation below 1%. In both cases, as the number of networks increases, the standard deviation decreases further.

These results show that, now that our scaling law is known, it is possible to obtain an accurate (and stable) estimation using far less data than we used to evaluate the quality of the fit in Section 4. This implies that translating our scaling law to new settings will be far less data-intensive than developing and evaluating it in the first place. Moreover, the results in this section reflect a particularly naive way of selecting points: doing so randomly; we made no effort to ensure that, for example, the networks represented a diverse range of widths, depths, dataset sizes, and densities. By selecting these networks in a strategic way, it may be possible to further reduce the number of networks necessary to obtain a similarly accurate fit.

6. Discussion: Selecting a Functional Form

We have shown that our proposed functional form $\hat{\epsilon}(d, l, w, n)$ accurately approximates the error when pruning families of neural networks. In this section, we discuss some of the key criteria that led us to select this particular functional form. We intend this section to provide insight into our choices in the context of the broader design space and to highlight opportunities for further refinement.

Criterion 1: Transitions. In Section 3, we observe that, when pruning a neural network with IMP, error has a low-error plateau, a power-law region, and a high-error plateau. Between these regions are *transitions* where error varies smoothly from one region to the next. Matching the shape of these transitions was a key consideration for selecting our function family. To illustrate the importance of properly fit-

ting the transitions, Figure 6 shows two possible functional families for fitting the relationship between density and error for the CIFAR-10 ResNets. Actual error is in black, and the functional form from Section 3 is in blue. In red is the fit for a functional form adapted from the one that Rosenfeld et al. (2020) use to model the relationship between width and error. The difference between these functional families is the way they model transitions, and the one we choose in this paper better models the transitions in our setting. For further discussion of this comparison, see Appendix G.

Criterion 2: Few, interpretable parameters. Selecting a functional form is not merely a curve-fitting exercise. We seek the underlying structure that governs the relationships between d, l, w, n , and error in a manner akin to a law of physics. As such, our functional form should have a small number of parameters that are *interpretable*. In our functional form (Eq. 2), each parameter has a clear meaning. The parameters $\epsilon^\dagger, p',$ and γ control the high-error plateau, the transition to the power-law region, and the slope of the power-law region. ϕ and ψ control the interchangeability of width and depth with density. We approximate error over multiple orders of magnitude and 4,301 configurations of ResNet-20 on CIFAR-10 with just five parameters, indicating we have distilled key information about the behavior of pruning into our functional form.

Sources of systemic error and limitations. By seeking to minimize the number of parameters in our functional form, we leave some phenomena unmodeled. In particular, there are two phenomena we have chosen *not* to model that introduce systemic error. First, the low-error plateau is not a plateau. Error often improves slightly at high densities before returning to ϵ_{np} during the transition to the power-law region. Our model treats the region as flat and treats error as monotonically increasing as density decreases. This source of error accounts for a bias of $\sim 1\%$ relative error in our estimation (Appendix H). Second, we model both transitions (between the power-law region and each plateau) with a single shape and the same transition rate. If we treated them separately and used higher-order terms in the rational form, we could potentially reduce some of the residual error in our estimation at the cost of additional complexity.

7. Implications and Conclusions

Our main contribution is a functional form $\hat{\epsilon}(\epsilon_{np}, d, l, w, n)$ that accurately predicts the error when pruning members of a network family using IMP. There are several broader implications of our ability to characterize pruning in this way. The mere existence of this functional form means there is indeed structure to the way pruning affects error. Although prior work (Cai et al., 2019) has implicitly relied on the existence of structure for a different pruning method, we are the first to explicitly describe such structure. This




Figure 7. Estimated error from our scaling law (blue) and the scaling law adapted from Rosenfeld et al. (2020) (red) for CIFAR-10 ResNet as w varies. Actual error is in black. Our scaling law better captures transitions between the low-error plateau and the power-law region.

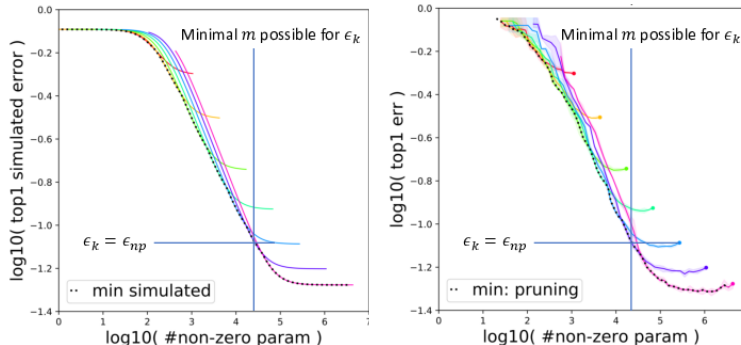


Figure 8. Estimated error as width varies for the CIFAR-10 ResNets (left). Actual error as width varies for the CIFAR-10 ResNets (right). The dotted black line is the minimal number of parameters necessary to reach each error ϵ_k among all of the pruned networks. Reaching this point requires starting with a particular lower-error network (purple) and pruning until error increases to ϵ_k . Starting too large (pink) will miss this point.

functional form enables a framework in which we can reason conceptually and analytically about pruning. In doing so, we can make new observations about pruning that are non-obvious or costly to exhaustively demonstrate empirically. For example, recall our motivating question:

Given a family of neural networks, which should we prune (and by how much) to obtain the network with the smallest parameter-count such that its error does not exceed some threshold ϵ_k ?

This is an increasingly common question as the community begins to think about pruning as a way of seeking the optimal pruned member of a *family* of networks rather than as a technique applied to a specific network in isolation. At its heart, the question is really an optimization problem: find the configuration of d , l , and w that minimizes parameter-count m subject to an error constraint: $\operatorname{argmin}_{w,l,d} m$ s.t. $\hat{\epsilon} = \epsilon_k$. For ResNets, the parameter-count m is proportional to (dlw^2) .⁹ Hence, this yields the following optimization problem:

$$l, w, d = \operatorname{argmin}_{l,w,d} l w^2 d \quad \text{s.t.} \\ \epsilon_{np} \left\| \frac{l^\phi w^\psi d - j p' (\epsilon^\dagger / \epsilon_{np})^{1/\gamma}}{l^\phi w^\psi d - j p'} \right\|^\gamma = \epsilon_k$$

This optimization problem is solvable directly without running any further experiments.

Studying this optimization problem reveals a useful insight about—in this case—the CIFAR-10 ResNets. In the pruning literature, it is typical to report the minimum density where the pruned network matches the error $\epsilon_{np}(l, w)$ of the

⁹Increasing the depth linearly increases the number of parameters, but increasing the width quadratically increases the number of convolutional filters and thereby the parameter-count.

unpruned network (Han et al., 2015). However, our scaling law suggests this is not the smallest model to achieve error $\epsilon_{np}(l, w)$. Instead, it is better to train a larger network with depth l' and width w' and prune until error reaches $\epsilon_{np}(l, w)$, despite the fact that error will be higher than $\epsilon_{np}(l', w')$. This analytic result parallels and extends the findings of Li et al. (2020) on NLP tasks. However, unlike Li et al., our scaling law suggests starting too large is detrimental for the CIFAR-10 ResNets, leading to a higher parameter-count at error ϵ_k .

Figure 8 (left) illustrates this behavior concretely: it shows the error predicted by our scaling law for CIFAR-10 ResNets with varying widths. The dotted black line shows the minimal parameter-count at which we predict it is possible to achieve each error. Importantly, none of the low-error plateaus intersect this black dotted line, meaning a model cannot be minimal until it has been pruned to the point where it increases in error. This occurs because the transitions of our functional form are gradual. On the other hand, if we start with a model that is too large, it will no longer be on the black line when it has been pruned to the point where its error reaches $\epsilon_{np}(l, w)$; this behavior occurs because error decreases as a function of the invariant m^* rather than the parameter-count m and because $m \not\propto m^*$. In Figure 8 (right), we plot the same information from the actual CIFAR-10 data and see the same phenomena occur in practice. The difference between the estimated and actual optimal parameter count is no more than 25%.

Looking ahead, there are many directions for future work. Further studying sources of systematic error (transition shape and error improvements on the low-error plateau) is a promising avenue for making it possible to *extrapolate* from small-scale settings to large-scale settings (see Appendix F for a forward-looking discussion). Furthermore, while we focus on CIFAR-10 and ImageNet ResNets in the

main body, it is important to understand the generality of our functional form for other networks and tasks (see Appendix E). Finally, now that we have described the structure of the error of IMP-pruned networks, it will be valuable to study the nature of scaling laws that capture the behavior of the plethora of other pruning methods that achieve different tradeoffs between parameter-count and error.

Acknowledgements

Jonathan Rosenfeld was funded by the Efficient NN Model Scaling grant from Analog Devices Incorporated. This work was partially supported by NSF grants CCF-1563880 and IIS-1607189

We gratefully acknowledge the support of Google, which provided us with the TPU resources necessary to conduct our experiments through the TPU Research Cloud.

We gratefully acknowledge the support of IBM, which provided us with access to the GPU resources necessary to conduct our experiments on CIFAR-10 through the MIT-IBM Watson AI Lab.

This work was supported in part by cloud credits from the MIT Quest for Intelligence.

This work was supported in part by the Office of Naval Research (ONR N00014-17-1-2699).

This work was supported in part by DARPA award #HR001118C0059.

We thank Yonatan Belinkov and Roy Shilkrot for their constructive comments.

References

- Blalock, D., Ortiz, J. J. G., Frankle, J., and Gutttag, J. What is the state of neural network pruning? *Conference on Machine Learning and Systems*, 2020.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners, 2020.
- Cai, H., Gan, C., and Han, S. Once for all: Train one network and specialize it for efficient deployment. *arXiv preprint arXiv:1908.09791*, 2019.
- Frankle, J., Dziugaite, G. K., Roy, D. M., and Carbin, M. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, 2020.
- Gale, T., Elsen, E., and Hooker, S. The state of sparsity in deep neural networks. *arXiv preprint arXiv:1902.09574*, 2019.
- Gordon, M. A., Duh, K., and Andrews, N. Compressing bert: Studying the effects of weight pruning on transfer learning. *arXiv preprint arXiv:2002.08307*, 2020.
- Han, S., Pool, J., Tran, J., and Dally, W. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, pp. 1135–1143, 2015.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hestness, J., Narang, S., Ardalani, N., Diamos, G., Jun, H., Kianinejad, H., Patwary, M., Ali, M., Yang, Y., and Zhou, Y. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409*, 2017.
- Janowsky, S. A. Pruning versus clipping in neural networks. *Phys. Rev. A*, 39:6600–6603, Jun 1989. doi: 10.1103/PhysRevA.39.6600. URL <https://link.aps.org/doi/10.1103/PhysRevA.39.6600>.
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and Amodei, D. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- LeCun, Y., Denker, J. S., and Solla, S. A. Optimal brain damage. In *Advances in neural information processing systems*, pp. 598–605, 1990.
- Li, Z., Wallace, E., Shen, S., Lin, K., Keutzer, K., Klein, D., and Gonzalez, J. E. Train large, then compress: Rethinking model size for efficient training and inference of transformers. *arXiv preprint arXiv:2002.11794*, 2020.
- Reed, R. Pruning algorithms—a survey. *IEEE transactions on Neural Networks*, 4(5):740–747, 1993.
- Renda, A., Frankle, J., and Carbin, M. Comparing rewinding and fine-tuning in neural network pruning. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=S1gSj0NKvB>.
- Rosenfeld, J. S., Rosenfeld, A., Belinkov, Y., and Shavit, N. A constructive prediction of the generalization error across scales. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=ryenvpEKDr>.
- Tan, M. and Le, Q. V. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019.