

A HITCHHIKER’S GUIDE TO SCALING LAW ESTIMATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Scaling laws predict the loss of a target machine learning model by extrapolating from easier-to-train models with fewer parameters or smaller training sets. This provides an efficient way for practitioners and researchers alike to compare pre-training decisions involving optimizers, datasets, and model architectures. Despite the widespread use of scaling laws to model the dynamics of language model training, there has been little work on understanding how to best estimate and interpret them. We collect (and release) a large-scale dataset containing losses and downstream evaluations for 485 previously published pretrained models. We use these to estimate more than 1000 scaling laws, then derive a set of best practices for estimating scaling laws in new model families. We find that fitting scaling laws to intermediate checkpoints of training runs (and not just their final losses) substantially improves accuracy, and that—all else equal—estimates of performance are generally most accurate when derived from other models of similar sizes. However, because there is a significant degree of variability across model seeds, training multiple small models is sometimes more useful than training a single large one. Moreover, while different model families differ in scaling behavior, they are often similar enough that a target model’s behavior can be predicted from a single model with the same architecture, along with scaling parameter estimates derived from other model families.

1 INTRODUCTION

Substantial effort and cost are required to train even a single large language model (LLM).¹ There is thus an acute need for efficient decision-making aids that can evaluate the effectiveness of proposed changes to language models’ architecture or training data without full-scale training runs. While there is a large body of work that motivates or evaluates these changes using small models (Warstadt et al., 2023; Hillier et al., 2024), synthetic tasks (Akyürek et al., 2024; Wortsman et al., 2023) or theory (Jelassi et al., 2024), one of the most important tools for current practitioners is the estimation of **scaling laws** for LLMs (Ivgi et al., 2022; Dubey et al., 2024).

A scaling law extrapolates the performance of a target model from the performance of a set of models with fewer parameters or smaller training sets. Typically, this extrapolation requires models to belong to the same **model family**, differing only in parameter count and training set size, but using the same architecture and training distribution. A high-quality scaling law accurately predicts the target model’s test performance (Rosenfeld et al.; Kaplan et al., 2020; Hoffmann et al., 2022).

Most past work describing and characterizing scaling laws has begun by exhaustively training models in a family across a full range of dataset sizes and parameter counts. One question that has received comparatively little attention is how, when training a new LLM, a practitioner with limited computational resources should choose *which* small-scale models to train in order to best estimate a target model’s final performance. This paper offers a practical guide to when, and how, to use small models to efficiently obtain meaningful predictions about large models’ behavior—maximizing prediction reliability while minimizing the budget for preliminary experimentation, which necessarily involves tradeoffs between the number of preliminary models trained, the size of the largest preliminary model, and size of the dataset used to train it.

¹Code, data and full numbers are found in our repository

We begin by collecting diverse model data to perform a large-scale meta-analysis of scaling laws (§3). Usually, scaling law research relies on a single collection of closely related models, or alters only a minimal aspect of pretraining (e.g. data size; Muennighoff et al., 2024). Instead, we gather data from as diverse a set of scaled families as possible, to allow this and future meta-analysis of scaling laws that generalize across architectures, datasets and settings.

The rest of the paper uses this data to analyze a number of key questions around scaling law estimation:

1. **How reliably may we expect scaling laws to extrapolate?** Variation between random parameter initializations can produce changes of up to 4% in loss. Most published improvements in pretraining procedures, when performing minimal controlled experiments, report loss changes between 4% and 50% (§4).
2. **How much does the shape of scaling laws vary across model families?** Different model families have scaling laws with a different functional dependence on model size (§5). However, transformer LMs are similar enough that, with a single model from a target family and a scaling law from a different model family, it is sometimes possible to accurately estimate target model performance (§5.1).
3. **Must scaling laws be estimated only from fully trained models?** Even though optimization procedures are typically sensitive to the full size of a training run, estimating scaling laws from intermediate training checkpoints greatly improves scaling law fit (§6). It is generally possible to estimate a model’s final loss beginning roughly $1/3$ of the way through training.
4. **How large must models be to produce reliable scaling estimates?** All else equal, experimenting with large models is typically more useful than with small models (§7), but may be outweighed by the benefits of reduced variance from training more, smaller models (§8).
5. Taken together, **cost-effective estimation** of a scaling law should consider the **number** of models, the **size** of the models, and the number of **training tokens** for each model. We highlight those size, tokens and number of models effects in Fig. 2.

Our experiments also provide insight into the functional form of scaling laws themselves, suggesting that they may have fewer degrees of freedom (§9) than typically assumed. We conclude with discussion of other work on scaling law estimation that may be of interest to practitioners §10.

2 DEFINING A SCALING LAW

A scaling law estimates the loss of a costly model by training cheaper ones (see Fig. 1) which share a pretraining procedure and differ by some hyperparameters, typically model size (#params) and number of tokens seen during training (#toks). A scaling law is a function that predicts a target model’s loss on held-out data when setting the value of one hyperparameter (Kaplan et al., 2020) or both (Rosenfeld et al.; Hoffmann et al., 2022). Comparing laws’ predictions about different pretraining choices (e.g. data Ge et al., 2024) allows informed decisions about which large-scale model to train. A scaling law also enables finding the optimal choice of hyperparameters under computational constraints on pretraining (Hoffmann et al., 2022) or inference (Touvron et al., 2023; Sardana et al.).

Formally, we will call a **model** f any single concrete neural language model with a specific set of parameters. Different seeds, or even different checkpoints from the same training run, correspond to different models. We define a **scaled model family** f as a set of models, with each $f \in F$ differing only in size $\#params(f)$ and number of tokens $\#toks(f)$.

There are two specific subsets of scaled model families that will be useful in our experiments. First, the **maximal parameter family** $\max_{\#params}(F)$ contains only models in F with the largest number of parameters. Formally, define $m = \max_{f \in F} \#params(f)$; then $\max_{\#params}(F) =$

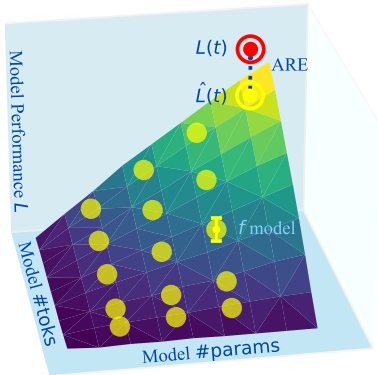


Figure 1: Illustration of a scaled family, an estimated scaling law, and its prediction error for a target model.

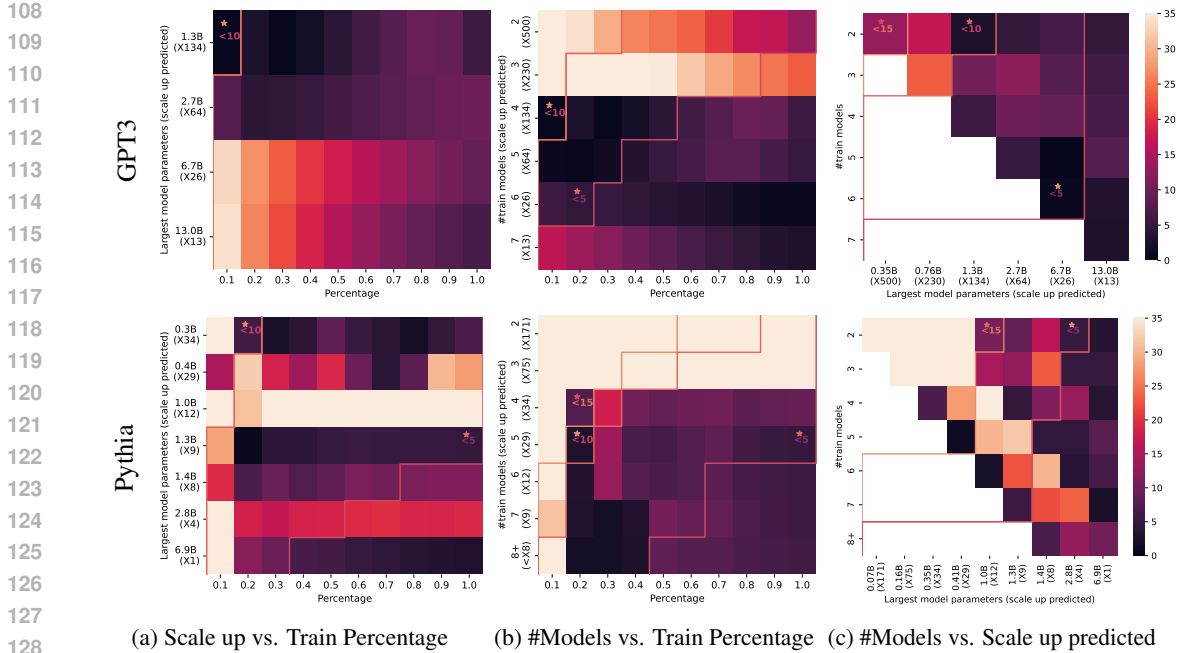


Figure 2: The effects of three variables on scaling law accuracy. Each cell corresponds to a single scaling law estimated from a set F_{train} of model checkpoints, with the color denoting that scaling law’s error when predicting the largest model in a family. Each column shows a subset of the three axes along which these training sets differ: (1) the number of tokens used to train each LM in F_{train} (expressed as a fraction of the full training corpus), (2) the number of distinct models trained; and (3) the size of the largest model trained (expressed as a scale-up factor—the ratio between the target model and the largest model in F_{train}). In (a), all laws are estimated from four models. In (c) all laws use the full corpus. Orange lines show iso-FLOP contours (sets of scaling laws whose training sets require the same computational cost to produce). \star represent the most efficient ways to obtain 15%, 10% and 5% ARE. One of the most immediate conclusions from these plots is that scaling law estimation is quite noisy—the inclusion of a single badly-behaved model in the estimation procedure can produce large errors, and in small model families error does not reliably decrease with additional computation. However—because of noise—it is often preferable to extrapolate from a large number of small, partially trained models rather than a small number of large models.

$\{f \in F : \#\text{params}(f) = m\}$. This family will generally contain the **target** model(s) whose behavior we wish to predict $t \in F_{\text{target}}$. Second, the **q-maximal token family** $\max_{\#\text{toks}}(F, q)$ contains all models in f trained on at least a q -sized fraction of the training set. Formally, define $t = q \cdot (\max_{f \in F} \#\text{toks}(f))$; then $\max_{\#\text{toks}}(F, q) = \{f \in F : \#\text{toks}(f) \geq t\}$. Note that this definition does not distinguish between partially trained models on one hand, and models trained to convergence on a subset of the largest training set used in a family on the other. Throughout this paper, we will not in general distinguish between these two types of models, a decision evaluated in Section 6. Indeed, except where noted, $\max_{\#\text{toks}}(F, q)$ should be thought of as containing the checkpoints from the last $q\%$ of a training run.

A **scaling law** $\hat{L}(f | F)$ estimates the performance of a new model f given a model family F . (We will simply write $\hat{L}(f)$ when the family is clear from the context.) All experiments in this paper use the common functional form from the literature (Hoffmann et al., 2022):

$$\hat{L}(f) := E + \frac{A}{\#\text{params}(f)^\alpha} + \frac{B}{\#\text{toks}(f)^\beta}. \tag{1}$$

Here E captures the scaled family’s general performance; A, α and B, β describe the scaling effect of $\#\text{params}$ and $\#\text{toks}$ respectively.² These parameters are estimated by first collecting a set of

²We believe many of the findings in this paper apply to other functional forms that have been proposed for scaling laws (Caballero et al.), and even suggest new parameterizations as described in §9.

162 **training models** F_{train} , then minimizing the reconstruction error

$$163 \arg \min_{E, A, \alpha, B, \beta} \sum_{f \in F_{\text{train}}} (\hat{L}(f) - L(f))^2$$

164 where $L(f)$ denotes the empirical negative log-likelihood of some held-out data under the model f .

165 In this sense, a scaling law is an ordinary parametric machine learning model, and we may ask many
166 of the same questions about \hat{L} that we ordinarily ask about LLMs —what training data (F_{train}) should
167 we collect? How do we estimate accuracy? We seek to provide empirical answers to these questions,
168 for which we first require data.

169 3 DATA FOR 1000+ SCALING LAWS AND MORE

170 As part of this work, we have collected and released the largest-scale public dataset describing scaling
171 behavior across model families. This dataset aggregates information from a large number of LLM
172 training efforts that have released information about the behavior of multiple models of different
173 sizes or scales. While experiments in this paper focus on scaling laws that measure loss, the dataset
174 also includes information about model performance on downstream evaluation benchmarks where
175 available. We have focused on language models where the largest one is more than 3B parameters
176 and where data was shared publicly or in private correspondence. Our repository accepts further
177 contributions and requests for additions. In addition to those, we have manually extracted some data
178 from papers that did not release models but reported losses in figures.

179 3.1 DATA SOURCES

180 For each model in this dataset, we report any downstream evaluation and loss that was measured
181 during training, as well as calculated #toks for each, links to matching checkpoints when available,
182 links to data sources, and information about computational cost (in FLOPs) and number of training
183 epochs (i.e. passes over the training set). Each model is identified by a unique name, a type (e.g.
184 llama), #toks, #params, architecture type (e.g. encoder-decoder), and seed.

185 Models in this dataset include Pythia (Biderman et al., 2023, which provides the largest set of
186 models and variations in a family), OPT (Zhang et al., 2022, collected thanks to Xia et al., 2023;
187 Biderman et al., 2023), OLMO (Groeneveld et al., 2024), Amber (Liu et al., 2023), K2 (Team,
188 2024), Mamba (Liu et al., 2023) RedPajamas³ ModuleFormer mixture of experts (Shen et al., 2023),
189 overtrained models (Gadre et al., 2024), Mamba, Llama and hybrid architecture variations from
190 Poli et al. (2024), transformer architectures (Alabdulmohsin et al., 2022), Bloom (Le Scao et al.,
191 2023), T5-Pile (Sutawika et al., 2024), Pandey (2024) models, GPT-family models with different data
192 regimes (Muennighoff et al., 2024), Gopher (Hoffmann et al., 2022) and GPT3 (Brown et al., 2020).

193 The data consists of 1.9M steps of training evaluated on loss or perplexity, usually on multiple data
194 sources belonging to 485 unique pretrained models, and more than 40 scaled families. To combine
195 the data together, each loss or score was kept in its original form and an aggregating function to
196 normalize the scores to similar scales (perplexity and loss have a log relation) and average was
197 supplied. The data, saved as a compressed table with all the metadata and scores together is hence
198 large but comprehensive.

199 We hope this will provide a useful resource for the community and plan to extend it further as models
200 get released and their training dynamics are shared. We see such a resource as a facilitator to more
201 research on model development (e.g. A/B testing), scaling laws, downstream scaling laws (Gadre
202 et al., 2024; Ruan et al., 2024; Owen, 2024; Isik et al., 2024), training dynamics (Choshen et al.,
203 2022) and more.

204 3.2 SCALING LAW ESTIMATION

205 In the rest of the paper, we present findings from estimating hundreds of scaling laws as follows:

³<https://www.together.ai/blog/redpajama-models-v1>

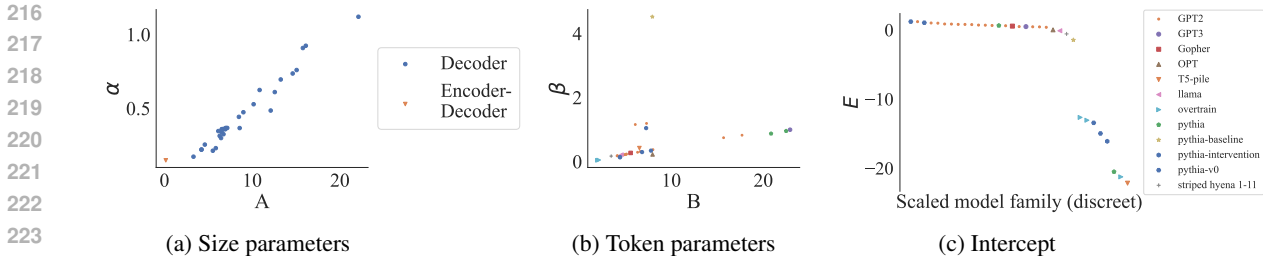


Figure 3: Parameters differ between scaled model families. Surprisingly, however, the pairs of parameters controlling the influence of model and training set size have similar ratios. The legend shows model architecture (left), scaling families (center) and per-family intercept (right).

Fitting For each model family F , we identify the maximal parameter family $F_{\max} = \max_{\#\text{params}}(F)$, and estimate a scaling law \hat{L} using the remaining models $F_{\text{train}} = F \setminus F_{\max}$. Estimation of scaling law parameters uses the `curve_fit` function in `scikit-learn` (Pedregosa et al., 2011). We additionally experimented with an L-BFGS-based solver but found it to be less stable. We only estimate scaling laws for model families that contain at least three models.

Evaluation To evaluate estimated scaling laws reliably, we need to account for loss fluctuations during large-scale model training. Thus, we test against a few checkpoints near the end of training: we choose as target models F_{target} the 30%-maximal token family from the set F_{\max} defined in the previous paragraph—that is, we take $F_{\text{target}} = \max_{\#\text{toks}}(F_{\max}, 0.3)$. We then report the mean **absolute relative error (ARE)** $\mathbb{E}_{f \in F_{\text{target}}} |L(f) - \hat{L}(f | F_{\text{train}})| / L(f)$ between the empirical loss L and the loss \hat{L} predicted by the scaling law. We reran our experiments with huber loss see §E.

4 HOW WELL CAN I EXPECT A SCALING LAW TO PREDICT?

4% is the best ARE typically obtained; ARE up to 20% can still distinguish between many modeling choices.

To establish how accurate a scaling law must be to be *useful* to practitioners, we first assess what changes in model accuracy have been considered meaningful in past work. We have surveyed experiments in the literature where an A/B test was performed, i.e., two models were trained similarly, manipulating one attribute to see how it affects scores. Empirically, we found no widely adopted modeling changes that were motivated with less than a 4% relative difference between models. Additionally, reported variance across random restarts of the same model architecture reaches up to 3.5% (c.f., §8; Sellam et al., 2021). We take this to mean that this is approximately the minimal effect-size experimenters care about and possibly the minimal effect one can reliably measure. Accordingly, this bounds the best goodness of fit we should expect or require of scaling laws.

To offer several concrete points of comparison: Pythia 6.9B models fixed inconsistencies in their code and hence have two versions (c.f. App. B; Biderman et al., 2023) which differ in loss by 40%. They also provide data deduplication A/B test that had a minor effect on the loss of about 5%. Gadre et al. (2024) tested the effect of training 400M parameter models for different $\#\text{toks}$. The most similar (double the training tokens) has approximately 4% change and can reach a 50% loss difference with 30 times more training. Training on a constant $\#\text{toks}$ but repeating the same data resulted in almost no changes for up to 4 repetitions (epochs), and later in about 8%, 50% on 14.44 repetitions of the data (Muennighoff et al., 2024). Instead of varying the amount of data or epochs, Ge et al. (2024) found that training on a different kind of data incurred ARE of approximately 10% and different data mixes led to 6% changes or less.

270 5 WHEN I TRAIN A NEW MODEL, DO I EVEN NEED A NEW SCALING LAW? 271

272 *Different model families exhibit different scaling behavior, but performance can sometimes be estimated*
273 *using a single model in a new family.*
274

275 Scaling laws relate performance to scalar training parameters like model or dataset size. For discrete
276 decisions (whether the choice of nonlinearity or data preprocessing scheme), it is not immediately
277 obvious how to pool information across models that differ in these traits (see Ruan et al., 2024;
278 Maia Polo et al., 2024, for concurrent work that performs this pooling based on downstream task
279 behavior). Clearly, different pretrained models with the same #params and #toks still show
280 different loss, so these differences can be consequential. But how do discrete choices of architecture,
281 training procedure, or dataset, affect the form of scaling laws?

282 One way to answer this question is to look at the parameter estimates for scaling law parameters E ,
283 α , A , β and B differ across model families. These results are shown in Fig. 3, where it can be seen
284 that there are often dramatic differences in all five parameters across families. In this sense, even
285 the rate at which additional data or parameters improve model performance depend on underlying
286 architectural details, suggesting that understanding the behavior of a new model family may require a
287 new scaling law.

288 But another way to answer this question is to ask how reliably we can predict final model accuracy
289 when borrowing (or pooling) some parameters of scaling laws between families—even if these
290 result in poor parameter estimates, they may predict large-scale model behavior within the range of
291 meaningful differences identified in Section 4. To do so, we set the #params scaling parameters
292 (A , α) to fixed values reported in past work, and estimate remaining parameters for individual model
293 families. We take the variable values found by Muennighoff et al. (2024) (see Besiroglu et al., 2024;
294 Porian et al., 2024 for a discussion of estimates from earlier work including Hoffmann et al., 2022).
295 We find (see Fig. 6 in App. A) that in some cases only a single training run in a new model family is
296 necessary to obtain accurate scaling law predictions. In the OLMO family, for example, we obtain
297 less than 1% error estimating the accuracy of a 7B model from a collection of 1B model checkpoints.
298 We find that predictions generalize, and a constant #params scaling factor is enough for most models
299 (except the encoder-decoder T5-Pile). However, error rates are larger than in the source family, and
300 predictions for larger models are worse (most conspicuous in OPT’s error of 37%, 25% and 15%
301 when extrapolating from 8.7B, 13B and 30B to 175B).

302 5.1 CAN I JUST TRAIN THE TARGET MODEL A BIT INSTEAD OF MANY SMALL MODELS? 303

304 *Yes, but obtaining reliable estimates in this way requires up to 30% of the full training run.*
305

306 The above results (last row of Fig. 6 in App. A) also suggest the possibility of predicting losses
307 not with just smaller models, but with partially trained versions of the target model itself. When
308 predicting inside the same #params family—that is, estimating $\hat{L}(f | F_{\text{target}} \setminus \{f\})$ —the #params
309 term in Eq. (1) is constant, and extrapolation is only required for #toks. As seen in the figures, this
310 form of estimation is informative if permitted by computational constraints. Beyond the immediate
311 usefulness of this approach, it is a promising avenue for future research. Better adjusting the scaling
312 laws for predicting through training might improve this efficiency.

314 5.2 ARE EVEN SIMPLER BASELINES ENOUGH? 315

316 *Some extrapolation is necessary: scaling laws can produce accurate estimates even when the target model*
317 *vastly outperforms any training model.*

318 To provide another form of comparison for the predicted scaling laws, we compute two baselines.
319 Both baselines adopt a pessimistic evaluation assuming that the target model is no better than the best
320 model in the small model family used to estimate a scaling law. Specifically, the baselines are the
321 *best performance* $\hat{L}(\emptyset | F_{\text{train}}) = \min_{f \in F_{\text{train}}} L(f)$ and the performance of the *most-trained model*,
322 consuming the most compute for training, i.e. $\hat{L}(\emptyset | F_{\text{train}}) = \arg \max_{f \in F_{\text{train}}} \#params(f) \times$
323 $\#toks(f)$. Those baselines might be the best one can expect without fitting a law to scaling.




Figure 4: The effect of fitting on more of the training trajectory. Each cell represents the absolute relative error estimating scaling laws from a given number of models (vertical axis) trained on a given subset of the *final* checkpoints from a training run (so scaling laws on the left are estimated using all checkpoint, and laws on the right are estimated using only the final 10% of checkpoints). White cells failed to fit. As long as the first $\approx 10\%$ of checkpoints are discarded, final loss can often be predicted accurately.

We find (See App. 5.2) that out of the two, the *best performance* baseline is closer to $L(F_{\text{target}})$, which is to be expected, as the target model performance is better than any other model in F and this is the better of the two. In both cases, even with the full F , the baselines suffer more than 15% error, mostly above 10%, almost never get below 5%, and 18% ARE on average across all scaled families we study.

6 I HAVE SOME DATA, WHAT PORTIONS SHOULD I USE?

Estimate scaling laws from intermediate checkpoints, not just fully trained models!

Most past work on scaling behavior of language models (e.g., Gadre et al., 2024; Muennighoff et al., 2024) has trained a *separate* model for each value of $\#\text{toks}$ studied. This is based on the assumption that changes in the learning rate schedule, which depend on the size of the full dataset that will be used for training, render losses from intermediate checkpoints uninformative.

However, some recent work has demonstrated the effectiveness of learning schedules that do not require prior access to the size of the training set (Hu et al., 2024), and some work has questioned whether careful choice of the learning rate decay is even necessary for reliable scaling laws (Porian et al., 2024). Together, these findings motivate revisiting the assumption that only a single useful datapoint may be obtained from each training run. In the final portion of §5.1, we observed the value of intermediate checkpoints when only a single $\#\text{params}$ family is used to fit a scaling law. Mathematically, there has to be some difference between models trained on the same number of tokens and sizes depending on the scheduler, it is unknown however, how big is this difference. We now test whether this finding extends to larger families—i.e. whether including intermediate checkpoints from all models in a model family reduces ARE.

Results are shown in Fig. 4, which plots ARE for scaling laws estimated from data subsets of the form $\max_{\#\text{toks}}(F, q)$ for varying q . We find that including full training curves in scaling law estimation can predict losses well. In fact, relying merely on the end of training produces significantly worse performance across the board. Our remaining experiments thus fit scaling laws using all these intermediate checkpoints, and not final performance alone.

6.1 SHOULD I USE ALL INTERMEDIATE CHECKPOINTS?

Almost all, but drop checkpoints from the beginning of training.

In Fig. 4, we plot the ARE for different q -maximal token families serving as F , i.e., when fitting only with the end of training runs. There is not a clear trend indicating whether we should use all

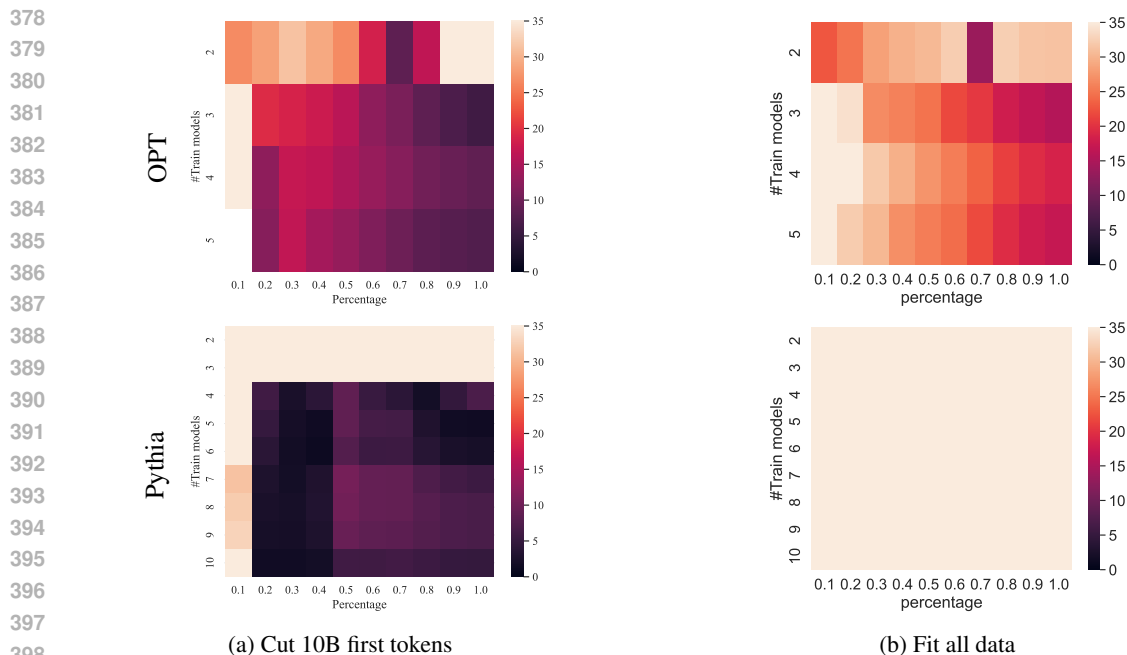


Figure 5: The effect of fitting with all the training losses and without the beginning 10B tokens seen. Each cell represents the absolute relative error when estimating a scaling law from a given number of models (vertical axis) trained on a given subset of checkpoints from the beginning of training (horizontal axis).

data (as might be suggested by GPT-3 results alone) or only some of it. But it is rarely the case that best estimates are obtained from the end of training alone.

There is, however, a distinctly uninformative phase at the beginning of training, as can be seen in the loss curves (App. B) and noted in the literature (e.g., Chen et al.). We observe that this period is more likely to contain significant spikes or an increase in loss (worse performance) despite additional training. We hence hypothesize this part should always be removed from the scaling law.

Indeed, our experiments depicted in Fig. 5 compare scaling law AREs with and without including models trained on less than 10B tokens in F . Evidently, the very beginning of training (often not even reported in logs and graphs) is sometimes harmful to the prediction and is perhaps more noisy. Specifically, we run the same experiments with and without ignoring the first 10B tokens seen. We find that for some models (e.g., OPT and Pythia) the ARE exceeds 15% even when using the whole data, but drops to 4-10% when ignoring those tokens. In preliminary experiments, we found that cutting fewer tokens gave noisier results, and cutting more had a negligible effect.

7 HOW BIG A MODEL SHOULD I TRAIN?

Larger models are better, but not necessary. Mainly, beware of specific models that might give noisy results.

In Fig. 2 we compare scaling laws when controlling the amount, percentage, or size of the models (2 at a time). We find that choosing models closer in #params to the target model is generally effective (e.g., Fig. 2a, 2c), but the effect is neither strong nor monotonic. For example, in all cases fitting on all F provides one of the lowest ARE. However, in GPT, Gopher and OPT, predicting with the smallest 4 models available is already enough to achieve less than 10% error. In Pythia, the smallest models are not predictive but the rest of the models provide a similar fit. While relying on a larger model is beneficial, predicting many scales up (e.g., the behavior of a $34\times$ larger model in Pythia) is still reliable, especially if accounting for other factors we discuss next.

In fact, training additional, larger models before fitting a scaling law may sometimes decrease accuracy due to increased variance in large model performance—see, for example, Pythia 2.8B in Fig. 1. Unfortunately, it is difficult to identify whether a seed is exceptionally high or low-performing without additional information. For example, cross-validation on F fails to detect it (see App. D).

Instead, this instability can be addressed by accounting for seed variability. A wasteful way to do so would be to train every model several times. A better alternative is to diversify and train each model on as differing hyperparameters (here, seed, $\#params$, $\#toks$) as possible and to maximize the information gained (a common practice in efficiency-coverage scenarios, e.g., Perlitz et al., 2024). Hence, we suggest training more models of differing sizes each accounting for both size and seed changes, rather than training multiple seeds. We further discuss the effects of number of models ($|F|$) in §8.

Selection of $\#params$ values to optimize statistical and computational efficiency is a problem we leave for future work. Given the choice of the largest model and the number of models, it is unclear how to space the model sizes, whether linearly, log-scale, or otherwise.

8 HOW MANY MODELS ARE NEEDED FOR RELIABLE PREDICTIONS?

5 models is a safe bet, more would improve the results' robustness. These models can be small.

We have seen that predicting with larger models and hence extrapolating less yields better results. However, given compute constraints (and additional hardware constraints like memory), practitioners may generally wish to use smaller models when possible. Consider for example Fig. 2b where we compare fitting on 4 models but vary their size. We find that *more* models reduce ARE even without being *bigger* models. As discussed in §7, adding a larger model to a current scaled family serves two goals, it increases the proximity to the predicted model, as well as increases the number of models seen.

We separate the contribution of size and number of models effect. In Fig. 2c, we predict with the largest model being held constant and add (at minimal cost) smaller models. We see again that larger models do benefit predictions. For example, the small models part (left) of the graph indicates large errors (bright). However, we also see again the unwanted effects a single model may have on the overall prediction. Consider for example the figure's diagonal in Pythia. Cells in a diagonal share a group of models and each row adds another one to F . Evidently this specific group hurts results, even when larger models are added to F . With enough models (bottom of diagonal), the negative decreases. Switching the model (next column) also removes the negative effect. Moreover, across all rows the tendency is never monotonic, implying larger models do not ensure better predictions.

But in general, we see that increasing the number of models tends to improve prediction. For example, in GPT3 the best predictions are with many models. Perhaps intuitively, adding a larger model and improving both $\#params$ and number of models aspects improves quite consistently (Fig. 2b and diagonals of Fig. 2c).

9 WHAT PARAMETERS DO I ACTUALLY NEED TO ESTIMATE?

Scaling laws might have fewer degrees of freedom than described in the literature.

Assuming we do not try to account for aspects other than $\#toks$ and $\#params$ (see §10), one might wonder if some of the observed errors come from model misspecification—an incorrect functional form for \hat{L} , which (with a small number of exceptions including Caballero et al.) has generally gone uncontested since it was first proposed (Rosenfeld et al.; Hoffmann et al., 2022). Here we specifically evaluate whether scaling laws empirically exhibit fewer degrees of freedom than has been proposed. First, we compute the principal components of the 5 learned parameters and find that 3 components explain 99.49% of the variance between the 5 parameters. Inspection reveals that two of these components tightly couple the pairs of parameters dealing with the same training parameter ($\#params$ and $\#toks$). Plotting values of A against α and of B against β (Fig. 3), we see a clear linear relationship between these variables despite their non-linear interaction in Eq. 1. There are a

486 few exceptions: the Encoder-Decoder model T5-Pile shows a different behavior from the rest of the
487 scaled families, and four additional scaled families show a different relationship between B and β . In
488 fact, all these families share the common feature that they were trained using multiple passes over
489 a single training set Gadre et al. (2024). The outlier point with $\beta > 4$ is a 70m baseline of Pythia
490 for a continual training intervention experiment (Biderman et al., 2023). Future work may consider
491 different function forms tying some of the parameters or introducing other ones instead.

492 Another change for the function form that future work should consider is accounting for the learning
493 rate schedule, as our experiments assumed it was negligible. A mismatch between the form and the
494 real dependence might explain the inconsistencies in using the beginning of training. As noted in §6.1,
495 the beginning is not fitting as well as later on, which we also see to some extent in the percentages
496 axis of Fig.4. This might also be expected as previous works did not take the training trajectory (and
497 loss schedule) into account and ignored this data.

499 10 RELATED WORK

500 This work builds on a large number of recent studies relating scaling law estimation and decision-
501 making about model training. Among the aspects studied are total training costs including inference
502 (Sardana et al.), effects of sophisticated data selection (Sorscher et al., 2022; Ge et al., 2024), training
503 time (Inbar & Sernau, 2024), transfer of learned skills (Hernandez et al., 2021), behavior of models in
504 other modalities (Mikami et al., 2022; Abnar et al.; Alabdulmohsin et al., 2024; Hesslow et al., 2022)
505 mixtures of experts (Ludziejewski et al.), data mixing (Ge et al., 2024), downstream performance
506 (Muennighoff et al., 2024), vocabulary size (Tao et al., 2024), and architecture comparisons (Tay
507 et al., 2023; Poli et al., 2024) including small models (Muckatira et al., 2024) or other phenomena
508 like finetuning (Zhang et al.) and the loss in different positions in the training sequences (Xiong et al.,
509 2024). Especially relevant to our context is Ruan et al. (2024); Maia Polo et al. (2024) that rely on
510 multiple pretraining settings for creating scaling laws that generalize across models or kinds of losses.

511 Another line of works that can be seen as a scaling law discusses the relation between model width
512 and hyperparameter choices (rather than loss) (Yang et al., 2022; 2021; Blake et al., 2024; Lingle,
513 2024).

516 11 LIMITATIONS

517 Our use of ARE as a primary evaluation metric does not distinguish between over-estimation or
518 under-estimation of performance. When using scaling laws to choose between candidate models to
519 train, these error estimates may be unnecessarily conservative (e.g. if both families’ laws are biased
520 in the same direction).

521 Another major limitation in this study is the difficulty of aggregating information across model
522 families. As most published families evaluate models of incomparable scales, often over incomparable
523 ranges, we were unable to produce an informative version of Fig. 2 that aggregated information across
524 all models available, and was thus able to give general recommendations about compute-optimal
525 choice of preliminary experiments.

528 12 DISCUSSION

529 This paper provides a first study of open questions in the estimation of scaling laws and their relation
530 to large-scale pretraining decisions. We expect that many of these conclusions could be sharpened
531 or extended with the availability of additional information about model training, and we call on
532 other leaders of large-scale training efforts to share training losses and evaluation results from
533 multiple checkpoints during pretraining—even in cases where model parameters themselves cannot
534 be released.

535 Our findings leave open many important questions, from performing efficient predictions by fitting
536 on many model families to scaling laws of the deltas between a/b test for a change in attribute (e.g.
537 optimizer) or generalize from one a/b test to another, and to other methods of efficiently compare
538 architectures that do not rely on multiple models (e.g. continual learning). In addition, our results in
539 §9 suggest other scaling law parameterizations might better fit data.

Practical recommendations	
§4	Set an estimation goal and a budget
§5.1	If budget allows, train the whole model for 30%
§A	If extremely constrained, consider predicting with one model
§6	Use all training losses (except the beginning)
§7	Train as big as possible, but limit tokens
§7	Train more models, not only larger

REFERENCES

- Samira Abnar, Mostafa Dehghani, Behnam Neyshabur, and Hanie Sedghi. Exploring the limits of large scale pre-training. In *International Conference on Learning Representations*.
- Ekin Akyürek, Bailin Wang, Yoon Kim, and Jacob Andreas. In-context language learning: Architectures and algorithms. *arXiv preprint arXiv:2401.12973*, 2024.
- Ibrahim M Alabdulmohsin, Behnam Neyshabur, and Xiaohua Zhai. Revisiting neural scaling laws in language and vision. *Advances in Neural Information Processing Systems*, 35:22300–22312, 2022.
- Ibrahim M Alabdulmohsin, Xiaohua Zhai, Alexander Kolesnikov, and Lucas Beyer. Getting vit in shape: Scaling laws for compute-optimal model design. *Advances in Neural Information Processing Systems*, 36, 2024.
- Tamay Besiroglu, Ege Erdil, Matthew Barnett, and Josh You. Chinchilla scaling: A replication attempt, 2024. URL <https://arxiv.org/abs/2404.10102>.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, Usven Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar Van Der Wal. Pythia: A suite for analyzing large language models across training and scaling. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 2397–2430. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/biderman23a.html>.
- Charlie Blake, Constantin Eichenberg, Josef Dean, Lukas Balles, Luke Y Prince, Björn Deiseroth, Andres Felipe Cruz-Salinas, Carlo Luschi, Samuel Weinbach, and Douglas Orr. u-mu p: The unit-scaled maximal update parametrization. *arXiv preprint arXiv:2407.17465*, 2024.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020.
- Ethan Caballero, Kshitij Gupta, Irina Rish, and David Krueger. Broken neural scaling laws. In *The Eleventh International Conference on Learning Representations*.
- Angelica Chen, Ravid Shwartz-Ziv, Kyunghyun Cho, Matthew L Leavitt, and Naomi Saphra. Sudden drops in the loss: Syntax acquisition, phase transitions, and simplicity bias in mlms. In *The Twelfth International Conference on Learning Representations*.
- Leshem Choshen, Guy Hacohen, Daphna Weinshall, and Omri Abend. The grammar-learning trajectories of neural language models. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8281–8297, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.568. URL <https://aclanthology.org/2022.acl-long.568>.

594 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
595 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn,
596 Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston
597 Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron,
598 Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris
599 McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton
600 Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, David
601 Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes,
602 Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip
603 Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme
604 Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu,
605 Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov,
606 Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah,
607 Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu
608 Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph
609 Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani,
610 Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz
611 Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence
612 Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas
613 Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri,
614 Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis,
615 Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov,
616 Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Celebi, Patrick Alrassy, Pengchuan
617 Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan,
618 Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit
619 Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou,
620 Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia
621 Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan,
622 Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla,
623 Stephane Collet, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek
624 Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao,
625 Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent
626 Gonguet, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu,
627 Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia,
628 Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen
629 Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe
630 Papakipos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya
631 Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex
632 Vaughan, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei
633 Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew
634 Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley
635 Gabriel, Ashwin Barambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin
636 Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu,
637 Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt
638 Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changan Wang, Changkyu Kim, Chao
639 Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon
640 Civin, Dana Beaty, Daniel Kreymer, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide
641 Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le,
642 Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily
643 Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix
644 Kreuk, Feng Tian, Firat Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank
645 Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern,
646 Govind Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid
647 Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen
Suk, Henry Aspegren, Hunter Goldman, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat,
Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff
Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin,
Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh

- 648 Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal,
649 Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun
650 Huang, Kunal Chawla, Kushal Lakhotia, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender
651 A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca
652 Wehrstedt, Madian Khabisa, Manav Avalani, Manish Bhatt, Maria Tsimpoukelli, Martynas Mankus,
653 Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi,
654 Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov,
655 Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat,
656 Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White,
657 Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning
658 Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli,
659 Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux,
660 Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao,
661 Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li,
662 Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Ruty Rinott,
663 Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru
664 Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Sheng
665 Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang,
666 Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen,
667 Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny
668 Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara
669 Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou,
670 Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish
671 Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov,
672 Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaofang
673 Wang, Xiaojian Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia,
674 Yendi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and
675 Zhiwei Zhao. The llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.
- 676 Samir Yitzhak Gadre, Georgios Smyrnis, Vaishaal Shankar, Suchin Gururangan, Mitchell Wortsman,
677 Rulin Shao, Jean Mercat, Alex Fang, Jeffrey Li, Sedrick Keh, et al. Language models scale reliably
678 with over-training and on downstream tasks. *arXiv preprint arXiv:2403.08540*, 2024.
- 679 Ce Ge, Zhijian Ma, Daoyuan Chen, Yaliang Li, and Bolin Ding. Data mixing made efficient: A
680 bivariate scaling law for language model pretraining, 2024. URL <https://arxiv.org/abs/2405.14908>.
- 681
- 682 Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord,
683 Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson,
684 Russell Authur, Khyathi Raghavi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu,
685 Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik,
686 Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk,
687 Saurabh Shah, Will Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep
688 Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Sol-
689 daini, Noah A. Smith, and Hannaneh Hajishirzi. Olmo: Accelerating the science of language
690 models, 2024. URL <https://arxiv.org/abs/2402.00838>.
- 691
- 692 Danny Hernandez, Jared Kaplan, Tom Henighan, and Sam McCandlish. Scaling laws for transfer.
693 *arXiv preprint arXiv:2102.01293*, 2021.
- 694 Daniel Hesslow, Niccoló Zanichelli, Pascal Notin, Iacopo Poli, and Debora Marks. Rita: a study on
695 scaling up generative protein sequence models. *arXiv preprint arXiv:2205.05789*, 2022.
- 696
- 697 Dylan Hillier, Leon Guertler, Cheston Tan, Palaash Agrawal, Ruirui Chen, and Bobby Cheng. Super
698 tiny language models. *ArXiv*, abs/2405.14159, 2024. URL <https://api.semanticscholar.org/CorpusID:269982112>.
- 699
- 700 Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza
701 Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al.
Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.

- 702 Shengding Hu, Yuge Tu, Xu Han, Chaoqun He, Ganqu Cui, Xiang Long, Zhi Zheng, Yewei Fang,
703 Yuxiang Huang, Weilin Zhao, et al. Minicpm: Unveiling the potential of small language models
704 with scalable training strategies. *arXiv preprint arXiv:2404.06395*, 2024.
705
- 706 Itay Inbar and Luke Sernau. Time matters: Scaling laws for any budget. *arXiv preprint*
707 *arXiv:2406.18922*, 2024.
- 708 Berivan Isik, Natalia Ponomareva, Hussein Hazimeh, Dimitris Pappas, Sergei Vassilvitskii, and
709 Sanmi Koyejo. Scaling laws for downstream task performance of large language models. *arXiv*
710 *preprint arXiv:2402.04177*, 2024.
711
- 712 Maor Ivgi, Yair Carmon, and Jonathan Berant. Scaling laws under the microscope: Predicting
713 transformer performance from small scale experiments. *arXiv preprint arXiv:2202.06387*, 2022.
714
- 715 Samy Jelassi, David Brandfonbrener, Sham M Kakade, and Eran Malach. Repeat after me: Trans-
716 formers are better than state space models at copying. *arXiv preprint arXiv:2402.01032*, 2024.
- 717 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott
718 Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models.
719 *arXiv preprint arXiv:2001.08361*, 2020.
720
- 721 Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman
722 Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. Bloom: A 176b-
723 parameter open-access multilingual language model. 2023.
724
- 725 Lucas Lingle. A large-scale exploration of mu-transfer. *arXiv preprint arXiv:2404.05728*, 2024.
- 726 Zhengzhong Liu, Aurick Qiao, Willie Neiswanger, Hongyi Wang, Bowen Tan, Tianhua Tao, Junbo
727 Li, Yuqi Wang, Suqi Sun, Omkar Pangarkar, Richard Fan, Yi Gu, Victor Miller, Yonghao Zhuang,
728 Guowei He, Haonan Li, Fajri Koto, Liping Tang, Nikhil Ranjan, Zhiqiang Shen, Xuguang Ren,
729 Roberto Iriondo, Cun Mu, Zhiting Hu, Mark Schulze, Preslav Nakov, Tim Baldwin, and Eric P.
730 Xing. Llm360: Towards fully transparent open-source llms, 2023.
731
- 732 Jan Ludziejewski, Jakub Krajewski, Kamil Adamczewski, Maciej Pióro, Michał Krutul, Szymon
733 Antoniak, Kamil Ciebiera, Krystian Król, Tomasz Odrzygóźdź, Piotr Sankowski, et al. Scaling
734 laws for fine-grained mixture of experts. In *Forty-first International Conference on Machine*
735 *Learning*.
- 736 Felipe Maia Polo, Seamus Somerstep, Leshem Choshen, Yuekai Sun, and Mikhail Yurochkin. Sloth:
737 scaling laws for llm skills to predict multi-benchmark performance across families. *arXiv preprint*
738 *arXiv:2410.11840*, 2024.
739
- 740 Hiroaki Mikami, Kenji Fukumizu, Shogo Murai, Shuji Suzuki, Yuta Kikuchi, Taiji Suzuki, Shin-ichi
741 Maeda, and Kohei Hayashi. A scaling law for syn2real transfer: How much is your pre-training
742 effective? In *Joint European Conference on Machine Learning and Knowledge Discovery in*
743 *Databases*, pp. 477–492. Springer, 2022.
- 744 Sherin Muckatira, Vijeta Deshpande, Vladislav Lialin, and Anna Rumshisky. Emergent abilities in
745 reduced-scale generative language models. In Kevin Duh, Helena Gomez, and Steven Bethard
746 (eds.), *Findings of the Association for Computational Linguistics: NAACL 2024*, pp. 1242–1257,
747 Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/
748 2024.findings-naacl.79. URL <https://aclanthology.org/2024.findings-naacl.79>.
- 749 Niklas Muennighoff, Alexander Rush, Boaz Barak, Teven Le Scao, Nouamane Tazi, Aleksandra
750 Piktus, Sampo Pyysalo, Thomas Wolf, and Colin A Raffel. Scaling data-constrained language
751 models. *Advances in Neural Information Processing Systems*, 36, 2024.
752
- 753 David Owen. How predictable is language model benchmark performance? *arXiv preprint*
754 *arXiv:2401.04757*, 2024.
755
- Rohan Pandey. gzip predicts data-dependent scaling laws. *arXiv preprint arXiv:2405.16684*, 2024.

- 756 Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier
757 Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn:
758 Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
759
- 760 Yotam Perlitz, Elron Bandel, Ariel Gera, Ofir Arviv, Liat Ein-Dor, Eyal Shnarch, Noam Slonim,
761 Michal Shmueli-Scheuer, and Leshem Choshen. Efficient benchmarking (of language models).
762 In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Confer-*
763 *ence of the North American Chapter of the Association for Computational Linguistics: Human*
764 *Language Technologies (Volume 1: Long Papers)*, pp. 2519–2536, Mexico City, Mexico, June
765 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.139. URL
766 <https://aclanthology.org/2024.naacl-long.139>.
- 767 Michael Poli, Armin W Thomas, Eric Nguyen, Pragaash Ponnusamy, Björn Deiseroth, Kristian
768 Kersting, Taiji Suzuki, Brian Hie, Stefano Ermon, Christopher Ré, Ce Zhang, and Stefano Massaroli.
769 Mechanistic design and scaling of hybrid architectures, 2024. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2403.17844)
770 [2403.17844](https://arxiv.org/abs/2403.17844).
- 771 Tomer Porian, Mitchell Wortsman, Jenia Jitsev, Ludwig Schmidt, and Yair Carmon. Resolving
772 discrepancies in compute-optimal scaling of language models. *arXiv preprint arXiv:2406.19146*,
773 2024.
- 774 Jonathan S Rosenfeld, Amir Rosenfeld, Yonatan Belinkov, and Nir Shavit. A constructive prediction
775 of the generalization error across scales. In *International Conference on Learning Representations*.
776
- 777 Yangjun Ruan, Chris J. Maddison, and Tatsunori Hashimoto. Observational scaling laws and the
778 predictability of language model performance, 2024.
- 779
- 780 Nikhil Sardana, Jacob Portes, Sasha Doubov, and Jonathan Frankle. Beyond chinchilla-optimal:
781 Accounting for inference in language model scaling laws. In *Forty-first International Conference*
782 *on Machine Learning*.
- 783 Thibault Sellam, Steve Yadlowsky, Ian Tenney, Jason Wei, Naomi Saphra, Alexander D’Amour,
784 Tal Linzen, Jasmijn Bastings, Iulia Raluca Turc, Jacob Eisenstein, et al. The multiberts: Bert
785 reproductions for robustness analysis. In *International Conference on Learning Representations*,
786 2021.
- 787 Yikang Shen, Zheyu Zhang, Tianyou Cao, Shawn Tan, Zhenfang Chen, and Chuang Gan. Mod-
788 uleformer: Learning modular large language models from uncurated data. *arXiv preprint*
789 *arXiv:2306.04640*, 2023.
- 790
- 791 Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari Morcos. Beyond neural
792 scaling laws: beating power law scaling via data pruning. *Advances in Neural Information*
793 *Processing Systems*, 35:19523–19536, 2022.
- 794 Lintang Sutawika, Aran Komatsuzaki, and Colin Raffel. Pile-t5, 2024. URL [https://blog.](https://blog.eleuther.ai/pile-t5/)
795 [eleuther.ai/pile-t5/](https://blog.eleuther.ai/pile-t5/). Blog post.
- 796
- 797 Chaofan Tao, Qian Liu, Longxu Dou, Niklas Muennighoff, Zhongwei Wan, Ping Luo, Min Lin, and
798 Ngai Wong. Scaling laws with vocabulary: Larger models deserve larger vocabularies. *arXiv*
799 *preprint arXiv:2407.13623*, 2024.
- 800 Yi Tay, Mostafa Dehghani, Samira Abnar, Hyung Chung, William Fedus, Jinfeng Rao, Sharan
801 Narang, Vinh Tran, Dani Yogatama, and Donald Metzler. Scaling laws vs model architectures:
802 How does inductive bias influence scaling? In Houda Bouamor, Juan Pino, and Kalika Bali
803 (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 12342–12364,
804 Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.
805 [findings-emnlp.825](https://aclanthology.org/2023.findings-emnlp.825). URL <https://aclanthology.org/2023.findings-emnlp.825>.
- 806
- 807 The LLM360 Team. Llm360 k2-65b: Scaling up fully transparent open-source llms. 2024.
- 808 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay
809 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation
and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

Alex Warstadt, Aaron Mueller, Leshem Choshen, Ethan Wilcox, Chengxu Zhuang, Juan Ciro, Rafael Mosquera, Bhargavi Paranjabe, Adina Williams, Tal Linzen, et al. Findings of the babyLM challenge: Sample-efficient pretraining on developmentally plausible corpora. In *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pp. 1–34, 2023. URL <https://aclanthology.org/2023.conll-babyLM.1/>.

Mitchell Wortsman, Peter J Liu, Lechao Xiao, Katie E Everett, Alexander A Alemi, Ben Adlam, John D Co-Reyes, Izzeddin Gur, Abhishek Kumar, Roman Novak, et al. Small-scale proxies for large-scale transformer training instabilities. In *The Twelfth International Conference on Learning Representations*, 2023.

Mengzhou Xia, Mikel Artetxe, Chunting Zhou, Xi Victoria Lin, Ramakanth Pasunuru, Danqi Chen, Luke Zettlemoyer, and Veselin Stoyanov. Training trajectories of language models across scales. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13711–13738, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.767. URL <https://aclanthology.org/2023.acl-long.767>.

Yizhe Xiong, Xiansheng Chen, Xin Ye, Hui Chen, Zijia Lin, Haoran Lian, Jianwei Niu, and Guiguang Ding. Temporal scaling law for large language models. 2024. URL <https://api.semanticscholar.org/CorpusID:269449894>.

Ge Yang, Edward Hu, Igor Babuschkin, Szymon Sidor, Xiaodong Liu, David Farhi, Nick Ryder, Jakub Pachocki, Weizhu Chen, and Jianfeng Gao. Tuning large neural networks via zero-shot hyperparameter transfer. *Advances in Neural Information Processing Systems*, 34:17084–17097, 2021.

Greg Yang, Edward J Hu, Igor Babuschkin, Szymon Sidor, Xiaodong Liu, David Farhi, Nick Ryder, Jakub Pachocki, Weizhu Chen, and Jianfeng Gao. Tensor programs v: Tuning large neural networks via zero-shot hyperparameter transfer. *arXiv preprint arXiv:2203.03466*, 2022.

Biao Zhang, Zhongtao Liu, Colin Cherry, and Orhan Firat. When scaling meets llm finetuning: The effect of data, model and finetuning method. In *The Twelfth International Conference on Learning Representations*.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. Opt: Open pre-trained transformer language models, 2022. URL <https://arxiv.org/abs/2205.01068>.

A SCALE UP WITH 1 MODEL

We bring errors of data from fitting from a single model on a given percentage of training to the largest model with full training. Scaling is constant and follows the literature (Muennighoff et al., 2024) and the largest model stands as target model (so the bottom line in each figure represents predicting from the beginning of training). In parallel to this paper, an even more efficient work on predicting with 1 model was suggested, and the two should be incorporated (Maia Polo et al., 2024).

B LOSS CURVES AND PREDICTIONS

We provide in Fig. 7 graphs of the loss during training of the target models per originating source (e.g., a paper) together with the predictions by using different percentage of the training.

C IS SCALING WORKING ONLY UPWARDS?

No. Small models usually show consistent and predicatable performance.

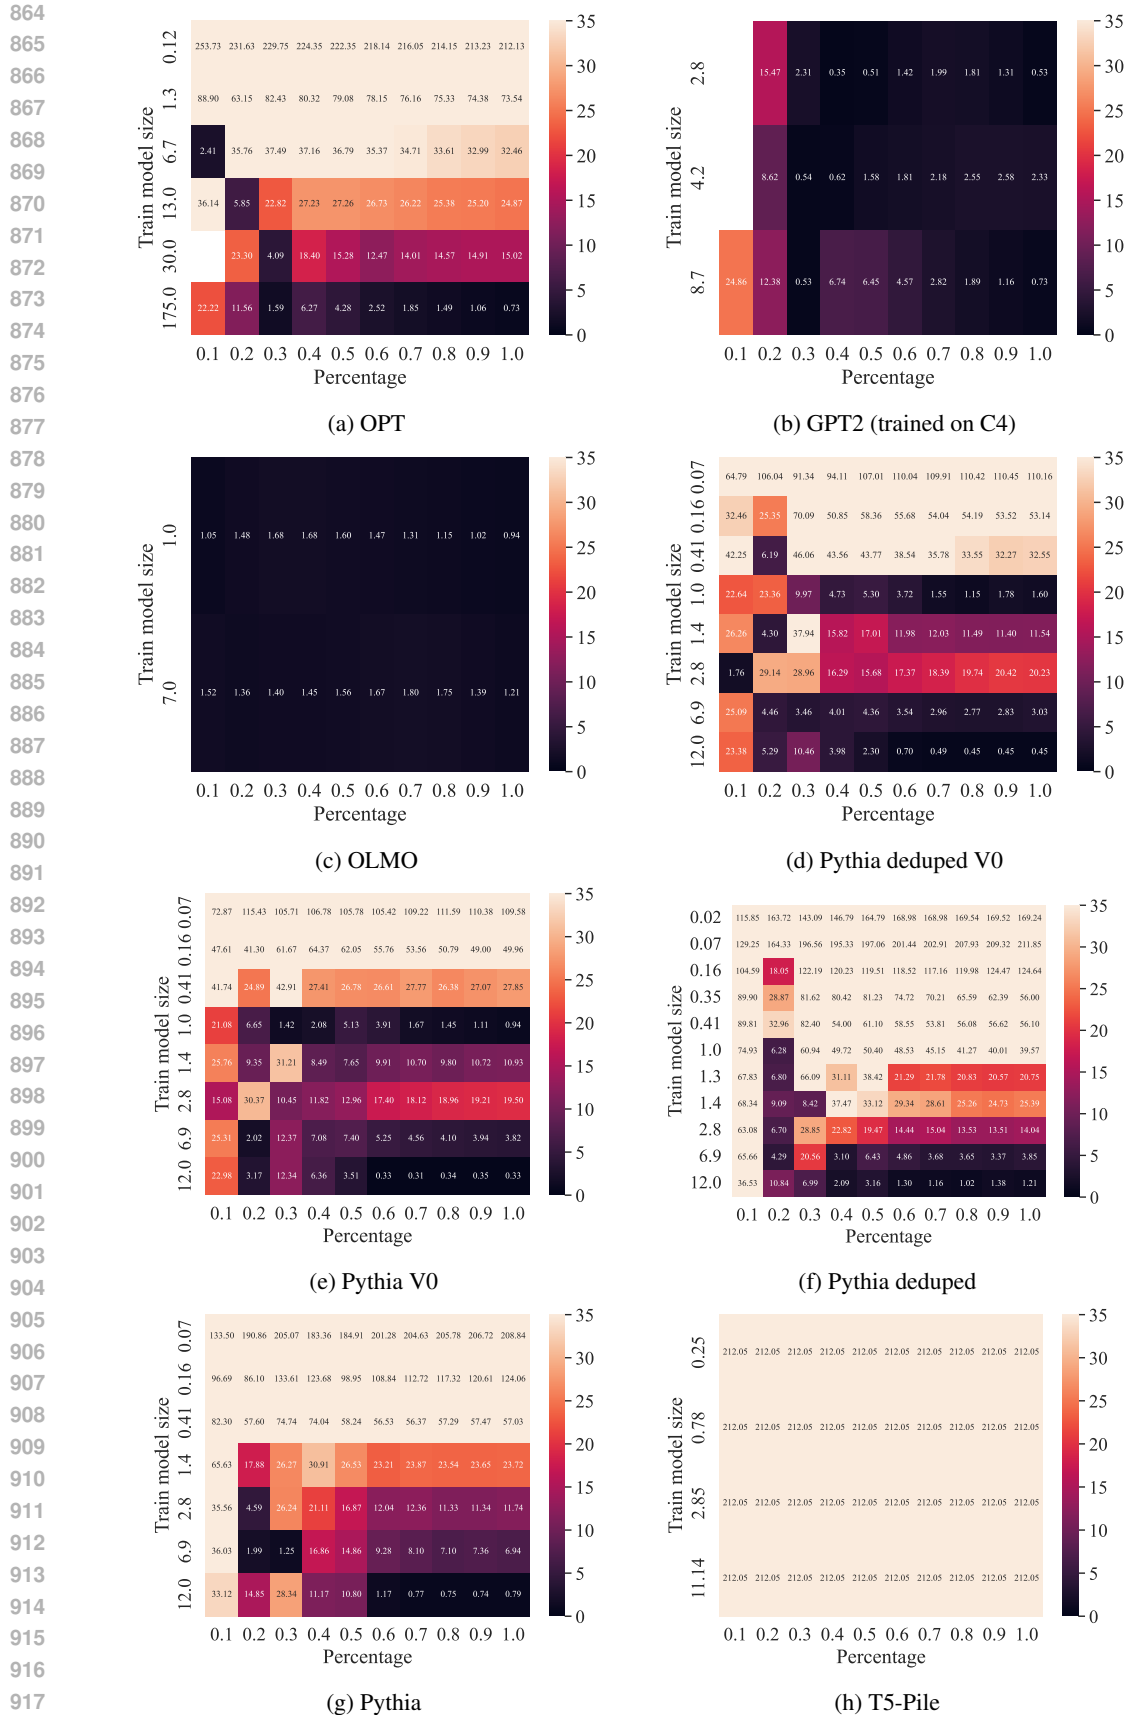


Figure 6: Fitting scaling laws under the assumption that all models scale similarly. Thus, a single model is needed to predict. The last row in each Figure represents predicting a model at the beginning of its training.

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

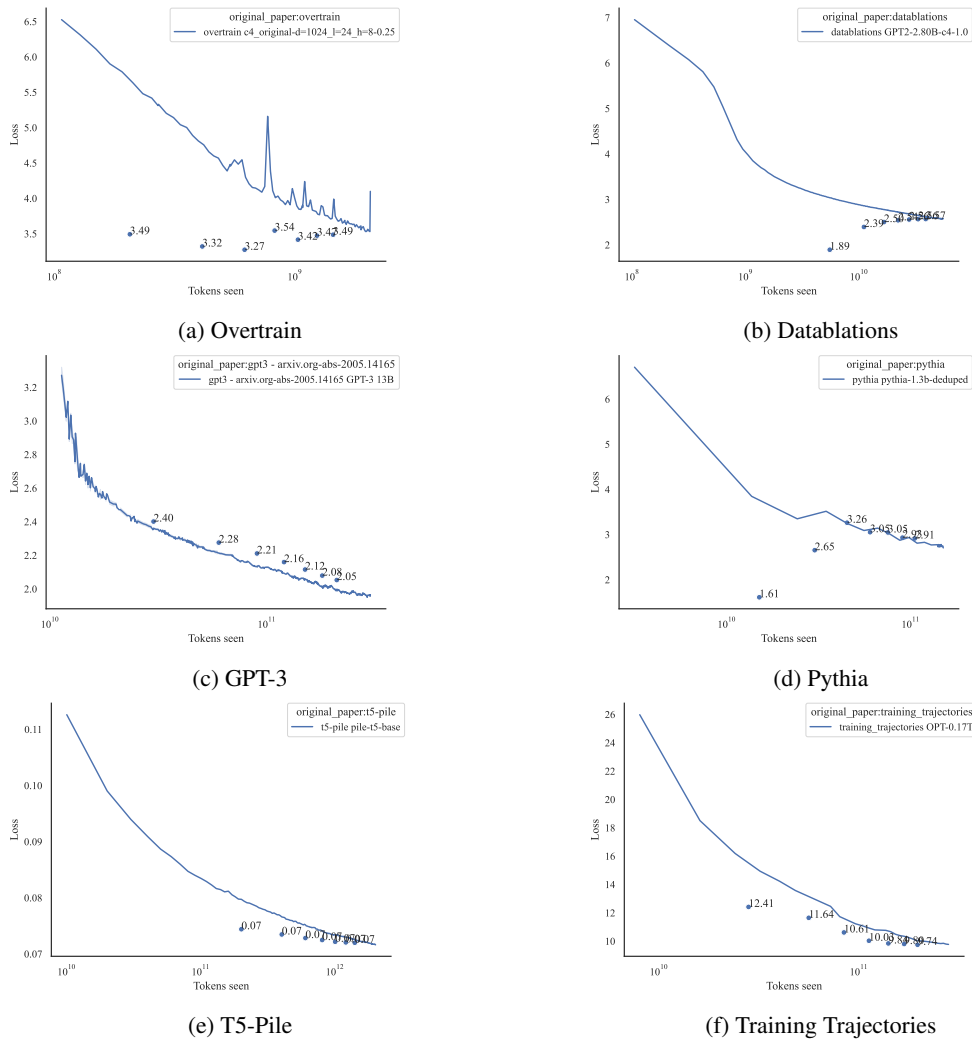


Figure 7: In each figure all losses from a specific source and predictions of the scaling loss with different percentage of the #toks and all models. Predictions are points where the X axis is the available data for prediction and y the prediction value, lines are the actual value. One scaled family per paper was sampled as a representative.

Usually, one does not use a scaling law to extrapolate to a smaller model as one can just train the small model. However, under observational scaling laws, where one wants to research a phenomenon without scaling at all (Ruan et al., 2024; Maia Polo et al., 2024), or when many models were trained and one wishes to create smaller models for various reasons (Hillier et al., 2024; Warstadt et al., 2023), scaling down might prove useful. Moreover, in the context of traditional scaling laws this may act as a baseline. Such an experiment may shed another light on the number of models $|F|$ versus their size #params. If large models are better because they are more stable or otherwise fit laws more robustly, few models will be enough, if the number of models or scale down difference from the prediction, it will show similar behaviour to scaling up. See more in §8.

To test this we reverse the order of models and predict with the largest models the loss on the smallest models. This means that for example in the case of 3 models, we predict the smallest model’s loss and fit the scaling law relying on the 3 largest models. As before, we break the results by the percentage of training done and do not reverse it.

As shown in Fig. 8, the number of models plays an important role in fitting well and a minimum of 30-40% of the training is necessary for good fit, more than that often improves further.

972
 973
 974
 975
 976
 977
 978
 979
 980
 981
 982
 983
 984
 985
 986
 987
 988
 989
 990
 991
 992
 993
 994
 995
 996
 997
 998
 999
 1000
 1001
 1002
 1003
 1004
 1005
 1006
 1007
 1008
 1009
 1010
 1011
 1012
 1013
 1014
 1015
 1016
 1017
 1018
 1019
 1020
 1021
 1022
 1023
 1024
 1025




Figure 8: Fitting scaling laws trying to predict the smallest model, with the largest (Y-axis) models trained on a percentage of the data (X-axis).

D CAN WE DETECT BAD MODELS TO FIT ON?

If so, not through cross validation.

In §7, we raise the issue of instability of scaling law predictions, with a single model vastly changing the results. We tried to see if, without knowing the ARE, we could remove bad models from the prediction. We hypothesized that models that we can't predict would mean models that would skew our predictions when fitted upon. We performed a cross-validation on the #params families in F each time setting the models with most #toks as target ans excluding the #params family from F . Our hypothesis was found to be incorrect. Such cases of hard-to-predict models were found to indicate that the models left in F are bad predictors and not that the target is very dissimilar (a "bad" training). In 58% of the cases removing that model from the scaling law created the worst ARE possible on the actual target, more than removing any other model.

E HUBER REPLICATION

Huber loss is sometimes used instead of ARE (Hoffmann et al., 2022). Huber loss is defined as

$$L_{\delta}(a) = \begin{cases} \frac{1}{2}a^2 & \text{for } |a| \leq \delta, \\ \delta \cdot (|a| - \frac{1}{2}\delta) & \text{otherwise.} \end{cases} \quad (2)$$

We use $\delta = e10^{-3}$ as done in (Hoffmann et al., 2022). The overall results are similar but for completeness report them in Fig. 9.

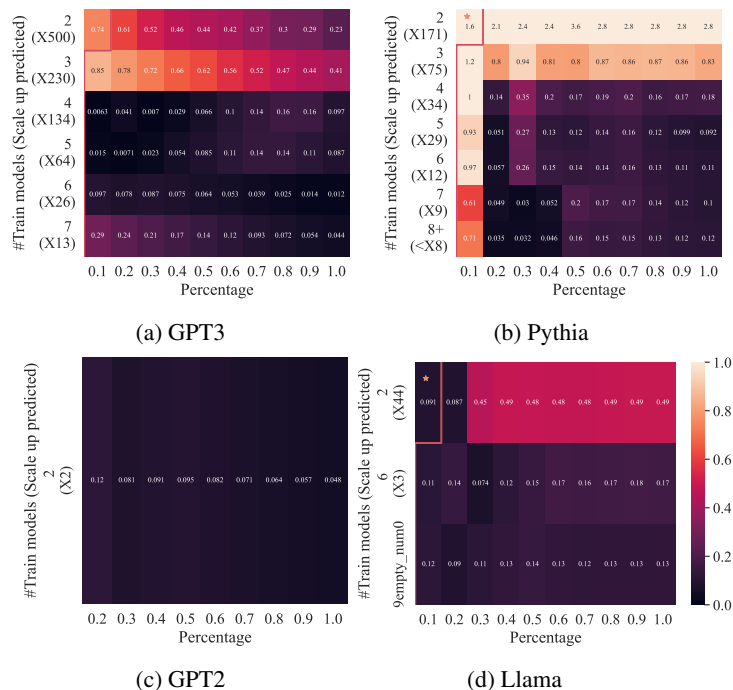


Figure 9: Scaling laws under a Huber loss. The line represents most efficient setting to receive <0.05. Comparison of several models given different amount of models and percentages of training.