

Scaling Laws Under the Microscope: Predicting Transformer Performance from Small Scale Experiments

Maor Ivgi Yair Carmon Jonathan Berant

The Blavatnik School of Computer Science, Tel-Aviv University

{maor.ivgi, joberant}@cs.tau.ac.il , ycarmon@tauex.tau.ac.il

Abstract

Neural scaling laws define a predictable relationship between a model’s parameter count and its performance after training in the form of a power law. However, most research to date has not explicitly investigated whether scaling laws can be used to accelerate model development. In this work, we perform such an empirical investigation across a wide range of language understanding tasks, starting from models with as few as 10K parameters, and evaluate downstream performance across 9 language understanding tasks. We find that scaling laws emerge at finetuning time in *some* NLP tasks, and that they can also be exploited for debugging convergence when training large models. Moreover, for tasks where scaling laws exist, they can be used to predict the performance of larger models, which enables effective model selection. However, revealing scaling laws requires careful hyperparameter tuning and multiple runs for the purpose of uncertainty estimation, which incurs additional overhead, partially offsetting the computational benefits.

1 Introduction

Transformer-based language models (LMs) (Vaswani et al., 2017; Devlin et al., 2019; Raffel et al., 2019; Radford et al., 2019; Brown et al., 2020) which are at the foundation of modern NLP systems, have been recently shown to exhibit scaling laws (Kaplan et al., 2020; Hernandez et al., 2021; Tay et al., 2021, 2022), that is, the test loss of LMs obeys a predictable power law with respect to the number of model parameters, dataset size, and computation budget. This finding ignited substantial research that demonstrated scaling laws in a wide range of areas including computer vision (Rosenfeld et al., 2020; Henighan et al., 2020; Zhai et al., 2021; Bahri et al., 2021; Abnar et al., 2021), acoustic models (Droppo and Elibol, 2021), and board games (Jones, 2021; Ben-Assayag and El-Yaniv, 2021), among others.

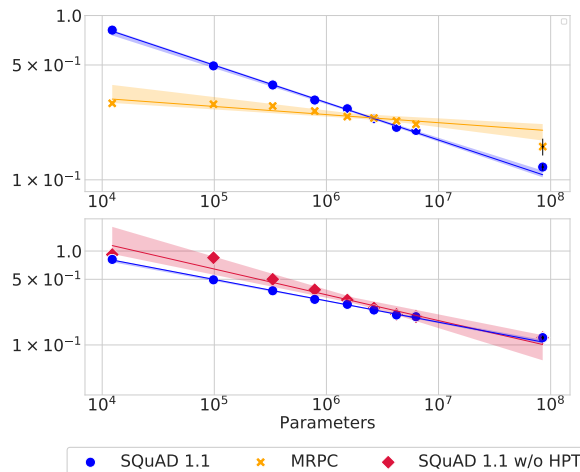


Figure 1: Performance of models when finetuned on SQuAD 1.1 (evaluated using $1 - F_1$) and MRPC (classification error). **Top:** Models exhibit a clean scaling law fit on SQuAD 1.1 ($R^2 = 0.998$) compared to MRPC ($R^2 = 0.763$). **Bottom:** Without hyperparameter tuning (w/o HPT) the goodness-of-fit is lower ($R^2 = 0.83$).

On top of being a fascinating phenomenon on its own, scaling laws can potentially be harnessed for more efficient development of models. Specifically, if scaling laws hold, one can perform modeling decisions at multiple small scales, and extrapolate to infer which model will perform best at a larger scale. While starting small is an established technique (Tan and Le, 2019), scaling laws can provide a more principled framework for this methodology. Moreover, scaling laws can potentially accelerate the development cycle, and reduce carbon footprint caused by training large neural models (Schwartz et al., 2020; Bender et al., 2021).

However, for this idea to materialize, several questions must be addressed, which have not been fully considered by past literature. First, do scaling laws consistently occur at finetuning time across a wide range of language understanding tasks? Second, do they manifest reliably at the small-scale regime? And third, what is the predictive power of

scaling laws when used to predict the behavior of large models for the purpose of model selection? While recent work touched upon these questions (Rosenfeld et al., 2020; Tay et al., 2021, 2022), most work focused on post-hoc analysis of highly-parameterized models, and looked at performance on aggregate benchmarks, such as GLUE (Wang et al., 2019b) and SuperGLUE (Wang et al., 2019a), without analyzing performance at the level of individual dataset.

In this work, we perform a thorough empirical evaluation of scaling laws from the perspective of NLP researchers with resource constraints, addressing the aforementioned questions. We analyze scaling laws at finetuning time across 9 different tasks with models ranging from merely 10K parameters up to roughly 100M (e.g. BERT-Base (Devlin et al., 2019)). Moreover, we move away from post-hoc analysis and evaluate whether scaling laws can be used to predict the performance of larger models. As a case study, we test whether scaling laws can be used for model selection, by comparing the performance of two LMs with two different pre-training objectives: Masked Language Modeling (MLM) (Devlin et al., 2019) and Pointwise Mutual Information (PMI) (Levine et al., 2021).

Our experiments reveal several findings.

1. We show near-perfect scaling laws with high goodness-of-fit at pretraining time for both MLM and PMI even with the smallest models and for multiple architectural choices (§3).
2. At finetuning time, scaling laws emerge only in some of the tasks (§4.1). Fig. 1 (top) shows an example of a good fit over SQuAD (Rajpurkar et al., 2016) and a less impressive fit over MRPC (Dolan and Brockett, 2005).
3. In some tasks, a certain minimal model size is required (§4.1). For example, on MRPC only models with at least four hidden layers performed significantly better than chance.
4. Careful hyperparameter tuning (HPT) can be crucial for exposing precise scaling laws at finetuning time, especially at smaller scales (§4.2). For example, in Fig. 1 (bottom), not only is the fit much better with HPT (blue), it also dramatically affects the prediction at larger scales.
5. As for using scaling laws for model selection (§4.3, §4.4), our MLM vs. PMI case study

shows that scaling laws can be used to perform model selection at larger scales whenever we observe high goodness-of-fit ($R^2 > 0.95$) of the scaling law at smaller scales.

Overall, our empirical findings paint a nuanced picture of the potential of scaling laws as a tool for model design. On one hand, we observe scaling laws at finetuning time for some NLP tasks, and show that they can be used to predict the performance of a model that is 10x larger. On the other hand, this does not happen consistently on all tasks, and revealing these scaling laws requires careful control over hyperparameters and convergence conditions, incurring additional overhead that might counteract the computational benefits.

2 Method

We describe our experimental setup in §2.1, and then our procedure for evaluating goodness-of-fit and the predictive power of scaling laws in §2.2.

2.1 Experimental setup

Architecture We consider encoder-only transformer models, similar to the architecture of BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019b). Encoder-only models are ubiquitous in NLP for a wide range of classification tasks such as natural language inference, question answering, information extraction, text classification, and more.

Model configurations Recent work (Tay et al., 2022, 2021; Hernandez et al., 2021; Zhai et al., 2021; Bahri et al., 2021) focused on parameter-rich transformers, ranging from 5–10M trainable parameters to 40B. We, instead, investigate smaller models, with as few as 10K trainable parameters, assuming a computationally-constrained environment. To preserve the architecture as we scale model parameters, we increase the number of layers (L) from one to twelve, while keeping the aspect ratio (AR) constant ($AR := \frac{L}{H}$ where H is the hidden layer width).

We experiment with two different aspect ratios, 32 and 64, and scale our models over four orders of magnitude, up to the $\sim 85M$ parameters of BERT-Base (Devlin et al., 2019). In particular, we train small-scale models with aspect ratio 32 with 1 to 8 layers, and small-scale models with aspect ratio 64 with 1 to 5 layers. Finally, we train from scratch a BERT-Base (Devlin et al., 2019) model which has 12 hidden layers and aspect ratio 64. A detailed list of all configurations is in Table 6 in App. A.1.

	Train	Eval	Metric	Top freq.
SQuAD 1.1	87,599	10,570	F_1	N/A
MNLI	98,175	9,815	Acc.	0.35
QNLI	104,743	5,463	Acc.	0.51
SST-2	67,349	872	Acc.	0.51
RACE	87,866	4,887	Acc.	0.27
CoLA	8,551	1,043	MCC	0
SQuAD 2.0	130,319	11,873	$Best-F_1$	N/A
MRPC	3,668	408	Acc.	0.68
BoolQ	9,427	3,270	Acc.	0.62

Table 1: Statistics on downstream tasks. **Top freq.** shows the performance when always predicting the most frequent class in the evaluation set, where applicable. See the original papers definitions of the evaluation metrics. In MNLI we use 25% of the training examples, and use the validation-matched set as evaluation data.

Pretraining We experiment with two pretraining objectives (i.e., “upstream”), as different objectives can affect models downstream performance (Devlin et al., 2019; Liu et al., 2019b; Raffel et al., 2019; Levine et al., 2021). As a test case, we compare the popular Masked Language Modeling (MLM) masking (Devlin et al., 2019), where random tokens are masked, to Pointwise Mutual Information (PMI) masking (Levine et al., 2021), where masking is performed over sequences of tokens that tend to co-occur in the training corpus. We test whether small scale experiments on these objectives can predict performance at larger scale and inform which objective is better for a particular task. To the best of our knowledge, prior work did not compare the effects of different pretraining objectives on the behavior of scaling laws.

Finetuning tasks Past work (Henighan et al., 2020; Abnar et al., 2021; Zhai et al., 2021; Bahri et al., 2021) evaluated performance on finetuning (i.e., “downstream”) tasks in computer vision, but less attention has been given to the relation between architectures and finetuning accuracy in NLP. In this work, we address this lacuna and report results after finetuning on 9 different datasets: SQuAD 1.1 (Rajpurkar et al., 2016), MNLI (Williams et al., 2018), QNLI (Wang et al., 2019b), SST-2 (Socher et al., 2013), RACE (Lai et al., 2017), CoLA (Warstadt et al., 2019), SQuAD 2.0 (Rajpurkar et al., 2018), MRPC (Dolan and Brockett, 2005) and BoolQ (Clark et al., 2019).

For all tasks, we use the common classification head (single layer MLP) on top of the prepended CLS token. In each task, we evaluate using the official metric and finetune on the training data suggested by the authors, except in MNLI where

we randomly sample a subset of 25% of the training data for efficiency. The full specification of datasets can be found in Table 1. For details on the finetuning procedure, refer to App. A.1.

Notation Similar to Kaplan et al. (2020), we denote the number of trainable parameters, not including word embeddings, by N , and estimate it with $N \approx 12LH^2$ where L is the number of layers and H is the hidden dimension.

2.2 Evaluation

While past work reported the scaling coefficients of a fitted power law, there is no consensus on a measure for estimating the *goodness-of-fit* of a scaling law to a set of points. Moreover, once the power law is computed, less attention has been dedicated to estimating its *predictive power* to larger scales. We propose evaluation metrics for these quantities.

Goodness-of-fit Given a task and a metric $F : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ to be minimized, we analyze an architecture across M different scales, where we finetune the architecture T times per scale. To find the power law coefficients, we fit a regression line by performing *least-squares* in log-log scale over all $M \cdot T$ points. We then define the goodness-of-fit as the R^2 measure given by the line and the points.

Because different seeds result in different performance, we wish to compute confidence intervals for the regression line. To this end, we employ a hierarchical bootstrap procedure, where we sample a set of data points $B = 1000$ times, use the points to produce B fitted lines, and compute confidence intervals ($[2.5, 97.5]$ percentiles) around the slope and around each point along the line. We provide full details on the hierarchical bootstrap procedure and the computation of R^2 in App. A.3.

Predictive power Given data on the performance of a model on a larger scale, we evaluate the predictive power of the inferred scaling law by computing the *Mean-Relative-Error (MRE)* between the predicted performance and the true one. In particular, given the true performance values of k experiments $y = (y_1, \dots, y_k)$ on k larger scales, and the corresponding predictions $\hat{y} = (\hat{y}_1, \dots, \hat{y}_k)$ we define

$$MRE(y, \hat{y}) := \frac{1}{k} \sum_{i=1}^k \left| \frac{y_i - \hat{y}_i}{y_i} \right| \quad (1)$$

When $k = 1$, we drop the absolute value to keep information on whether the model is overshooting or undershooting and call this *Relative Error (RE)*.

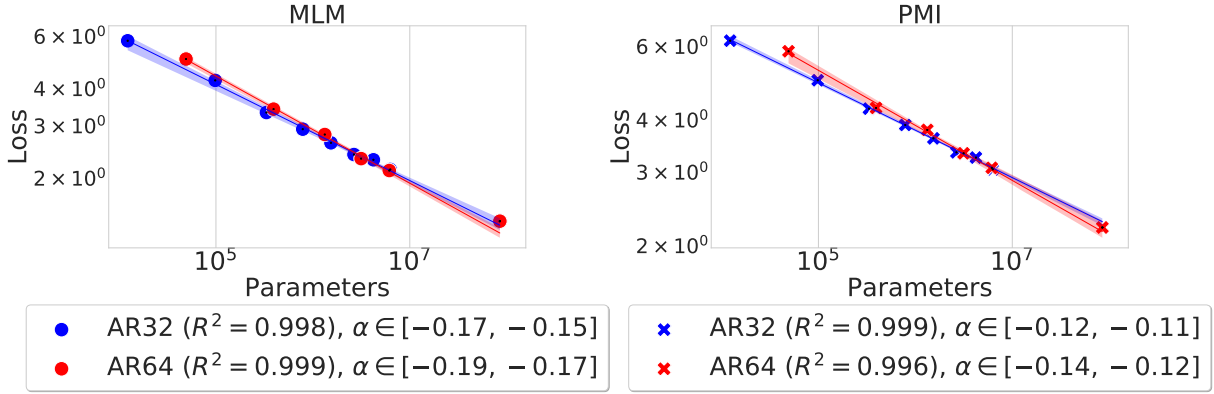


Figure 2: Evaluation loss w.r.t. models’ parameter count, colored by aspect ratio, along with goodness-of-fit and the slope (α). **Left:** MLM objective. **Right:** PMI objective. Note: we do not consider the BERT-Base model (rightmost point) when computing R^2 , since we later use it to examine predictive power. Here, $M = 8$ for AR = 32 and $M = 5$ for AR = 64, and $T = 1$ in both cases (see text for definitions of M and T).

3 Pretraining

We pretrain all models from scratch on the *Datasets* (Lhoest et al., 2021) Wikipedia dump, see experimental details in App. A.1. We use early stopping to declare convergence, but note that determining convergence is non-trivial, as we discuss in §3.1.

Fig. 2 shows the evaluation loss of each configuration as a function of the parameter count for both MLM and PMI. As can be seen from the clean linear relationships (in log-log scale), both aspect ratios and both objectives present a power law, with R^2 exceeding 0.99 in all cases. This is consistent with the language modeling results of Kaplan et al. (2020), and shows that scaling laws exist at pretraining time for the MLM and PMI objectives.

Past work (Kaplan et al., 2020; Tay et al., 2022) examined scaling laws of different architectures (e.g., Transformer (Vaswani et al., 2017), Reformer (Kitaev et al., 2020), Performer (Zhao and Deng, 2019)). They showed that while different architectures affect scaling laws, architectural hyperparameters, such as aspect ratio (AR), make little difference as long as they are within a reasonable range. Interestingly, in both MLM and PMI the slope of AR 64 is slightly better than the slope of AR 32, even when taking into account the slope’s 95% confidence intervals. The intersection of the two AR lines illustrates the potential of using scaling laws for model selection: choosing the AR based on small scale experiments would lead to choosing AR 32, while the performance of models with AR 64 seems comparable and perhaps better in the larger scale. However, the confidence intervals of the fit, depicted as sleeves in Fig. 2, intersect

at the larger scales, meaning we cannot predict a performance difference with confidence.

We note that our largest MLM model, which uses AR 64, 12 layers, and has 85M trainable parameters performs slightly worse than predicted by the power law, which might hint that it is under-trained. We leave verifying this and training larger models with different ARs to future work.

3.1 Debugging convergence with scaling laws

A useful side-effect of the clean scaling law behavior during pretraining is the ability to detect issues in pretraining convergence. In several cases, training stopped due to early stopping, but its loss was greater than predicted by the fit done on other scales. When investigated further, we found that increasing the patience hyperparameter led to further significant loss reduction on the evaluation data, finally converging at the predicted value.

This result points to a methodological issue in current literature, where researchers train models “until convergence”. Convergence is not well-defined, since it is affected by early stopping hyperparameters, such as *patience* and *minimal decrease*. This can lead to under-optimized models, as we observed here. We believe that being precise about the definition of a “converged model” is important for reproducibility of scaling laws research. We further illustrate this and provide the precise criteria we used to declare convergence in App. A.4.

4 Finetuning

Tay et al. (2022) recently showed that evaluation loss during pretraining does not necessarily corre-

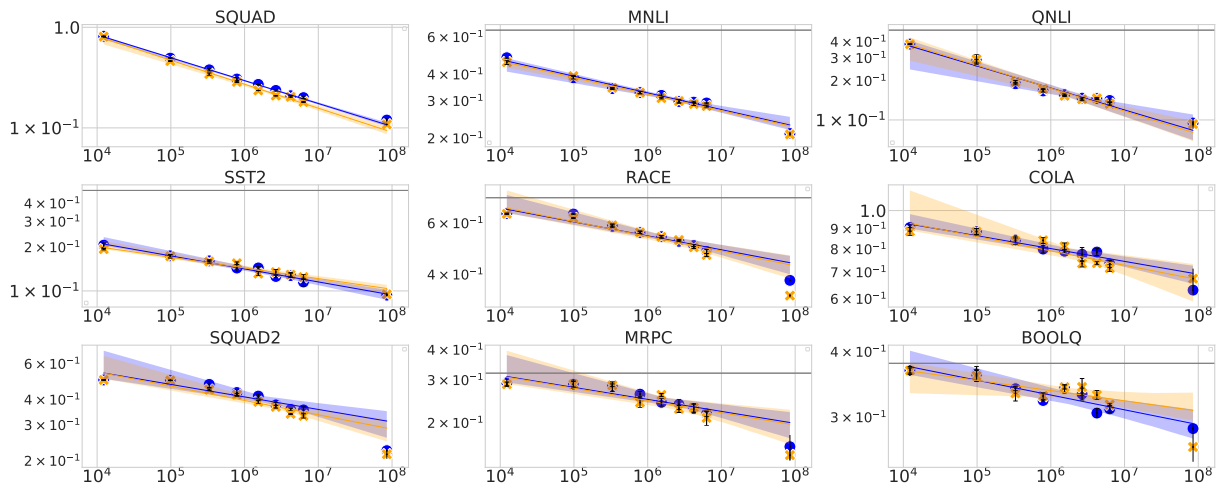


Figure 3: Evaluation performance w.r.t. number of parameters. The y-axis of each plot is $1 - \text{Metric}$ where *Metric* is listed in table 1. The horizontal gray line indicate the simple baseline achieved by a majority vote on the most frequent class, when applicable. **Blue**: models pretrained with MLM. **Orange**: models pretrained with PMI.

	R^2	MRE	$R^2_{2:}$
SQuAD 1.1	0.998	0.006	0.995
MNLI	0.971	0.020	0.961
QNLI	0.965	0.024	0.876
SST-2	0.951	0.005	0.906
RACE	0.916	0.045	0.974
CoLA	0.845	0.088	0.820
SQuAD 2.0	0.797	0.057	0.917
MRPC	0.763	0.020	0.849
BoolQ	0.749	0.023	0.669

Table 2: Goodness-of-fit on finetuning tasks. *MRE* measures the predictive power of scales 1-6 ($M = 6$), evaluated on scales 7,8. The line separates tasks with $R^2 \geq 0.95$ ($M = 8$). $R^2_{2:}$ measures goodness-of-fit on scales 2-8 ($M = 7$). In all cases, $T = 5$.

late with performance on downstream tasks. In this section, we revisit this finding. In particular, we investigate: a) differences in scaling laws across tasks, b) the effects of hyperparameter tuning, architectural design and pretraining objectives, and c) the predictive power of emerging scaling laws. All finetuning experiments use the final checkpoints of the pretrained models described in §3.

4.1 Downstream tasks are not born equal

Tay et al. (2021) and Tay et al. (2022) showed the behavior of transformers over the aggregated GLUE (Wang et al., 2019b) and SuperGLUE (Wang et al., 2019a) benchmarks. We now examine their behavior over a diverse set of NLP tasks.

Fig. 3 and the R^2 values in Table 2 show that different tasks vary in terms of the quality of the power law fit. As before, We do not consider the BERT-Base model with 85M parameters when comput-

ing R^2 since we use this model to test the predictive power of scaling laws. In some tasks, such as SQuAD 1.1, MNLI and QNLI, we observe a relatively clean power law ($R^2 > 0.96$), even though their evaluation metrics (e.g., F_1 and accuracy) differ from their finetuning loss. On the other end of the spectrum is BoolQ, which is not even monotonic w.r.t. number of parameters. Other tasks lie in different places along this spectrum.

We hypothesize that the two factors that play a role in determining the emergence of scaling laws during finetuning are (a) the proximity of the task to the pretraining objective, and (b) the amount of data to finetune on (Table 1). Namely, on all tasks where the training data contains less than 10K examples, R^2 was low (< 0.85). Furthermore, in tasks where $R^2 > 0.95$, the objective can be cast as language modeling with an implicit prompt (e.g. “The sentiment in this review is [MASK].”). Conversely, RACE is a multiple-choice classification task, and indeed its goodness-of-fit is relatively low, despite having a large number of training examples.

An exception to the above is SQuAD 2.0, which presents a worse scaling law compared to SQuAD 1.1, despite having almost 50% more training examples. The main difference between the two tasks is their metric and the existence of non-answerable questions in the latter. While SQuAD 1.1 measures simple F_1 , SQuAD 2.0 evaluates models based on *Best-F1*. That is, the F_1 score reachable if an optimal confidence threshold is chosen *post-hoc* to detect non-answerable questions. We conjecture that both the evaluation metric, and the task of

	R^2		Slope	
	HPT	w/o HPT	HPT	w/o HPT
SQD 1	0.998	0.830	-0.22 ± 0.01	-0.30 ± 0.07
MNLI	0.971	0.933	-0.07 ± 0.02	-0.10 ± 0.03
QNLI	0.965	0.867	-0.14 ± 0.05	-0.18 ± 0.08
SST-2	0.951	0.960	-0.10 ± 0.02	-0.11 ± 0.02
RACE	0.916	0.949	-0.06 ± 0.02	-0.06 ± 0.01
CoLA	0.845	0.735	-0.03 ± 0.01	-0.05 ± 0.02
SQD 2	0.797	0.765	-0.08 ± 0.03	-0.08 ± 0.04
MRPC	0.763	0.740	-0.06 ± 0.03	-0.07 ± 0.03
BoolQ	0.749	0.838	-0.03 ± 0.01	-0.03 ± 0.01

Table 3: Effect of performing hyperparameter tuning on the fit. SQD 1 and SQD 2 refer to SQuAD 1.1 and 2.0.

detecting non-answerable questions contributed to diverging from the LM objective and thus explain the degradation in the goodness-of-fit. To test this, we take the models that were finetuned on SQuAD 2.0 and evaluate them with the SQuAD 1.1 metric on the subset of answerable questions. This indeed increases R^2 from 0.797 to 0.915. We provide further details in App. A.7.

Finally, we note that we do not test whether evaluation loss exhibits a power law on finetuning tasks, since we tune hyperparameters based on the target metric. This is since log-loss can increase significantly even when the target metric is still improving, due to a single example in the evaluation set that incurs higher and higher loss during finetuning (Soudry et al., 2018). For further discussion, see App. A.6.

Critical size A possible reason for low R^2 scores is that models need a minimal “size” to handle a certain task. We define $R_{2_c}^2$ to be the goodness-of-fit when considering only models with depth of at least 2. As is evident from Table 2, the R^2 substantially improves in SQuAD 2.0, RACE and MRPC when omitting the smallest scale. This finding is in line with recent work (Chowdhery et al., 2022; Wei et al., 2022; Srivastava et al., 2022) which suggests some capabilities of language models may emerge only from a certain scale.

4.2 Hyperparameter tuning

When scaling models, one cannot assume that hyperparameters found for one scale would remain optimal for other scales. Despite this, prior work did not report a hyperparameter tuning (HPT) phase during finetuning. Table 3 highlights the difference in goodness-of-fit between models trained with the hyperparameters used by BERT-Base vs. when tuned for each scale individually. As can be seen

	Prediction	Actual	RE	RE _{2_c}
SQD 1	0.893	0.880	0.014	0.011
MNLI	0.769	0.790	-0.026	-0.040
QNLI	0.917	0.905	0.013	0.008
SST-2	0.905	0.906	-0.002	0.001
RACE	0.561	0.618	-0.092	-0.045
CoLA	0.305	0.369	-0.175	-0.125
SQD 2	0.685	0.776	-0.117	-0.076
MRPC	0.800	0.841	-0.049	-0.030
BoolQ	0.708	0.715	-0.009	0.000

Table 4: Predictive power based on 8 smaller models, predicting the performance of BERT-Base. The dashed line separates tasks with $R^2 \geq 0.95$, and RE_{2_c} is the RE when the fit ignores single layer models.

by the R^2 values in the table, when HPT is performed, the power law of scaled models tends to be cleaner. Moreover, because hyperparameters were originally tuned on larger models, smaller scales exhibit a large discrepancy from their optimal performance. This leads to an imprecise power law fit with a steeper slope compared to when HPT is performed (see Table 3), manifested by overshooting predictions. For details on HPT, see App. A.2.

4.3 Predictive power

To check whether scaling laws are useful, we need to evaluate their predictive power. To test that, we conduct the following experiments. First, we split the samples used to fit power laws, as discussed above, and test how well do models with 1-6 layers predict the performance of models with 7 or 8 layers (aspect ratio 32), and evaluate with MRE.

The column MRE in Table 2 shows the results of this experiment. In most tasks, the MRE is quite small ($\leq 2.5\%$), and is correlated with R^2 . For example, the six smaller models (with 1-6 hidden layers) finetuned on SQuAD 1.1 predict the F_1 of the two larger models (7-8 hidden layers) to a 0.6% relative difference (roughly half an F_1 point). One notable case is SST-2, where MRE is excellent (0.5%), but R^2 is lower than some other tasks (0.951). We attribute this to the fact that the slope of the fitted line in SST-2 is relatively gentle – all eight scales score in the range 79.4-88.5, see Fig. 3. Since R^2 measures the proportion of variance explained compared to a constant prediction, it is more sensitive to errors when the slope is close to zero. Similarly, BoolQ also presents a good MRE score even though its R^2 is low. Analyzing Fig. 3 shows that while the fit is poor, most scores lie in a small range and close to the naïve majority baseline, explaining this contradiction.

	Predicted Δ	Actual Δ	Sign(Δ)	R^2
SQD 1	0.013	0.011	Correct	≥ 0.95
MNLI	0.002	0.002	Correct	≥ 0.95
QNLI	0.003	0.002	Correct	≥ 0.95
SST-2	-0.009	-0.001	Correct	≥ 0.95
RACE	0.007	0.043	Correct	< 0.95
CoLA	0.021	-0.043	Wrong	< 0.95
SQD 2	0.024	0.009	Correct	< 0.95
MRPC	0.003	0.012	Correct	< 0.95
BoolQ	-0.017	0.022	Wrong	< 0.95

Table 5: Finetuning results for models pretrained with MLM/PMI. Positive Δ indicates PMI is better than MLM. Tasks above line have $R^2 \geq 0.95$.

Expanding on this, we pretrain a 12-layer BERT-Base model from scratch using the same setup as all models, and with MLM as the objective. We then finetune it on the different tasks. Table 4 shows the relative prediction error (RE) based on the fit from models with 1-8 layers. Note that the largest model, with 8 layers, has 14x less parameters compared to BERT-Base. In all cases where the goodness of fit is high ($R^2 \geq 0.95$), the prediction is accurate to less than 3% difference and in some cases to almost 0.1%. Following the discussion in §4.1 where we saw that RACE, SQuAD 2.0 and MRPC might be affected by a critical size limit, we compute the RE_2 : based on the models with 2-8 layers and report the results in Table 4. Indeed, all three datasets, which gain a significant boost in R_2^2 : compared to R^2 , also benefit improved predictive power. Interestingly, in 7 out of 9 tasks, our prediction is over-optimistic (i.e. RE is negative). This hints that our BERT-Base model might be under-trained.

4.4 Case study: MLM vs. PMI objectives

As a case study, we simulate the use of scaling laws from the perspective of a resource-constrained researcher introducing a new pretraining objective, such as PMI. Specifically, we imagine pretraining and finetuning small models with 1-8 layers, for the PMI and MLM objectives, and then predicting which model will perform better when scaling up to a model with 12 layers, i.e., BERT-Base.

Model selection Table 5 compares the predicted performance gap vs. the actual performance gap in the BERT-Base results (positive values indicate predicted/actual performance of PMI is higher than MLM and vice versa). We find that in all cases where R^2 is high, the predicted gap holds the same sign as the actual one, suggesting that the predictions are useful for performing model selection.

Moreover, in SQuAD 1.1, QNLI and MNLI, the predicted gap is very accurate. Overall, we conclude that when the goodness-of-fit, i.e., R^2 , is high enough, scaling laws present a viable approach for performing model selection without training a large model. We leave for future work to determine if predictions remain accurate when extrapolating over multiple orders-of-magnitude.

Computational efficiency We have shown that for *some* NLP tasks, scaling laws can be an effective tool for model selection, and be further used for debugging convergence. However, applying them requires multiple runs across scales for uncertainty estimation and HPT. Thus, a key question is how much resources are saved with this effort.

To examine this, we perform a theoretical analysis of the FLOPs required to pretrain and finetune the small-scale models vs. the larger ones in our experimental setup, where we extrapolate to a model that is one order-of-magnitude larger. Following Kaplan et al. (2020), we estimate the number of FLOPs for the forward and backward passes with $C \approx 6ND$ where D is the total number of tokens observed. Assuming all models observe the same number of tokens (ignoring early stopping as it requires extra FLOPs for the evaluation set), the difference in computation arises solely from the number of parameters. For example, the total count of parameters of the 8 small-scale models in our setup is 16M while BERT-Base contains 85M, suggesting a 5x improvement. Extrapolating to larger scales will yield more substantial savings, but we did not pretrain and evaluate any larger models.

In practice, we observed that smaller scale models require more epochs during finetuning, increasing their token count, D . Still, even if we sum all FLOPs performed for HPT and finetuning over all scales and *all 9 tasks*, we empirically observed a 2.5x decrease in FLOP count compared to training the larger model. We do not compare runtime because different models were trained on different types of nodes, but we expect savings in terms of runtime to be even greater, as training multiple models is trivial to parallelize.

Overall, one can expect compute savings of 2.5-5x when scaling to a model that is an order of magnitude larger, albeit at the cost of performing careful monitoring of convergence and HPT.

5 Related Work

Scaling laws in transformers Since Kaplan et al. (2020) demonstrated scaling laws for transformer-based language models, researchers have been investigating the extent of this phenomenon. Tay et al. (2022) investigated the effect of inductive bias on scaling laws and showed how different architectures affect the emerging scaling laws. Tay et al. (2021) showed that model shape matters, and that pretraining and finetuning losses are not necessarily correlated. Contrary to Kaplan et al. (2020), they showed that shape also plays a role in finetuning performance, rather than size alone. Hernandez et al. (2021) focused on python code, showing a trade-off between data and compute. Ghorbani et al. (2021) analyzed scaling laws in transformer models used in neural machine translation while Gordon et al. (2021) and Bansal et al. (2022) discussed practical implications of the predictive power of such results.

In parallel to work done in NLP, several works focused on scaling laws in other domains. Henighan et al. (2020) observed scaling laws in multi-modal settings. Zhai et al. (2021) gave a comprehensive review on scaling laws behavior of upstream computer vision tasks, and Abnar et al. (2021) investigated the relationships between upstream and downstream performance of vision transformers.

An important line of work was dedicated to explaining the emergence of scaling laws in neural models (Hashimoto, 2021). Bahri et al. (2021) connect the scaling exponent to the intrinsic dimension of the data-manifold realized by trained networks representations. Bordelon et al. (2020) and Bisla et al. (2021) connect scaling behavior to the spectrum of the kernel in the related NTK model. Theoretical explanations for neural scaling laws include analogy to kernel methods (Spigler et al., 2020; Bordelon et al., 2020), nearest neighbors methods (Sharma and Kaplan, 2022; Bisla et al., 2021), or a combination of the two (Bahri et al., 2021).

Harnessing scaling laws for model design

Rosenfeld et al. (2020) performed small-scale experiments to approximate the generalization error of larger models with a functional form accounting for model and dataset sizes. However, they focused on pretraining, while we also investigate the predictive power on downstream language tasks.

Hashimoto (2021) used scaling laws to predict the optimal composition of a training set from dif-

ferent data sources. Kirstain et al. (2021) investigated the effect of parameter count and data size on improving performance on various language tasks, while Johnson et al. (2018) designed a performance extrapolation task to estimate how much training data is needed to achieve the required performance.

A parallel line of work that tries to extrapolate optimal architectures based on small scale experiments is *Neural Architecture Search* (Zoph and Le, 2017). Such methods have outperformed human designed architectures (Zoph and Le, 2017; Liu et al., 2019a; Chen et al., 2018; Real et al., 2019). However, it has been shown that the resulting architectures, such as EfficientNet (Tan and Le, 2019), do not always scale well (Bello et al., 2021).

6 Discussion and Future Work

In this work, we show that scaling laws can be used as an effective tool for model selection and as a diagnostic tool to test convergence of large-scale models. Our *practical takeaways* are:

1. Scaling laws can be used as an effective predictive and diagnostic tool, as long as the fit is good. Specifically, we find that $R^2 \geq 0.95$ is a good indicator.
2. Performing independent HPT at every scale is crucial for model selection and for the emergence of scaling laws in particular.
3. Pretraining models to convergence is important for observing the scaling behavior of transformer models over downstream NLP tasks.

7 Conclusion

This work is motivated by a practical question: can scaling laws provide a principled method for developing models at very small scale and extrapolating to larger ones? We perform a thorough empirical analysis on the emergence of scaling laws on a wide range of language understanding tasks. We find that scaling laws emerge for *some* tasks, potentially as a function of the proximity between the downstream task and the pretraining objective, but revealing them incurs the overhead of hyperparameter tuning across multiple scales. Our results show that scaling laws are beneficial for debugging model convergence when training large models, and to predict model behavior when they emerge at small scale.

8 Limitations

We discuss four limitations left for future work. (a) we focused on small-scale models, and thus have no empirical evidence for models that are significantly larger; (b) we analyze encoder-only modes, and leave decoder-based models for future work; (c) we analyze 9 different downstream tasks, they are all based on English-only datasets, and none of them evaluate models' generation capabilities. (d) while we provide a rule-of-thumb for telling if scaling laws predictions are reliable, it remains unclear why scaling laws do not apply to all downstream tasks, which remains an area for future research.

Acknowledgements

We thank Mor Geva for her useful comments. This research was partially supported by The Yandex Initiative for Machine Learning, the Shashua Fellowship, the Len Blavatnik and the Blavatnik Family foundation, the Israeli Science Foundation (ISF) grant no. 2486/21, and the European Research Council (ERC) under the European Union Horizons 2020 research and innovation programme (grant ERC DELPHI 802800). This work was completed in partial fulfillment for the Ph.D. degree of the first author.

References

- Samira Abnar, Mostafa Dehghani, Behnam Neyshabur, and Hanie Sedghi. 2021. [Exploring the limits of large scale pre-training](#). *ArXiv preprint*, abs/2110.02095.
- Yasaman Bahri, Ethan Dyer, Jared Kaplan, Jaehoon Lee, and Utkarsh Sharma. 2021. [Explaining neural scaling laws](#). *ArXiv preprint*, abs/2102.06701.
- Yamini Bansal, B. Ghorbani, Ankush Garg, Biao Zhang, Maxim Krikun, Colin Cherry, Behnam Neyshabur, and Orhan Firat. 2022. Data scaling laws in nmt: The effect of noise and architecture. *ArXiv*, abs/2202.01994.
- Irwan Bello, William Fedus, Xianzhi Du, Ekin Dogus Cubuk, A. Srinivas, Tsung-Yi Lin, Jonathon Shlens, and Barret Zoph. 2021. [Revisiting resnets: Improved training and scaling strategies](#). *ArXiv preprint*, abs/2103.07579.
- Shai Ben-Assayag and Ran El-Yaniv. 2021. [Train on small, play the large: Scaling up board games with alphazero and gnn](#). *ArXiv preprint*, abs/2107.08387.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*.
- Devansh Bisla, Apoorva Nandini Saridena, and Anna Choromańska. 2021. A theoretical-empirical approach to estimating sample complexity of DNNs. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 3264–3274.
- Blake Bordelon, Abdulkadir Canatar, and Cengiz Pehlivan. 2020. [Spectrum dependent learning curves in kernel regression and wide neural networks](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 1024–1034. PMLR.
- 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. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Liang-Chieh Chen, Maxwell D. Collins, Yukun Zhu, George Papandreou, Barret Zoph, Florian Schroff, Hartwig Adam, and Jonathon Shlens. 2018. [Searching for efficient multi-scale architectures for dense image prediction](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 8713–8724.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek B Rao, Parker Barnes, Yi Tay, Noam M. Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Benton C. Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier García, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Díaz, Orhan Firat, Michele Catasta, Jason Wei, Kathleen S. Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. Palm: Scaling language modeling with pathways. *ArXiv*, abs/2204.02311.

- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Comet.ML. 2021. [Comet.ML home page](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Jasha Droppo and Oguz Elibol. 2021. [Scaling laws for acoustic models](#). In *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association, Brno, Czechia, 30 August - 3 September 2021*, pages 2576–2580. ISCA.
- B. Ghorbani, Orhan Firat, Markus Freitag, Ankur Bapna, Maxim Krikun, Xavier García, Ciprian Chelba, and Colin Cherry. 2021. [Scaling laws for neural machine translation](#). *ArXiv preprint*, abs/2109.07740.
- Mitchell A Gordon, Kevin Duh, and Jared Kaplan. 2021. [Data and parameter scaling laws for neural machine translation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5915–5922, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. 2020. [Array programming with NumPy](#). *Nature*, 585(7825):357–362.
- Tatsunori Hashimoto. 2021. [Model performance scaling with multiple data sources](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 4107–4116. PMLR.
- T. J. Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo Jun, Tom B. Brown, Prafulla Dhariwal, Scott Gray, Chris Hallacy, Benjamin Mann, Alec Radford, Aditya Ramesh, Nick Ryder, Daniel M. Ziegler, John Schulman, Dario Amodei, and Sam McCandlish. 2020. [Scaling laws for autoregressive generative modeling](#). *ArXiv preprint*, abs/2010.14701.
- Danny Hernandez, Jared Kaplan, T. J. Henighan, and Sam McCandlish. 2021. [Scaling laws for transfer](#). *ArXiv preprint*, abs/2102.01293.
- Mark Johnson, Peter Anderson, Mark Dras, and Mark Steedman. 2018. [Predicting accuracy on large datasets from smaller pilot data](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 450–455, Melbourne, Australia. Association for Computational Linguistics.
- Andrew Jones. 2021. [Scaling scaling laws with board games](#). *ArXiv preprint*, abs/2104.03113.
- Jared Kaplan, Sam McCandlish, T. J. Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeff Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#). *ArXiv preprint*, abs/2001.08361.
- Yuval Kirstain, Patrick Lewis, Sebastian Riedel, and Omer Levy. 2021. [A few more examples may be worth billions of parameters](#). *ArXiv preprint*, abs/2110.04374.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2020. [Reformer: The efficient transformer](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [RACE: Large-scale ReAding comprehension dataset from examinations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.
- Yoav Levine, Barak Lenz, Opher Lieber, Omri Abend, Kevin Leyton-Brown, Moshe Tennenholtz, and Yoav Shoham. 2021. [Pmi-masking: Principled masking of correlated spans](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gungjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François

- Lagunas, Alexander Rush, and Thomas Wolf. 2021. [Datasets: A community library for natural language processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hanxiao Liu, Karen Simonyan, and Yiming Yang. 2019a. [DARTS: differentiable architecture search](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. [Roberta: A robustly optimized bert pretraining approach](#). *ArXiv preprint*, abs/1907.11692.
- The pandas development team. 2020. [pandas-dev/pandas: Pandas](#).
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8024–8035.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam M. Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *ArXiv preprint*, abs/1910.10683.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V. Le. 2019. [Regularized evolution for image classifier architecture search](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 4780–4789. AAAI Press.
- Jonathan S. Rosenfeld, Amir Rosenfeld, Yonatan Belinkov, and Nir Shavit. 2020. [A constructive prediction of the generalization error across scales](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Roy Schwartz, Jesse Dodge, Noah Smith, and Oren Etzioni. 2020. Green ai. *Communications of the ACM*, 63:54 – 63.
- Utkarsh Sharma and Jared Kaplan. 2022. Scaling laws from the data manifold dimension. *Journal of Machine Learning Research*, 23(9):1–34.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, and Nathan Srebro. 2018. [The implicit bias of gradient descent on separable data](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Stefano Spigler, Mario Geiger, and Matthieu Wyart. 2020. Asymptotic learning curves of kernel methods: empirical data versus teacher–student paradigm. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(12).
- Aarohi Srivastava, Abhinav Rastogi, Abhishek B Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W. Kocurek, Ali Safaya, Ali Tazarv, Alice Xiang, Alicia Parrish, Allen Nie, Aman Hussain, Amanda Askell, Amanda Dsouza, Ameet Annasaheb Rahane, Anantharaman S. Iyer, Anders Johan Andreassen, Andrea Santilli, Andreas Stuhlmüller, Andrew M. Dai, Andrew D. La, Andrew Kyle Lampinen, Andy

Zou, Angela Jiang, Angelica Chen, Anh Vuong, Animesh Gupta, Anna Gottardi, Antonio Norelli, Anu Venkatesh, Arash Gholamidavoodi, Arfa Tabasum, Arul Menezes, Arun Kirubarajan, Asher Mullokandov, Ashish Sabharwal, Austin Herrick, Avia Efrat, Aykut Erdem, Ayla Karakacs, Bridget R. Roberts, Bao Sheng Loe, Barret Zoph, Bartłomiej Bojanowski, Batuhan Ozyurt, Behnam Hedayatnia, Behnam Neyshabur, Benjamin Inden, Benno Stein, Berk Ekmekci, Bill Yuchen Lin, Blake Stephen Howald, Cameron Diao, Cameron Dour, Catherine Stinson, Cedrick Argueta, C'esar Ferri Ram'irez, Chandan Singh, Charles Rathkopf, Chenlin Meng, Chitta Baral, Chiyu Wu, Chris Callison-Burch, Chris Waites, Christian Voigt, Christopher D. Manning, Christopher Potts, Cindy Tatiana Ramirez, Clara Rivera, Clemencia Siro, Colin Raffel, Courtney Ashcraft, Cristina Garbacea, Damien Sileo, Daniel H Garrette, Dan Hendrycks, Dan Kilman, Dan Roth, Daniel Freeman, Daniel Khashabi, Daniel Levy, Daniel Gonz'alez, Danny Hernandez, Danqi Chen, Daphne Ippolito, Dar Gilboa, David Dohan, D. Drakard, David Jurgens, Debajyoti Datta, Deep Ganguli, Denis Emelin, Denis Kleyko, Deniz Yuret, Derek Chen, Derek Tam, Dieuwke Hupkes, Diganta Misra, Dilyar Buzan, Dimitri Coelho Mollo, Diyi Yang, Dong-Ho Lee, Ekaterina Shutova, Ekin Dogus Cubuk, Elad Segal, Eleanor Hagerman, Elizabeth Barnes, Elizabeth P. Donoway, Ellie Pavlick, Emanuele Rodolà, Emma FC Lam, Eric Chu, Eric Tang, Erkut Erdem, Ernie Chang, Ethan A. Chi, Ethan Dyer, Ethan Jerzak, Ethan Kim, Eunice Engefu Manyasi, Evgenii Zheltonozhskii, Fan Xia, Fatemeh Siar, Fernando Mart'inez-Plumed, Francesca Happ'e, François Chollet, Frieda Rong, Gaurav Mishra, Genta Indra Winata, Gerard de Melo, Germán Kruszewski, Giambattista Parascandolo, Giorgio Mariani, Gloria Wang, Gonzalo Jaimovitch-L'opez, Gregor Betz, Guy Gur-Ari, Hana Galijasevic, Han Sol Kim, Hannah Rashkin, Hanna Hajishirzi, Harsh Mehta, Hayden Bogar, Henry Shevlin, Hinrich Schütze, Hiromu Yakura, Hongming Zhang, Hubert Wong, Ian Aik-Soon Ng, Isaac Noble, Jaap Jumelet, Jack Geissinger, John Kernion, Jacob Hilton, Jaehoon Lee, Jaime Fernández Fisac, J. Brooker Simon, James Koppel, James Zheng, James Zou, Jan Koco'n, Jana Thompson, Jared Kaplan, Jarema Radom, Jascha Narain Sohl-Dickstein, Jason Phang, Jason Wei, Jason Yosinski, Jekaterina Novikova, Jelle Bosscher, Jenni Marsh, Jeremy Kim, Jeroen Taal, Jesse Engel, Jesujoba Oluwadara Alabi, Jiacheng Xu, Jiaming Song, Jillian Tang, Jane W Waweru, John Burden, John Miller, John U. Balis, Jonathan Berant, Jorg Frohberg, Jos Rozen, José Hernández-Orallo, Joseph Boudeman, Joseph Jones, Joshua B. Tenenbaum, Joshua S. Rule, Joyce Chua, Kamil Kanclerz, Karen Livescu, Karl Krauth, Karthik Gopalakrishnan, Katerina Ignatyeva, Katja Markert, Kaustubh D. Dhole, Kevin Gimpel, Kevin Ochieng' Omondi, Kory Wallace Mathewson, Kristen Chialfallo, Ksenia Shkaruta, Kumar Shridhar, Kyle McDonnell, Kyle Richardson, Laria Reynolds, Leo Gao, Li Zhang, Liam Dugan, Lianhui Qin, Lidia Contreras-

Ochando, Louis-Philippe Morency, Luca Moschella, Luca Lam, Lucy Noble, Ludwig Schmidt, Luheng He, Luis Oliveros Col'on, Luke Metz, Lutfi Kerem cSenel, Maarten Bosma, Maarten Sap, Maartje ter Hoeve, Madotto Andrea, Maheen Saleem Farooqi, Manaal Faruqi, Mantas Mazeika, Marco Baturan, Marco Marelli, Marco Maru, M Quintana, Marie Tolkiehn, Mario Giulianelli, Martha Lewis, Martin Potthast, Matthew Leavitt, Matthias Hagen, M'aty'as Schubert, Medina Baitemirova, Melissa Arnaud, Melvin Andrew McElrath, Michael A. Yee, Michael Cohen, Mi Gu, Michael I. Ivanitskiy, Michael Starritt, Michael Strube, Michal Swkedrowski, Michele Bevilacqua, Michihiro Yasunaga, Mihir Kale, Mike Cain, Mimee Xu, Mirac Suzgun, Monica Tiwari, Mohit Bansal, Moin Aminnaseri, Mor Geva, Mozhdah Gheini, T MukundVarma, Nanyun Peng, Nathan Chi, Nayeon Lee, Neta Gur-Ari Krakover, Nicholas Cameron, Nicholas S. Roberts, Nicholas Doiron, Nikita Nangia, Niklas Deckers, Niklas Muennighoff, Nitish Shirish Keskar, Niveditha Iyer, Noah Constant, Noah Fiedel, Nuan Wen, Oliver Zhang, Omar Agha, Omar Elbaghdadi, Omer Levy, Owain Evans, Pablo Antonio Moreno Casares, Parth Doshi, Pascale Fung, Paul Pu Liang, Paul Vicol, Pegah Alipoormolabashi, Peiyuan Liao, Percy Liang, Peter W. Chang, Peter Eckersley, Phu Mon Htut, Pi-Bei Hwang, P. Milkowski, Piyush S. Patil, Pouya Pezeshkpour, Priti Oli, Qiaozhu Mei, QING LYU, Qinlang Chen, Rabin Banjade, Rachel Etta Rudolph, Raefar Gabriel, Rahel Habacker, Ram'on Risco Delgado, Raphaël Millièvre, Rhythm Garg, Richard Barnes, Rif A. Saurous, Riku Arakawa, Robbe Raymaekers, Robert Frank, Rohan Sikand, Roman Novak, Roman Sitelew, Ronan Le Bras, Rosanne Liu, Rowan Jacobs, Rui Zhang, Ruslan Salakhutdinov, Ryan Chi, Ryan Lee, Ryan Stovall, Ryan Teehan, Rylan Yang, Sahib J. Singh, Saif M. Mohammad, Sajant Anand, Sam Dillavou, Sam Shleifer, Sam Wiseman, Samuel Gruetter, Sam Bowman, Samuel S. Schoenholz, Sanghyun Han, Sanjeev Kwatra, Sarah A. Rous, Sariik Ghazarian, Sayan Ghosh, Sean Casey, Sebastian Bischoff, Sebastian Gehrmann, Sebastian Schuster, Sepideh Sadeghi, Shadi S. Hamdan, Sharon Zhou, Shashank Srivastava, Sherry Shi, Shikhar Singh, Shima Asaadi, Shixiang Shane Gu, Shubh Pachchigar, Shubham Toshniwal, Shyam Upadhyay, Shyamolima Deb-nath, Siamak Shakeri, Simon Thormeyer, Simone Melzi, Siva Reddy, Sneha Priscilla Makini, Soohwan Lee, Spencer Bradley Torene, Sriharsha Hattwar, Stanislas Dehaene, Stefan Divic, Stefano Ermon, Stella Rose Biderman, Stephanie C. Lin, Stephen Prasad, Steven T. Piantadosi, Stuart M. Shieber, Summer Mishnerghi, Svetlana Kiritchenko, Swaroop Mishra, Tal Linzen, Tal Schuster, Tao Li, Tao Yu, Tariq A. Ali, Tatsuo Hashimoto, Te-Lin Wu, Theo Desbordes, Theodore Rothschild, Thomas Phan, Tianle Wang, Tiberius Nkinyili, Timo Schick, T. N. Kornev, Timothy Telleen-Lawton, Titus Tunduny, Tobias Gerstenberg, Trenton Chang, Trishala Neeraj, Tushar Khot, Tyler O. Shultz, Uri Shaham, Vedant Misra, Vera Demberg, Victoria Nyamai, Vikas Rau-

- nak, Vinay Venkatesh Ramasesh, Vinay Uday Prabhu, Vishakh Padmakumar, Vivek Srikumar, William Fedus, William Saunders, William Zhang, W Vossen, Xiang Ren, Xiaoyu F Tong, Xinyi Wu, Xudong Shen, Yadollah Yaghoobzadeh, Yair Lakretz, Yang Song, Yasaman Bahri, Ye Ji Choi, Yichi Yang, Yiding Hao, Yifu Chen, Yonatan Belinkov, Yu Hou, Yu Hou, Yushi Bai, Zachary Seid, Zhao Xinran, Zhuoye Zhao, Zi Fu Wang, Zijie J. Wang, Zirui Wang, Ziyi Wu, Sahib Singh, and Uri Shaham. 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *ArXiv*, abs/2206.04615.
- Mingxing Tan and Quoc V. Le. 2019. [Efficientnet: Re-thinking model scaling for convolutional neural networks](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 6105–6114. PMLR.
- Yi Tay, Mostafa Dehghani, Samira Abnar, Hyung Won Chung, William Fedus, Jinfeng Rao, Sharan Narang, Vinh Quang Tran, Dani Yogatama, and Donald Metzler. 2022. Scaling laws vs model architectures: How does inductive bias influence scaling? *ArXiv*, abs/2207.10551.
- Yi Tay, Mostafa Dehghani, Jinfeng Rao, William Fedus, Samira Abnar, Hyung Won Chung, Sharan Narang, Dani Yogatama, Ashish Vaswani, and Donald Metzler. 2021. [Scale efficiently: Insights from pre-training and fine-tuning transformers](#). *ArXiv preprint*, abs/2109.10686.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. [SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python](#). *Nature Methods*, 17:261–272.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 3261–3275.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019b. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. Emergent abilities of large language models. *ArXiv*, abs/2206.07682.
- Wes McKinney. 2010. [Data Structures for Statistical Computing in Python](#). In *Proceedings of the 9th Python in Science Conference*, pages 56 – 61.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. 2021. [Scaling vision transformers](#). *ArXiv preprint*, abs/2106.04560.
- Jie Zhao and Yong Deng. 2019. Performer selection in human reliability analysis: D numbers approach. *Int. J. Comput. Commun. Control*, 14:437–452.
- Ruiqi Zhong, Dhruva Ghosh, Dan Klein, and Jacob Steinhardt. 2021. [Are larger pretrained language models uniformly better? comparing performance at the instance level](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3813–3827, Online. Association for Computational Linguistics.

Yukun Zhu, Ryan Kiros, Richard S. Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 19–27. IEEE Computer Society.

Barret Zoph and Quoc V. Le. 2017. [Neural architecture search with reinforcement learning](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

A Appendix

A.1 Experimental details

All experiments were done with the *transformers* (Wolf et al., 2020) library (version 4.4.0) and tracked using the *Comet.ML* infrastructure (Comet.ML, 2021). All pretraining and finetuning datasets were provided by the *Datasets library* (Lhoest et al., 2021) (version 1.4.1) and were left as is, except for MNLI (Williams et al., 2018) where a random subset of 25% of the training samples were used to finetune the models. For pretraining, we used the Wikipedia dump provided by *dataset library* in the English dataset configuration name *20200501.en*. Whenever possible, example recipes provided by *transformers* were used to tune the models. We trained all models until convergence (discussed further in App. A.4) and chose the checkpoint that performed best over the evaluation set (i.e., post-hoc early stopping). To support the training and analysis of the results, we used *numpy* (Harris et al., 2020), *scipy* (Virtanen et al., 2020), *pandas* (Wes McKinney, 2010; pandas development team, 2020) and *scikit-learn* (Pedregosa et al., 2011). All models ran using *PyTorch* (Paszke et al., 2019).

The complete configuration of the different scales can be found in Table 6. In all cases, we focused on the emergence of scaling laws in NLP tasks rather than achieving optimal results, and thus did not perform any “mid-training” (Wang et al., 2019b). To account for randomness, during finetuning, we used five different seeds for each model-task pair. However, since pretraining is more expensive, we only have a single random seed during pretraining.

A.2 Hyperparameter tuning

Finetuning As discussed in §4.2, we observed that hyperparameter tuning during finetuning has considerable impact. To determine the hyperparameters to tune, we first experimented with modifying the different options in various scales and tasks. In particular, we examined the effects of *weight decay*, *batch size*, *number of epochs*, *initial learning rate*, *warm-up*, *learning rate scheduler* and *dropout*. We found that by fixing the *learning rate scheduler* to *CONSTANT*, the dropout rate and warm-up to 10% and the *batch size* to high enough (64), we are able to outperform other configurations by only varying the *learning rate* and *total number of epochs*. Then, we started performing a grid-search to choose those

AR	Layers	Heads	N
32	1	4	12,288
32	2	4	98,304
32	3	4	331,776
32	4	4	786,432
32	5	4	1,536,000
32	6	4	2,654,208
32	7	4	4,214,784
32	8	4	6,291,456
64	1	4	49,152
64	2	4	393,216
64	3	4	1,327,104
64	4	4	3,145,728
64	5	4	6,144,000
64	12	12	84,934,656

Table 6: Model configurations and number of parameters. The last line represents BERT-Base (Devlin et al., 2019).

Algorithm 1: Hierarchical bootstrap

```
input : performance results  $p_{i,j}$  where  $i \in [1, M]$ 
        and  $j \in [1, T]$ , number of samples  $B$ 
output :  $(\alpha_1, \beta_1), \dots, (\alpha_B, \beta_B)$  #  $b$  fitted lines
for  $b = 1, \dots, B$  do
    # Sampling uniformly with replacement
    scales  $\leftarrow$  sample  $M$  scales from  $[1, M]$ 
    points  $\leftarrow$  []
    for  $i \in$  scales do
        sp  $\leftarrow$  sample  $T$  points from  $p_{i,1}, \dots, p_{i,T}$ 
        points.extend(sp)
     $(\alpha_i, \beta_i) \leftarrow$  FitLine(points)
return  $(\alpha_1, \beta_1), \dots, (\alpha_B, \beta_B)$ 
```

values for each model-task pair. Figure. 4 shows the effect of this hyperparameter tuning.

Pretraining In the case of pretraining, we have found that by using a large enough *batch size* (256) and a small enough initial *learning rate* (10^{-4}) with long training horizon (500K steps) all models achieve comparable results at convergence to those achieved when hyperparameters were tuned (though the convergence rate differed). The rest of the hyperparameters used were a *linear learning rate scheduler*, 0.1 *dropout rate* and 10K *warmup* steps.

A.3 Evaluation

As discussed and is visible in Fig.3, there is variance in the finetuning performance resulting from the different seeds. Thus, we fit the power laws based on multiple random seeds to capture the **expected** performance of a scale. Moreover, we com-

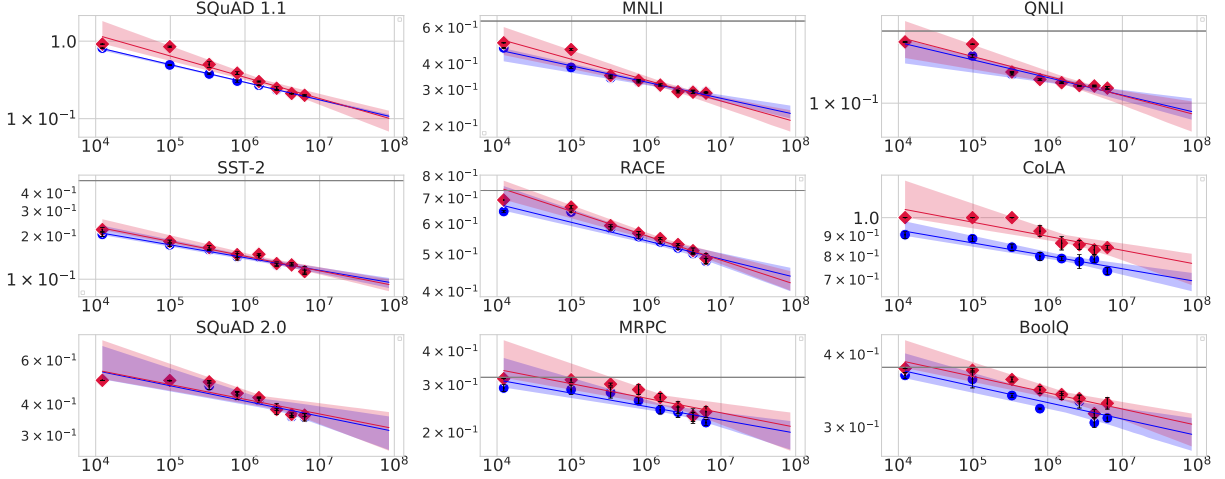


Figure 4: Effect of HPT on finetuning scaling laws (axes as in figure 3). **Blue:** with HPT. **Red:** w/o HPT.

pute the uncertainty of the fit in the form of a confidence interval. To do so, we use our data points to bootstrap $B = 1000$ fitted lines, and take the $[2.5, 97.5]$ percentiles. However, we find that the naïve bootstrapping approach in which one samples b points with replacements from a pool of size b (where $b := M \times T$) only accounts for the variance coming from the finetuning seeds. Since we expect more variance to come from different pre-training random seeds (as was shown by [Zhong et al. \(2021\)](#)), we use a hierarchical bootstrapping algorithm (Alg. 1) to compensate for this. Our hierarchical bootstrapping procedure works by first sampling M scales with replacement, and then for each scale, we sample T data points with replacement. While the number of observations in each fitted sample is the same, most samples will not include all scales and instead give more weight to specific scales when fitting a power law. When referring to the subroutine `FitLine(p_1, \dots, p_k)` where each point $p_i = (x_i, y_i)$, we fit a power law function such that $\ln(y) = \alpha \cdot \ln(x) + \beta$ and minimize the squared loss $\sum_i (\ln(y_i) - \ln(\hat{y}_i))^2$. When computing R^2 , we use the same procedure over all $M \times T$ points to fit a line, and compute the goodness-of-fit w.r.t. to its predictions. Specifically, given a fitted line $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, R^2 is defined as:

$$R^2 := 1 - \frac{\text{SS}_{\text{res}}}{\text{SS}_{\text{tot}}} \quad (2)$$

where

$$\text{SS}_{\text{res}} := \sum_i (y_i - f(x_i))^2 \quad (3)$$

$$\text{SS}_{\text{tot}} := \sum_i (y_i - \bar{y})^2 \quad (4)$$

such that $\bar{y} = \frac{1}{k} \sum_{i=1}^k y_i$.

As can be seen in Fig. 5, our approach results in considerably more conservative confidence intervals, where the only difference comes from pre-training variance. We further use this estimation to get the range (in the same percentiles) of the slope, and thus can measure how uncertain the prediction is.

A.4 Model convergence

We discussed the importance of controlling for variance and tuning hyperparameters when evaluating scaling laws. However, another potential pitfall is model under-training. While many past lines of work report that they “train until convergence”, they do not explicitly discuss the criteria for stopping training. In particular, when using a decaying learning rate schedule (e.g., *Linear*), the change in loss will tend to 0 as the learning rate approaches zero, which will give the appearance of convergence. Moreover, when using early-stopping, hyperparameters such as *patience* and *minimal decrease* may affect the final model and lead to under-optimization.

To optimize results, during finetuning we trained for the entire allocated epoch budget and chose post-hoc the best performing checkpoint w.r.t. the evaluation metric. However, as pretraining is considerably more compute-intensive, we used “early stopping”, requiring no decrease in evaluation loss over 1500 consecutive update steps. The only exception is BERT-Base ([Devlin et al., 2019](#)), where we increased the patience from 1500 to 7500 after suspecting the model may be under-trained. As can

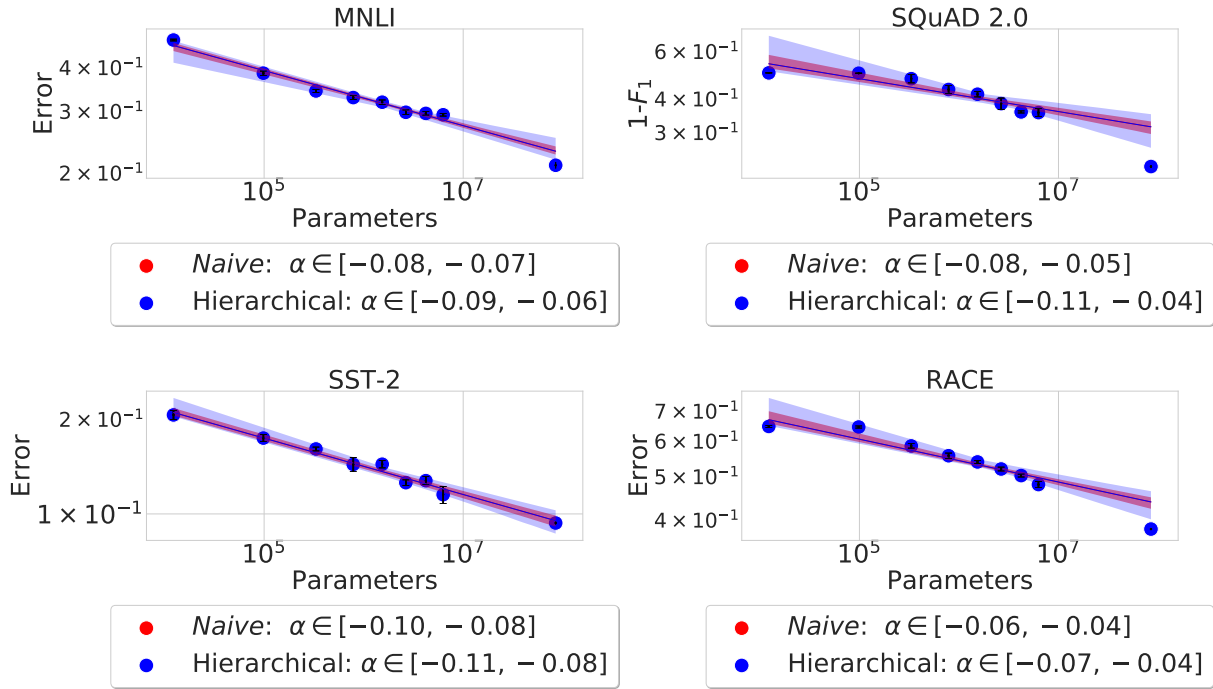


Figure 5: Comparison of our hierarchical bootstrapping method against the naive (i.e. “flat”) sampling approach.

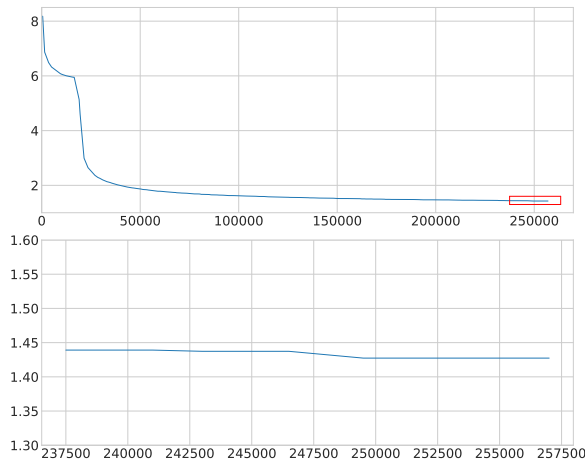


Figure 6: Evaluation loss of BERT-Base (Devlin et al., 2019) pretraining with MLM against the number of training steps. **Bottom:** Zoom-in on the final steps.

be seen in Fig. 6 which shows the evaluation loss of our BERT-Base over the pretraining data, there are many cases in which the model stops improving for a considerable amount of time (and indeed stops if the early stopping patience is set too low) despite not converging yet. When “zooming-out”, it is clear that the model is still training, and has not reached convergence. This supports our hypothesis that our large scale model is under-trained and accounts for some of the error in the predictions from Table 2 and Fig. 3.

A.5 PMI vs. MLM results

Table 7 shows the full set of scaling law results for the various models and tasks, comparing pretraining with MLM and PMI. While §4.4 discusses the potential benefits of using scaling laws as a method for model selection, we observe a slight disparity between our BERT-Base scores and the ones reported in literature. In particular, Levine et al. (2021) reported their BERT-Base models trained with PMI to achieve 81.4 *Best-F₁* score on SQuAD 2.0 and 70.1 accuracy on RACE when pretraining for 1M steps. We attribute the difference to the significantly more update steps taken (1M vs. our 250K) and usage of the *Book Corpus* dataset (Zhu et al., 2015) during pretraining. This conjecture is supported by the scores they present when pretraining an additional 1.4M steps (83.3 and 72.3 respectively) as well as increasing the pretraining corpora significantly (83.9 and 74.8 respectively).

A.6 Fitting Evaluation Loss

As mentioned in §4.1, the goodness-of-fit for evaluation loss can be quite poor, especially for classification tasks. This is since a single outlier that the model is wrong on with high confidence (and confidence tends to increase during training) can lead to high evaluation log-loss even when the target metric keeps improving; this is a known phenomenon (Soudry et al., 2018). As is visible from Fig. 7,

	R^2		Slope		Pred. (act.)		$R^2_{2:}$		$RE_{2:}$	
	MLM	PMI	MLM	PMI	MLM	PMI	MLM	PMI	MLM	PMI
SQD 1	1.00	0.99	[-0.23, -0.21]	[-0.25, -0.21]	0.89 (0.88)	0.91 (0.89)	1.00	0.99	0.01	0.02
MNLI	0.97	0.99	[-0.09, -0.06]	[-0.08, -0.07]	0.77 (0.79)	0.77 (0.79)	0.96	0.98	-0.04	-0.02
QNLI	0.96	0.96	[-0.20, -0.09]	[-0.20, -0.12]	0.92 (0.91)	0.92 (0.91)	0.88	0.89	0.01	0.02
SST-2	0.95	0.93	[-0.11, -0.08]	[-0.09, -0.07]	0.90 (0.91)	0.90 (0.91)	0.91	0.89	0.00	-0.01
RACE	0.92	0.90	[-0.07, -0.04]	[-0.08, -0.04]	0.56 (0.62)	0.57 (0.66)	0.97	0.96	-0.05	-0.04
CoLA	0.85	0.75	[-0.05, -0.02]	[-0.07, -0.02]	0.30 (0.37)	0.33 (0.33)	0.82	0.84	-0.12	-0.01
SQD 2	0.80	0.86	[-0.11, -0.04]	[-0.11, -0.05]	0.69 (0.78)	0.71 (0.79)	0.92	0.98	-0.08	-0.04
MRPC	0.76	0.68	[-0.09, -0.03]	[-0.10, -0.03]	0.80 (0.84)	0.80 (0.85)	0.85	0.77	-0.03	-0.02
BoolQ	0.75	0.60	[-0.04, -0.02]	[-0.03, -0.00]	0.71 (0.71)	0.69 (0.74)	0.67	0.37	0.00	-0.03

Table 7: Comparison of finetuning results for models pretrained with MLM/PMI. R^2 is the goodness-of-fit observed on the models with 1-8 layers and the *Slope* is the 95% confidence interval estimated for the slope of the fit. *Pred.* refer to the prediction of the BERT-Base performance based on the fit, and *act.* is the actual value observed. $R^2_{2:}$ and $RE_{2:}$ is the goodness-of-fit and *RE* (see §2.2) respectively, when fitting based only on models with 2-8 layers. SQD 1 and 2 refer to SQuAD 1.1 and 2.0 respectively.

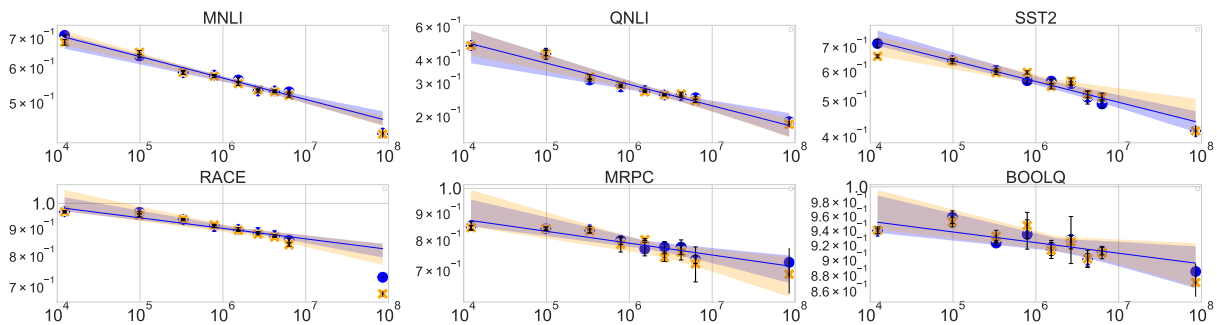


Figure 7: Evaluation loss w.r.t. number of parameters. The y-axis of each plot is $\frac{\mathcal{L}_{eval}}{-\ln(c)}$ where \mathcal{L}_{eval} is the best evaluation loss and c is the number of classes in each c -way classification task. **Blue**: models pretrained with MLM. **Orange**: models pretrained with PMI.

fitting a power-law on the best evaluation loss have no advantage over using the target metric. Specifically, when fitting the evaluation loss rather than evaluation accuracy causes an average drop of over 10% in goodness-of-fit as measured by R^2 .

A.7 SQuAD 2.0

In §4.1 we discussed the sub-optimal fit exhibited by SQuAD 2.0 compared to SQuAD 1.1. To test our hypothesis that this is due to the change in metric, we evaluated the models that were trained on SQuAD 2.0 using the SQuAD 1.1 objective. As expected, Fig. 8 shows they exhibit a much cleaner scaling law (compare red to orange line), with a slope similar to the one observed for SQuAD 1.1 (blue line). The slight difference in the intercept and uncertainty may be attributed to the difference in the dataset itself and the finetuning objectives. As mentioned, the goodness-of-fit measured by R^2 increased significantly as well, supporting our hypothesis.

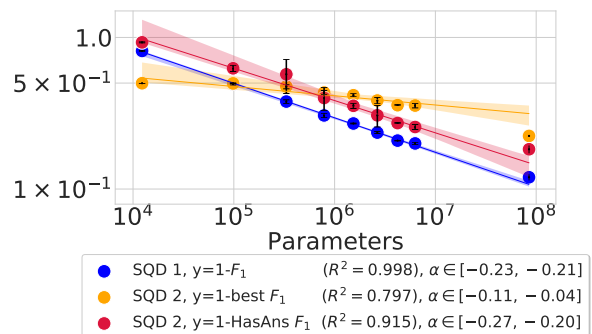


Figure 8: Effect of evaluation metric on scaling laws in SQuAD 1.1 and 2.0 (SQD 1 and SQD 2 respectively). SQuAD 2.0 models had their hyperparameters tuned to minimize $1 - \text{Best-}F_1$, and evaluated also with F_1 on the subset of answerable questions from SQuAD 1.1. ($1 - \text{HasAns } F_1$)