

Downstream Datasets Make Surprisingly Good Pretraining Corpora

Kundan Krishna, Saurabh Garg, Jeffrey P. Bigham, Zachary C. Lipton

Carnegie Mellon University

{kundank, sgarg2, jbigham, zlipton}@andrew.cmu.edu

Abstract

For most natural language processing tasks, the dominant practice is to finetune large pretrained transformer models (e.g., BERT) using smaller downstream datasets. Despite the success of this approach, it remains unclear to what extent these gains are attributable to the massive background corpora employed for pretraining versus to the pretraining objectives themselves. This paper introduces a large-scale study of *self-pretraining*, where the same (downstream) training data is used for both pretraining and finetuning. In experiments addressing both ELECTRA and RoBERTa models and 10 distinct downstream classification datasets, we observe that self-pretraining rivals standard pretraining on the BookWiki corpus (despite using around $10\times-500\times$ less data), outperforming the latter on 7 and 5 datasets, respectively. Surprisingly, these task-specific pretrained models often perform well on other tasks, including the GLUE benchmark. Self-pretraining also provides benefits on structured output prediction tasks such as question answering and common-sense inference, often providing more than 50% improvements compared to standard pretraining. Our results hint that often performance gains attributable to pretraining are driven primarily by the pretraining objective itself and are not always attributable to the use of external pretraining data in massive amounts. These findings are especially relevant in light of concerns about intellectual property and offensive content in web-scale pretraining data.¹

1 Introduction

For training predictive models operating on natural language data, the current best practice is to *pre-train* models on large unlabeled *upstream* corpora to optimize self-supervised objectives, for example, masked language modeling (MLM); the resulting weights are then used to initialize models that

are subsequently trained (*finetuned*) on the labeled *downstream* data available for the task at hand. Large-scale pretrained models typically provide significant performance boosts when compared to models trained directly on the downstream task (with random initializations) (Peters et al., 2018; Devlin et al., 2019; Chiang and Lee, 2020; Krishna et al., 2021). Upstream corpora tend to be significantly larger than the downstream corpora and the success of this approach is often attributed to its ability to leverage these massive upstream corpora (Liu et al., 2019; Yang et al., 2019). For example, the seminal BERT model (Devlin et al., 2019) was pretrained using the BookWiki corpus which is a combination of English Wikipedia and BooksCorpus (Zhu et al., 2015), totaling 13GB of plain text. Subsequent models have moved on to web-scale data. For example, XLNet (Yang et al., 2019), RoBERTa (Liu et al., 2019), and T5 (Raffel et al., 2020), were trained on 158GB, 160GB and 750GB of data, respectively.

As upstream corpus size and downstream performance have gone up, popular attempts at explaining these gains have focused on themes of “knowledge transfer” from the upstream corpus, attributing them to shared linguistic structure, semantics (Lina et al., 2019; Tenney et al., 2019), and facts about the world (Petroni et al., 2019). However, since the introduction of large-scale pretraining corpora occurred together with the invention of self-supervised pretraining objectives (e.g. masked language modeling (Devlin et al., 2019) and replaced token detection (Clark et al., 2019)), it remains unclear to what extent large-scale corpora are integral to these leaps in performance. For several tasks, especially summarization, recent works achieved surprising performance gains in settings where the upstream corpus is created synthetically with arbitrary symbols, but the pretraining objective is designed to capture some of the structure of the task (Krishna et al., 2021; Wu et al., 2022).

¹Pretrained models can be downloaded from <https://github.com/acmi-lab/self-pretrain>

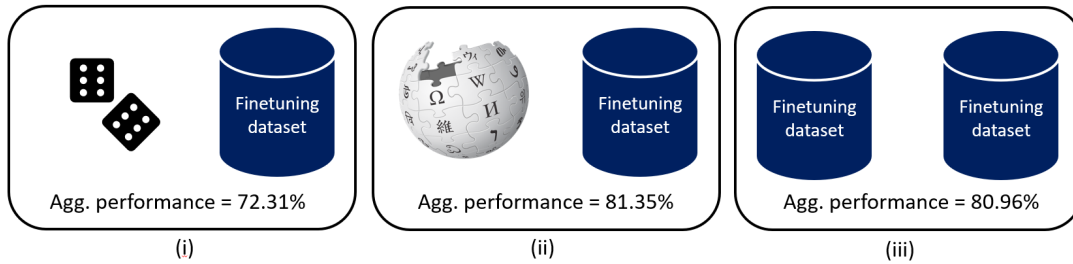


Figure 1: Aggregate performance of an ELECTRA model across 10 finetuning datasets when it is (i) randomly initialized (ii) pretrained on upstream corpus (BookWiki) (iii) pretrained on the finetuning dataset itself

In this work, we ask just how much of pretraining’s benefits could be realized in the absence of upstream corpora by pretraining directly on the downstream corpora (with the same self-supervised objectives). We find that this approach, which we call *self-pretraining*, often rivals the performance boosts conferred by *off-the-shelf* models pretrained on large upstream corpora (Figure 1), even outperforming them on 7 out of 10 datasets. Prior research has shown that *additional* self-supervised pretraining of off-the-shelf models using the downstream data can give further gains (Gururangan et al., 2020). Our study goes further, showing that even when starting from random initializations, and without using any external data beyond the downstream data itself, self-pretraining can rival standard practices. Since self-pretraining requires the same data that must already be available for finetuning, the benefits of pretraining in this case cannot be attributed to *transfer* of knowledge from the upstream corpus. Instead, these benefits can only be attributed to the pretraining objective, which is possibly able to learn some inductive biases better than the finetuning objective (e.g. linguistic knowledge (Tenney et al., 2019)), or perhaps simply initializes network parameters such that their statistics lead to better optimization during finetuning (Wu et al., 2022). While similar observations were made in the computer vision community (El-Nouby et al., 2021), we argue that it is especially important to establish these phenomena in the language domain, for which building on self-supervised pre-trained models is now a ubiquitous practice.

To understand differences in predictions with different pretraining strategies (i.e., between self-pretrained and off-the-shelf models), we analyse the errors made by these models on the same downstream data (Sec. 6). Despite similar performance of these models, we find that self-pretrained and off-the-shelf models make significantly less corre-

lated errors when compared to two independently finetuned models pretrained with either strategy.

We find that models pretrained on one downstream dataset often perform surprisingly well when finetuned to other downstream datasets (Sec. 5), including the GLUE benchmark. Even though the downstream datasets in our study come from a wide variety of domains (e.g., news, online forums, tweets), we find that pretraining on any of these downstream datasets delivers significant performance gains on most datasets (greater than half of off-the-shelf model’s gains in 88% of cases) irrespective of domain. However, the best performance on a downstream dataset is usually achieved by the model pretrained on that dataset itself.

In addition to classification tasks, we also experiment with tasks such as span-based question answering, named entity recognition, and grounded commonsense inference (Sec. 8). Self-pretraining delivers around 40-80% of the performance boost compared to models pretrained on the BookWiki corpus across ELECTRA and RoBERTa models. Hence, self-pretraining can perform better than finetuning randomly initialized models even for tasks that require prediction of more complex structured output than a single label, and for tasks whose solution relies on commonsense knowledge.

Our contributions can be summarized as follows:

- Comparison of self-pretrained and off-the-shelf pretrained models (both with ELECTRA and RoBERTa architectures) across 10 downstream classification tasks.
- Analysis of out-of-distribution performance of models pretrained on one downstream dataset and finetuned on other downstream datasets, including the GLUE benchmark.
- Demonstration of self-pretraining’s efficacy on more complex tasks than classification such as tasks requiring structured output prediction or commonsense reasoning.

2 Related Work

Self-pretraining in Computer Vision Most relevant to our work, recent/concurrent works in computer vision explore self-pretraining (He et al., 2022; El-Nouby et al., 2021). In a contemporary work, He et al. (2022) showed that pretraining with a Masked AutoEncoder (MAE) objective (analogue of MLM objective for images) boosts the performance of ViT models on the Imagenet-1K dataset. El-Nouby et al. (2021) showed that pretraining solely on downstream datasets for object detection and segmentation tasks reaches the performance of Imagenet-pretrained models. Our work establishes that a similar phenomenon is observed for NLP tasks too across a wide range of datasets.

Pretraining on Downstream Data in NLP *Task-Adaptive PreTraining* (TAPT (Gururangan et al., 2020)) consists of taking off-the-shelf pretrained models like BERT and RoBERTa and engaging in further pretraining on the downstream datasets before finetuning them to the task at hand. TAPT has been shown to improve performance of off-the-shelf models in a variety of works (Logeswaran et al., 2019; Han and Eisenstein, 2019; Chakrabarty et al., 2019). Another way in which downstream data has been used is for retrieval to create a small pretraining corpus for efficient pretraining (Yao et al., 2022). By contrast, our work pretrains models *only* on the downstream dataset, enabling a head-to-head comparison between the performance of off-the-shelf and self-pretrained models, and (in some situations) challenging the necessity of upstream corpora altogether.

Claims about Knowledge transfer Many works claim that pretraining extracts generally useful *knowledge* from the upstream corpus such as linguistic patterns (Lina et al., 2019; Tenney et al., 2019; Manning et al., 2020) and facts (Petroni et al., 2019), and that this accounts for the performance gains that they enjoy on downstream tasks. Several works, e.g., in the *probing* literature (Tenney et al., 2019; Manning et al., 2020; Petroni et al., 2019), demonstrate that from the internal representations of a model, it is easy (e.g., via linear models) to predict certain linguistic features or real-world facts. However, these studies do not clarify the mechanism by which these observations relate to performance gains on downstream tasks. Tenney et al. (2019) recognizes this limitation, stating “*the observation of a (linguistic) pattern does not tell us how it is used*”. Our work suggests that to the

extent that such knowledge extraction plays a role in pretraining’s benefits, sufficient knowledge is often present in the downstream dataset and need not be *transferred* from huge upstream corpora.

Challenges to the Knowledge Transfer Narrative Multiple previous works have questioned whether knowledge transfer can fully account for the efficacy of pretraining. Improvements in performance on downstream NLP tasks have resulted from pretraining on other modalities like music and code (Papadimitriou and Jurafsky, 2020), sequences of meaningless symbols (Chiang and Lee, 2020; Krishna et al., 2021; Wu et al., 2022), and language denatured via shuffling of words (Sinha et al., 2021). On the other hand, models pretrained on language have shown improved performance on tasks dealing with other modalities such as image classification (Lu et al., 2021) and reinforcement learning for games (Reid et al., 2022). By contrast, we show that without surplus upstream data of any modality, self-pretraining alone can often perform comparably or even better than standard pretraining with a large upstream corpus. In a similar vein with these papers, our work suggests that a large portion of pretraining’s success may come from alternative, unexplored mechanisms which have more to do with the pretraining objective than knowledge transfer from upstream corpora.

3 Experimental Setup

Our experiments center around the ELECTRA model (Clark et al., 2019) and the RoBERTa-base model (Liu et al., 2019). On the broadest set of experiments, for which we can only afford to train one model, we employ ELECTRA because it performs better than RoBERTa given comparable compute budgets (Clark et al., 2019). In particular, we use the small variant of ELECTRA (14 million parameters), which performs similarly to BERT-base on GLUE (difference of ≈ 2 points) while training much faster (Clark et al., 2019). However, we replicate many of these results on the larger RoBERTa-base model revealing similar results and thus establishing the generality of our findings.

During pretraining, a text sequence is fed into the model with some tokens masked out. While MLM-only models like RoBERTa only have a *generator* network that predicts the content of the masked tokens, ELECTRA has an additional discriminator module that predicts if those predictions were correct. Both the generator and the discriminator

Dataset	Size (MB)	Classes	Domain	Task
AGNews (Zhang et al., 2015)	27	4	News	topic classification
QQP (Wang et al., 2018)	43	2	Online forum questions	paraphrase detection
Jigsaw Toxicity (Kaggle.com, 2018)	59	6	Wikipedia comments	toxicity detection
MNLI (Williams et al., 2018)	65	3	Diverse	natural language inference
Sentiment140 (Go et al., 2009)	114	5	Tweets	sentiment classification
PAWS (Zhang et al., 2019)	139	2	Wikipedia	paraphrase detection
DBPedia14 (Zhang et al., 2015)	151	14	Wikipedia	topic classification
Discovery (Sileo et al., 2019)	293	174	Web crawl	discourse marker prediction
Yahoo Answerstotics (Zhang et al., 2015)	461	10	Online forum answers	topic classification
Amazon Polarity (Zhang et al., 2015)	1427	2	Product reviews	sentiment classification

Table 1: The suite of downstream datasets used in this work along with their training set sizes

networks’ parameters are updated simultaneously during pretraining. After pretraining, the generator is discarded and the discriminator is used as an encoder for finetuning on downstream tasks.

We experimented with 10 different downstream datasets (Table 1). We chose these datasets in our testbed to span different dataset sizes ranging from 27 megabytes to about 1.4 gigabytes of text in the training split. These datasets are for different tasks such as topic classification, sentiment classification, natural language inference etc., and are created using data sourced from diverse domains. Most of them are multi-class classification tasks except Jigsaw Toxicity which is a multi-label classification task, and Sentiment140 which is modeled as a regression task. For finetuning a pretrained model on any dataset, we passed the input through the model, took the vector representation of the CLS token in the final layer, and passed it through a classification head with one hidden layer to get the output.

4 Self-pretraining Performance

In our first set of experiments, we compare self-pretraining’s performance with other pretraining techniques. For each dataset, we pretrain an ELECTRA model on text from its training split and then finetune it on the same training data using the associated labels. To create a pretraining corpus from a downstream dataset, we concatenate the input text from each of the examples, assembling them in random order. We evaluate the performance of each finetuned model on the corresponding dataset’s test split. For QQP and MNLI we just use the validation split because test set labels are private. For all datasets, we evaluate performance by accuracy, except for Sentiment140 and Jigsaw Toxicity, for which we use Pearson correlation and micro-averaged AUC scores, respectively (these are not

multi-class classification problems).

Notably, all self-pretrained models deliver significant performance boosts on their respective datasets (Table 2), and over half of them perform even better than the off-the-shelf model. We measured a model’s *benefit* as the increase in performance metric that it achieves over a randomly initialized model, divided by the increase in performance metric achieved by the off-the-shelf ELECTRA model against the same baseline. The average benefit of self-pretraining across all datasets is 103.70%. We do not see a clear correlation between the size of the dataset and the performance of self-pretraining. For example, the highest benefit of 131.33% is achieved for the smallest dataset (AGNews), which is merely 27MB in size, while the minimum benefit is achieved on the Discovery dataset, which is the third largest dataset measuring 293MB. For each downstream dataset, we also pretrain a model on a randomly sampled subset of Wikipedia of the same size as the dataset’s training corpus, and finetune it on the downstream task. This approach (called WikiSub) provides a size-adjusted comparison between using separate upstream data vs the downstream data for pretraining. We see that self-pretraining performs better than WikiSub in the majority of cases (Table 2).

We also evaluated the alternate pretraining technique *TAPT* as described in Gururangan et al. (2020). In this technique, we take the off-the-shelf ELECTRA model, which has already been pretrained on the upstream BookWiki corpus, and further pretrain it on the downstream dataset for 100 epochs. Self-pretraining outperforms TAPT on 6 datasets, notably including the two datasets where it outperformed the off-the-shelf models by the greatest benefit margin - *AGNews* and *Yahoo Answerstotics*. Interestingly, TAPT performs worse

than off-the-shelf model on the same 3 datasets where self-pretraining performs worse than off-the-shelf model (except Sentiment140). None of the three pretraining approaches seem to be uniformly better than any other.

Finally, we also evaluate the self-pretrained models on the GLUE benchmark and report results on the dev set ². The performance of the models on their pretraining dataset does not correlate strongly with its GLUE score. The GLUE score also does not monotonically go up with increasing dataset size, indicating that the data domain makes some difference. For example, the Amazon Polarity corpus scores just 66.14 on GLUE despite being about 1.4GB in size, while AGNews which is 27MB in size, scores 74.30. The highest GLUE score is achieved by pretraining on Yahoo Answertopics.

5 Cross Dataset Finetuning

In this set of experiments, we investigated if the models pretrained on a dataset are only useful for that specific task, or are they useful across the whole spectrum of tasks that we consider. We took each model pretrained on a dataset in our testbed and finetuned and evaluated it on all other datasets in the testbed. The performance benefits provided in all cases are shown as a heatmap in Figure 2.

We found that for almost all downstream datasets, pretraining on any other dataset provides significant advantage (Figure 2). In most cases, pretraining on the downstream dataset itself performs the best. Among datasets where self-pretraining performs better than off-the-shelf model (i.e. the diagonal entry is greater than 1), pretraining on datasets of larger size does not help further. However, for the datasets where self-pretraining’s benefit is much less than 100% (i.e. MNLI and Discovery), pretraining on a larger dataset (e.g., Yahoo Answertopics) performs better than self-pretraining.

Among all the pretrained models, a few models perform consistently good or bad across different downstream datasets (Figure 2). For example, the model pretrained on Yahoo Answertopics gets the highest average score of 0.90 across all datasets, while the PAWS-pretrained model gives the lowest aggregate score of 0.64. Similarly, there are downstream datasets that are benefited consistently by either a large or a small margin by pretraining on different datasets. For example, performance on

QQP and PAWS receives huge boosts by pretraining on most datasets. In contrast, performance on sentiment140 is mostly low, even dropping below 20% for 3 pretrained models.

We perform an ablation to investigate that given a fixed dataset to finetune on, is it better to pretrain on the *exact* same data (i.e., using the same set of inputs), or is it better to pretrain on different data with an identical distribution. To test this hypothesis, we divided the training splits of the downstream datasets randomly into two equal subsets (denoted as A and B). We pretrained one model on each subset and then finetuned them on both subsets separately. The validation and test sets used for finetuning are the same as in the original dataset.

We do not see any consistent benefits with pretraining and finetuning on the same dataset (Table 3). Instead, we found consistent patterns where models pretrained on one split (either A or B) outperformed models pretrained on the other, irrespective of the split used for finetuning. This suggests that the pretraining data has greater influence on the final performance than the finetuning data. Additionally, we observe that finetuning the superior pretrained model, using the downstream split other than the one used for pretraining, performs the best, suggesting overall exposure to more data helps.

6 Difference in Outputs of Self-pretrained and Off-the-shelf Models

Since self-pretrained models and off-the-shelf models perform similarly in terms of classification accuracy, a natural question to ask is: *do these models make errors on the same set of inputs?* To answer this question, we investigate the difference in predictions made by models pretrained with different strategies across all multi-class classification tasks. In particular, given model f_A and f_B , we compute *error inconsistency*, defined as follows:

$$\sum_{i=1}^n \frac{\mathbb{1}[f_A(x_i) \neq y_i \wedge f_B(x_i) = y_i]}{n} + \frac{\mathbb{1}[f_A(x_i) = y_i \wedge f_B(x_i) \neq y_i]}{n},$$

where $\{x_i, y_i\}_{i=1}^n$ is the test set. Intuitively, error inconsistency captures the fraction of examples where exactly one model is correct. This definition has been commonly used to estimate diversity in model prediction (Gontijo-Lopes et al., 2022; Geirhos et al., 2020). Across all the multi-class

²Following Clark et al. (2019) we exclude the WNLI task from the results.

Dataset	Size(MB)	RandInit	SelfPretrain	Offshelf	Benefit%	WikiSub	TAPT	GLUE
AGNews	27	91.75	94.34	93.72	131.33	93.51	94.07	74.30
QQP	43	82.93	90.66	90.34	104.34	89.16	90.64	75.43
Jigsaw Toxicity	59	97.83	98.49	98.53	94.99	98.35	98.48	76.65
MNLI	65	65.49	78.39	82.29	76.77	78.64	79.26	78.28
Sentiment140	114	63.75	67.04	66.95	102.91	65.52	65.65	72.67
PAWS	139	50.00	97.53	97.30	100.49	97.42	97.85	74.65
DBpedia14	151	98.59	99.22	99.11	121.17	99.18	99.23	70.38
Discovery	293	17.00	22.38	24.55	71.22	22.47	23.58	77.26
Yahoo Answersttopics	461	61.94	65.26	64.55	127.31	64.37	65.05	79.53
Amazon Polarity	1427	93.86	96.27	96.13	106.49	95.82	96.16	66.14

Table 2: Performance of ELECTRA-small models pretrained with different techniques on various downstream datasets. We also report their performance on the GLUE benchmark (dev set). For reference, a randomly initialized model scores 53.20 and the off-the-shelf model scores 79.43 on GLUE.

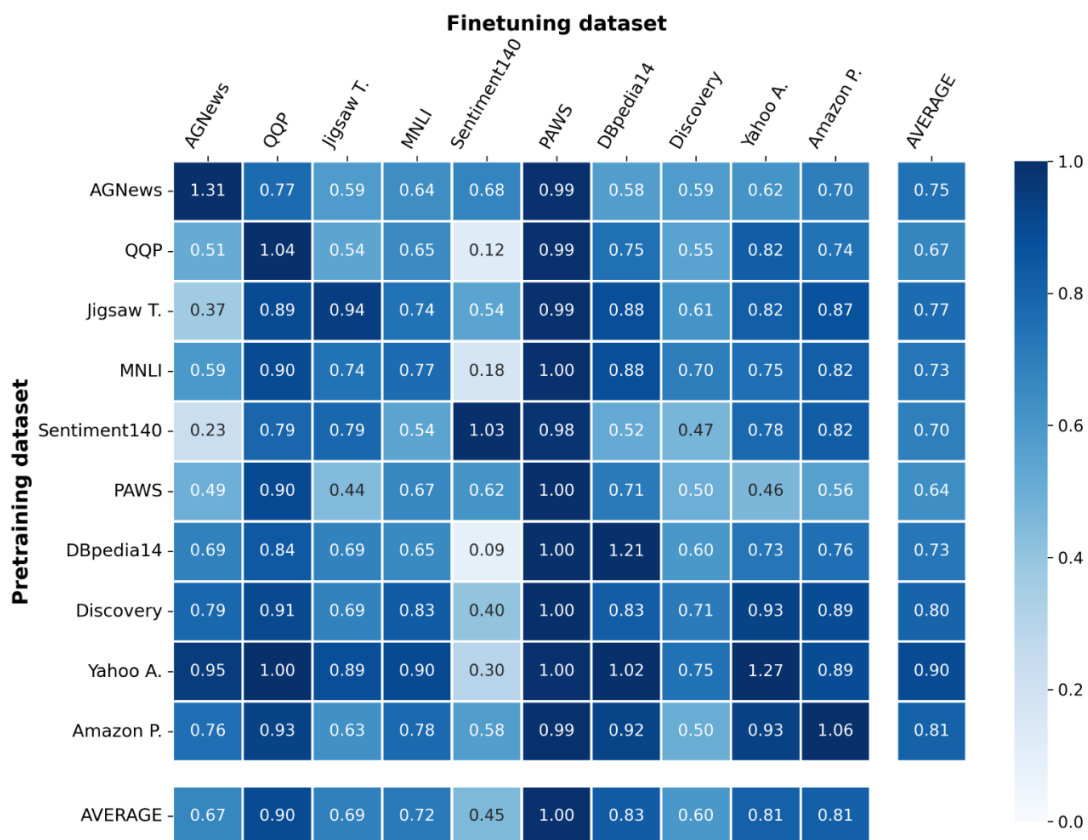


Figure 2: Performance benefits of models pretrained on each dataset, upon finetuning on each downstream dataset. Each value is the ratio of performance gains achieved by model pretrained on the row’s dataset vs off-the-shelf model, relative to random initialization, upon finetuning on the column’s dataset.

classification tasks, in addition to computing error inconsistency between self-pretrained and off-the-shelf model, for baseline comparison, we also tabulate error inconsistency between: (i) two independently finetuned versions of a self-pretrained model; and (ii) two independently finetuned versions of the off-the-shelf model.

Compared to error inconsistency between two models with the same pretraining dataset, we ob-

serve that models trained with different pretraining datasets have high error inconsistency in predictions (Table 4). For models with comparative performance, high error inconsistency highlights the high disagreement in predictions. This demonstrates that while different pretraining datasets produce similarly performing models in terms of overall accuracy, the model predictions are relatively dissimilar. Our observations here align with inves-

	MNLI		QQP			Discovery			Yahoo Answertopics		
	A	B	A	A	B	A	B	A	A	B	
A	76.00	76.42	A	84.28	84.79	A	18.78	18.61	A	64.18	64.34
B	75.93	75.05	B	88.73	88.41	B	19.99	19.98	B	64.09	64.18

Table 3: Performance when splitting the dataset into two equal-sized subsets A and B and then pretraining on one (row) and finetuning on another (column)

tigations in vision tasks, where Gontijo-Lopes et al. (2022) observed that models trained with different pretraining datasets produced uncorrelated errors.

Since different pretraining datasets produce models with uncorrelated errors, we ensemble these models to check if uncorrelated mistakes lead to a correct prediction. When the models make different predictions, in particular, when one model is correct and another is incorrect, the ensemble prediction will be dominated by the model with higher confidence in their prediction. As before, we consider ensembles of (i) two independently finetuned self-pretrained models; (ii) two independently finetuned off-the-shelf models; and (iii) a finetuned version, each of the self-pretrained and off-the-shelf models.

We make the following observations: First, as expected we observe that ensembling improves model performance as compared to a single model (Table 4). Second, despite having larger error inconsistency, we do not observe any significant improvements in ensembles of self-pretrained and off-the-shelf model as compared to ensembles of two models with the same pretraining strategy (Table 4). This is in contrast with findings on vision tasks where Gontijo-Lopes et al. (2022) observed that larger error inconsistency led to larger improvement in ensemble performance.

7 Ablations with Other Pretraining Architectures

We conducted our experiments so far with ELECTRA-small architecture because it is faster to pretrain than other popular models, yet delivers good downstream performance (Clark et al., 2019) (e.g. comparable to BERT-base on GLUE benchmark). Here, we conduct experiments with a larger model and a different pretraining objective to test the efficacy of self-pretraining more broadly.

We experiment with the RoBERTa model which uses the masked language modeling objective, rather than ELECTRA’s objective. We use the RoBERTa-base architecture, which has a much

larger parameter count of 110 million, compared to ELECTRA-small’s 14 million. Due to resource constraints, we pretrained the RoBERTa models for fewer iterations as in Warstadt et al. (2020). We pretrain a RoBERTa-base model on the BookWiki corpus for the same number of iterations. Our results show that self-pretraining performs comparably to pretraining on BookWiki corpus, delivering over 85% of pretraining benefit on 9 out of 10 datasets, and outperforming the model pretrained on BookWiki corpus (Table 5) on 5 datasets.

8 Performance on Structured Prediction and Commonsense NLI

While the bulk of our experiments were on a variety of classification tasks, we also experiment with some tasks beyond simple classification. We experiment with three types of tasks: (i) span based question answering, (ii) named entity recognition (NER), and (iii) grounded commonsense inference. For question answering we use the SQuAD dataset (Rajpurkar et al., 2016) (v1.1) and report the F1-score. For NER, we use the CONLL-2012 NER task which uses annotations from Ontonotes v5.0 (Weischedel et al., 2013) involving 18 kinds of named entities. To measure performance, we use the overall F1 score. We use the seqeval library for evaluation (<https://github.com/chakki-works/seqeval>). We include SWAG (Zellers et al., 2018) and HellaSwag (Zellers et al., 2019) for multiple-choice sentence completion.

For Electra-small models, we see that for each of these datasets self-pretraining achieves more than 70% pretraining benefit, and for RoBERTa-base model the benefit is 40-80% (Table 6). Even for the SWAG and HellaSwag datasets, which are designed to use rely on *commonsense inference* of pretrained models, we see performance boosts by pretraining using only the task’s training set.

Dataset	Ensemble Accuracy			Error Inconsistency		
	2×SelfPretrain	2×Offshelf	SelfPretrain + Offshelf	2×SelfPretrain	2×Offshelf	SelfPretrain + Offshelf
AGNews	94.66	94.17	94.54	1.76	3.50	4.01
QQP	90.92	90.74	91.63	4.57	5.27	8.91
MNLI	78.51	82.37	82.31	6.94	6.42	14.82
PAWS	97.70	97.45	97.75	0.96	1.30	2.07
DBPedia14	99.28	99.19	99.24	0.38	0.48	0.51
Discovery	22.98	25.25	25.02	7.85	9.18	12.66
Yahoo	65.32	64.69	65.64	5.27	5.49	9.55
Amazon	96.40	96.24	96.51	1.26	1.58	2.48

Table 4: Performance of ensemble models of self-pretrained and off-the-shelf models. For ensembling, we aggregate predictions of models after calibration with Temperature Scaling (Guo et al., 2017). We observe that in most of the datasets, SelfPretrain + Off-the-shelf ensembling does not improve over ensembles of two models with the same pre-training strategy, despite relatively higher error inconsistency of SelfPretrain + Off-the-shelf models.

Dataset	RandInit	SelfPretrain	BookWiki	Benefit %	TAPT
AGNews	91.91	94.28	94.22	102.27	94.07
QQP	76.50	88.68	90.18	89.05	90.64
Jigsaw Toxicity	97.32	97.72	98.03	56.02	98.48
MNLI	31.82	75.12	80.90	88.23	79.26
Sentiment140	56.68	68.55	60.19	338.26	65.65
PAWS	50.00	97.34	97.08	100.55	97.85
DBPedia14	98.57	99.21	99.24	95.98	99.23
Discovery	17.36	25.85	26.30	94.91	23.58
Yahoo Answerstoppers	61.11	65.96	64.58	139.80	65.05
Amazon Polarity	89.02	96.68	96.11	108.13	96.16

Table 5: Performance of RoBERTa-base models pretrained with different techniques on downstream datasets.

Datasets	Size(MB)	ELECTRA-small				RoBERTa-base			
		RI	SP	OS	Benefit%	RI	SP	BW	Benefit%
SQuAD	19	15.82	63.01	75.96	78.47	14.93	67.23	81.89	78.11
SWAG	22	27.55	60.56	73.76	71.43	27.95	45.18	70.37	40.62
HellaSwag	30	29.27	39.14	42.91	72.36	24.53	31.03	34.28	66.67
CONLL-2012	6.4	54.49	75.66	82.65	75.18	63.65	72.64	86.25	39.78

Table 6: Performance of ELECTRA and RoBERTa models pretrained with different techniques. RI: random initialization, SP: self-pretraining, OS: off-the-shelf; BW: pretrained on BookWiki by us.

9 Conclusion and Future Work

In this work, we showed that pretraining models only on text from the downstream dataset performs comparably to pretraining on a huge upstream corpus for a wide variety of datasets. The errors made by such *self-pretrained* models on the downstream tasks are significantly different from the ones made by the *off-the-shelf* models pretrained on upstream corpora. Our results suggest that the importance of learning from surplus upstream data for improving downstream task performance may have been overestimated. Crucially, our experiments also do

not show that upstream data does not help at all or knowledge transfer does not occur, but simply questions to what extent it is responsible for downstream gains. For example, the impressive zero-shot performance very large language models such as GPT-3 (Brown et al., 2020) clearly suggests knowledge transfer is involved. One direction of future work would be to investigate how the performance of self-pretraining compares of pretraining on upstream corpora as the model sizes go up by orders of magnitude.

We found that the quantity and quality of data

required for pretraining to provide significant benefit (over a randomly initialized model trained only with a supervised loss) is quite low. Downstream datasets which are tiny in comparison to typical upstream corpora, still function as useful pretraining corpora for getting performance gains across a wide range of datasets.

Since self-pretraining does not involve any upstream corpus, it prevents exposure of the model to potentially undesirable contents in the large upstream corpus, while still delivering large performance benefits. Research has demonstrated the negative influence of web-sourced pretraining corpora on models, such as generating toxic language (Gehman et al., 2020) or reflecting racial biases in predictions (Ahn and Oh, 2021). For use cases that require avoiding such issues, self-pretraining can provide a viable alternative to standard pretraining. In future work, we hope to compare how self-pretrained models and off-the-shelf models perform on these negative measures such as toxicity and social biases.

10 Limitations

Due to limited availability of compute resources, we were unable to scale up the model architecture to the large sizes becoming increasingly mainstream today. Similarly, the upstream corpus we used (BookWiki) is 16GB in size, and while it is large enough such that it was used to pretrain BERT (Devlin et al., 2019), much larger pretraining datasets are in use today such as the Colossal Common Crawl Corpus (Raffel et al., 2020). The relative performance achieved by using self-pretraining vs pretraining on upstream corpus can likely vary with the size of the model and upstream corpus, and more compute-heavy large scale experiments are needed to characterize it.

11 Acknowledgements

This work was funded by UPMC and Abridge AI Inc. We also gratefully acknowledge support from Amazon AI, the PwC Center, the Software Engineering Institute, and NSF (Award no. 2211955) for the compute resources used in this project. SG acknowledges Amazon Graduate Fellowship and JP Morgan AI Ph.D. Fellowship for their support.

References

- Jaimeen Ahn and Alice Oh. 2021. Mitigating language-dependent ethnic bias in bert. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 533–549.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Tuhin Chakrabarty, Christopher Hidey, and Kathleen McKeown. 2019. Imho fine-tuning improves claim detection. In *Proceedings of NAACL-HLT*, pages 558–563.
- Cheng-Han Chiang and Hung-yi Lee. 2020. Pre-training a language model without human language. *arXiv preprint arXiv:2012.11995*.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2019. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*.
- 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.
- Alaaeldin El-Nouby, Gautier Izacard, Hugo Touvron, Ivan Laptev, Hervé Jegou, and Edouard Grave. 2021. Are large-scale datasets necessary for self-supervised pre-training? *arXiv preprint arXiv:2112.10740*.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A Smith. 2020. Realtocixityprompts: Evaluating neural toxic degeneration in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369.
- Robert Geirhos, Kristof Meding, and Felix A Wichmann. 2020. Beyond accuracy: quantifying trial-by-trial behaviour of cnns and humans by measuring error consistency. *Advances in Neural Information Processing Systems*, 33:13890–13902.
- Alec Go, Richa Bhayani, and Lei Huang. 2009. Twitter sentiment classification using distant supervision. *CS224N project report, Stanford*, 1(12):2009.
- Raphael Gontijo-Lopes, Yann Dauphin, and Ekin D Cubuk. 2022. No one representation to rule them all: Overlapping features of training methods. In *International Conference on Learning Representations (ICLR)*.

- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith. 2020. Don’t stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360.
- Xiaochuang Han and Jacob Eisenstein. 2019. Unsupervised domain adaptation of contextualized embeddings for sequence labeling. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4238–4248.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009.
- Kaggle.com. 2018. Toxic comment classification challenge: Identify and classify toxic online comments. <https://www.kaggle.com/c/jigsaw-toxic-comment-classification-challenge>.
- Kundan Krishna, Jeffrey P Bigham, and Zachary C Lipton. 2021. Does pretraining for summarization require knowledge transfer? In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3178–3189.
- Yoav Levine, Noam Wies, Daniel Jannai, Dan Navon, Yedid Hoshen, and Amnon Shashua. 2021. The inductive bias of in-context learning: Rethinking pretraining example design. In *International Conference on Learning Representations*.
- Yongjie Lina, Yi Chern Tana, and Robert Frankb. 2019. Open sesame: Getting inside bert’s linguistic knowledge. *ACL 2019*, page 241.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Lajanugen Logeswaran, Ming-Wei Chang, Kenton Lee, Kristina Toutanova, Jacob Devlin, and Honglak Lee. 2019. Zero-shot entity linking by reading entity descriptions. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3449–3460.
- Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Kevin Lu, Aditya Grover, Pieter Abbeel, and Igor Mordatch. 2021. Pretrained transformers as universal computation engines. *arXiv preprint arXiv:2103.05247*.
- Christopher D Manning, Kevin Clark, John Hewitt, Urvasi Khandelwal, and Omer Levy. 2020. Emergent linguistic structure in artificial neural networks trained by self-supervision. *Proceedings of the National Academy of Sciences*, 117(48):30046–30054.
- Isabel Papadimitriou and Dan Jurafsky. 2020. Learning music helps you read: Using transfer to study linguistic structure in language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6829–6839.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. **Deep contextualized word representations**. 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 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- 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.
- Machel Reid, Yutaro Yamada, and Shixiang Shane Gu. 2022. Can wikipedia help offline reinforcement learning? *arXiv preprint arXiv:2201.12122*.
- Damien Sileo, Tim Van De Cruys, Camille Pradel, and Philippe Muller. 2019. **Mining discourse markers for unsupervised sentence representation learning**. 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 3477–3486, Minneapolis, Minnesota. Association for Computational Linguistics.
- Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela. 2021.

- Masked language modeling and the distributional hypothesis: Order word matters pre-training for little. *arXiv preprint arXiv:2104.06644*.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. Bert rediscovers the classical nlp pipeline. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355.
- Alex Warstadt, Yian Zhang, Xiaocheng Li, Haokun Liu, and Samuel Bowman. 2020. Learning which features matter: Roberta acquires a preference for linguistic generalizations (eventually). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 217–235.
- Ralph Weischedel, Martha Palmer, Mitchell Marcus, Eduard Hovy, Sameer Pradhan, Lance Ramshaw, Nianwen Xue, Ann Taylor, Jeff Kaufman, Michelle Franchini, et al. 2013. Ontonotes release 5.0 ldc2013t19. *Linguistic Data Consortium, Philadelphia, PA*, 23.
- 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.
- Yuhuai Wu, Felix Li, and Percy Liang. 2022. Insights into pre-training via simpler synthetic tasks. *arXiv preprint arXiv:2206.10139*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.
- Xingcheng Yao, Yanan Zheng, Xiaocong Yang, and Zhilin Yang. 2022. Nlp from scratch without large-scale pretraining: A simple and efficient framework. In *International Conference on Machine Learning*, pages 25438–25451. PMLR.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. Swag: A large-scale adversarial dataset for grounded commonsense inference. In *EMNLP*.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *NIPS*.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. Paws: Paraphrase adversaries from word scrambling. 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 1298–1308.
- Yukun Zhu, Ryan Kiros, Rich 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 *Proceedings of the IEEE international conference on computer vision*, pages 19–27.

A Appendix

A.1 The Role of Sentence Order in Pretraining Corpora

For virtually all pretrained models like BERT, ELECTRA, XLNet, the sentences in the pretraining corpora are ordered as they naturally occur in some document such as Wikipedia article. Devlin et al. (2019) mention in their work : “*It is critical to use a document-level corpus rather than a shuffled sentence-level corpus (...) in order to extract long contiguous sequences.*” However, for many of our pretraining corpora made from downstream datasets, the sentence taken in order do not form a coherent document or narrative text. For example, in the MNLI or QQP corpora, neighboring sentences will simply be premise-hypothesis pairs or potential paraphrase candidates.

Despite the sentence order not forming a coherent document, many pretraining corpora achieve high performance boosts on the GLUE language understanding benchmark (Table 7). For example, MNLI achieves around 96% of the performance boost of the off-the-shelf model (Table 7). Interestingly, shuffling the sentences in these corpora leads to a large drop in performance (Table 7). This suggests that there is some value to keeping the sentence order in a way that puts sentences from the same example in datasets like MNLI and QQP next to each other. A likely explanation of this is in Levine et al. (2021) where authors showed that including similar sentences in the same input sequence when pretraining should lead to improved performance via theoretical analysis and empirical experiments.

We test if GLUE performance can be improved by artificially re-ordering a set of sentences to promote the occurrence of similar sentences together. We rearrange the sentences in the sentence-shuffled versions of pretraining corpora to encourage content overlap among neighboring sentences, and see if this can recover some of the drops in performance that occurred due to shuffling. Our algorithm creates the corpus by iteratively appending sentences to it, such that at each step the new sentence is the one with maximum TF-IDF similarity with the previous sentence. Such a way of constructing a corpus by similarity based retrieval has been used in past works (Levine et al., 2021; Yao et al., 2022), with the main difference that they retrieved sentences from external corpora similar to the ones present in the downstream dataset, whereas we sim-

ply use it to reorder sentences already present in the downstream dataset for pretraining. We also make sure that the algorithm does not accidentally recover the original order of sentences (e.g. by matching the premise-hypothesis pairs originally in the MNLI dataset).

We experiment with 5 different datasets and find that the sentence-reordering scheme improves performance compared to random sentence order for all of them except QQP. For Discovery and DBPedia14 datasets, it scores even higher than our *standard* sentence ordering scheme which preserves the adjacency and order of sentences within each datapoint. This shows that re-ordering sentences to promote content similarity between neighboring sentences, can potentially improve GLUE score, without introducing any new information or narrative structure.

A.2 Experiments with Smaller ELECTRA Models

In addition to experimenting with a *base*-sized architecture (110M parameters), we also experiment with architectures which are even smaller than ELECTRA-small. We train ELECTRA models of smaller size by either reducing the number of layers in the generator and discriminator, or reducing the hidden dimension of the discriminator³. As the models get smaller, self-pretraining continues to significantly outperform random initialization and often outperforms pretraining on BookWiki corpus (Figure 3). Interestingly, the relative performance of self-pretrained and BookWiki-pretrained models tends to stay the same across model size. For example, for QQP self-pretraining is always best and for MNLI BookWiki-pretraining is always best irrespective of number of layers or hidden size.

A.3 Implementation Details for Pretraining and Finetuning

Hyperparameters for Pretraining For pretraining ELECTRA-small models, we use the standard hyperparameters (Table 8) as described in Clark et al. (2019). For the Roberta-base models, training with the standard hyperparameters with our computing resources would be prohibitively slow, and so we used hyperparameters from Warstadt et al. (2020) which require lesser time to train (Ta-

³In ELECTRA, the generator’s hidden size is already much smaller than that of the discriminator by design. So we do not reduce it further, in order to have a reasonably well-performing generator.

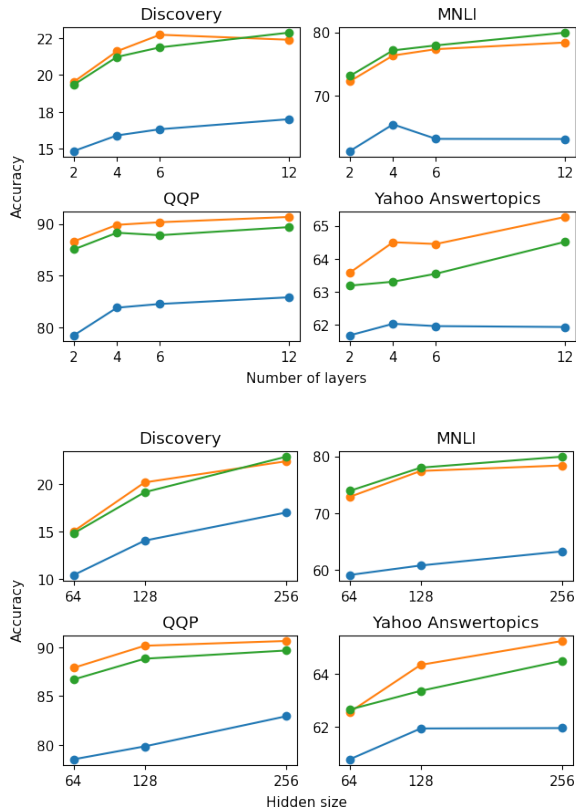


Figure 3: Variation in performance of ELECTRA models with change in number of layers and hidden size (— randomly initialized, — self-pretrained, — BookWiki-pretrained)

ble 8). For task-adaptive pretraining (TAPT), we follow Gururangan et al. (2020) and further pre-train off-the-shelf models for 100 epochs on the downstream task’s training set, with the first 6% of the resulting total updates used for learning rate warmup.

Hyperparameters for Finetuning For finetuning the models on the 10 downstream datasets, we use hyperparameters as shown in Table 9. We use the AdamW optimizer (Loshchilov and Hutter, 2018) for finetuning. We use early stopping based on validation set performance. The validation metric used is mean squared error for the sentiment140 dataset (regression), average binary crossentropy for the jigsaw dataset (multi-label classification), and accuracy for all other datasets (multi-class classification). The patience parameter for early stopping is set to 3 epochs. For finetuning ELECTRA-small models on the GLUE datasets, we use the standard learning rate of $1e-4$ following Clark et al. (2019).

Details about Use of Downstream Datasets All downstream datasets used in this paper were

sourced from the Huggingface library⁴. For the Yahoo Answertopics dataset, we use only the text from the answer (not the question) as input to the models (both for pretraining and finetuning). For the PAWS dataset, we use the version called “Unlabeled PAWS_{wiki}” in Zhang et al. (2019), which is actually *not* unlabeled but has silver labels. We preferred that version over others because of its larger size. For datasets which had a train and test split but no validation split (e.g. Yahoo Answertopics), we extracted 5000 random datapoints from the the train split to make the validation split. If a dataset had a train and validation split but no test split (e.g. Unlabeled PAWS_{wiki}), we designated the validation split to be the test split, and created a new validation set by extracting 5000 random datapoints from the train set.

A.4 Hardware and Software Packages Used

For pretraining ELECTRA models, we used Nvidia’s implementation of the ELECTRA code-base⁵, run using Nvidia’s Tensorflow cotainer image 21.07⁶. For pretraining Roberta models, we used the official implementation in the Fairseq library⁷. For finetuning experiments, we used the AllenNLP library for training and evaluation routines, coupled with the Huggingface library for the model architectures.

We used a collection of Nvidia V100 (32GB) and A6000(48GB) GPUs for our experiments. Pretraining an ELECTRA-small model takes around 1.5 days on 2 GPUs while pretraining a Roberta-base model takes around 1.5 days on 4 GPUs.

⁴<https://huggingface.co/docs/datasets/index>

⁵<https://github.com/NVIDIA/DeepLearningExamples/tree/master/TensorFlow2/LanguageModeling/ELECTRA>

⁶https://docs.nvidia.com/deeplearning/frameworks/tensorflow-release-notes/rel_21-07.html

⁷<https://github.com/facebookresearch/fairseq>

Pretraining Dataset	Random	Standard	TF-IDF(Ours)
None (RandomInit)	-	53.20	-
Sentiment140	-	72.67	75.29
DBpedia14	72.82	70.38	75.44
Discovery	71.79	77.26	78.94
MNLI	62.80	78.28	76.33
QQP	71.09	75.43	69.57
BookWiki (Off-the-shelf)	-	79.43	-

Table 7: GLUE scores achieved by different strategies for ordering sentences from the downstream dataset used for pretraining. Random: randomly ordered sentences; Standard: sentences within a datapoint occur contiguously in original order; TF-IDF: sentences reordered using content similarity.

Hyperparameter	ELECTRA	Roberta
Size (Parameter count)	Small (14M)	Base (110M)
Training steps	1M	100K
Warmup steps	10K	6K
Batch size	128	512
Peak learning rate	5e-4	5e-4
Sequence length	128	512

Table 8: Hyperparameters used for pretraining models

Hyperparameter	ELECTRA	Roberta
Training epochs	20	20
Batch size	32	32
Learning rate	{1e-4,1e-5}	2e-5
Max sequence length	512	512

Table 9: Hyperparameters used for finetuning models on 10 downstream tasks

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 10
- A2. Did you discuss any potential risks of your work?
Our method actually reduces the risks of pretraining which arise by the typical use large unscrutinized pretraining corpora, by relying on only the downstream dataset for pretraining.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section 1 summarizes the claims.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 3

- B1. Did you cite the creators of artifacts you used?
Section 3, Section 8, Table 1
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. All the models and datasets used are popular public resources made available for NLP research.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. All the models and datasets used are popular public resources made available for NLP research.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Not applicable. Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Table 1, Section A.3 in Appendix

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

C Did you run computational experiments?

Section 3 describes the experimental setup and Sections 4-8 describe results from those experiments.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

Section A.2 in appendix, Section A.4 in appendix, Table 8 in appendix

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Appendix section A.3

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

We use a single run to report results because pretraining is expensive. However, consistent trends across different datasets (more than 10 in our study) and two different pretraining approaches (Roberta and ELECTRA), do show the significance in our analysis

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section A.3, A.4 in appendix, Section 8

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Not applicable. Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Not applicable. Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Not applicable. Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Not applicable. Left blank.