

Sharpness-Aware Minimization Improves Language Model Generalization

Dara Bahri, Hossein Mobahi, Yi Tay

Google Research

{dbahri, hmobahi, yitay}@google.com

Abstract

The allure of superhuman-level capabilities has led to considerable interest in language models like GPT-3 and T5, wherein the research has, by and large, revolved around new model architectures, training tasks, and loss objectives, along with substantial engineering efforts to scale up model capacity and dataset size. Comparatively little work has been done to improve the generalization of these models through better optimization. In this work, we show that Sharpness-Aware Minimization (SAM), a recently proposed optimization procedure that encourages convergence to flatter minima, can substantially improve the generalization of language models without much computational overhead. We show that SAM is able to boost performance on SuperGLUE, GLUE, Web Questions, Natural Questions, Trivia QA, and TyDiQA, with particularly large gains when training data for these tasks is limited.

1 Introduction

Over the last several years, remarkable progress has been made within the domain of natural language understanding, with machine-learned models able to solve some tasks at near or above human-level performance. This progress has, by and large, been fueled by research centered around 1) better *inductive biases*, such as the attention-enabled Transformer architecture (Vaswani et al., 2017), 2) the clever leverage of massive corpora of textual data that was historically disregarded as “unlabeled,” usually in the form of pre-training objectives that strive to teach the model the structure of language (Radford et al., 2019; Devlin et al., 2018), 3) scaling up model capacity and the methods to support it (Shazeer and Stern, 2018), 4) multi-task learning (Raffel et al., 2019), and lastly, 5) larger and more diverse datasets along with ever-improving benchmarks that attempt to test the true capabilities of these models. Although these efforts

all share the single goal of improving the model’s generalization, doing so by explicit changes to the optimization of the loss function has received less attention in comparison.

Recently, motivated by the both empirical and theoretical findings that flatter minima lead to better generalization (Kleinberg et al., 2018; Shirish Keskar et al., 2016; Chaudhari et al., 2019; Smith and Le, 2017), Foret et al. (2020) proposed a novel modification to vanilla stochastic gradient descent they term “Sharpness-Aware Minimization,” or SAM. They show theoretically and empirically that optimizing with SAM encourages convergence to flatter points in the loss landscape and with it comes the anticipated improvement in out-of-sample error. While their empirical findings are limited to computer vision tasks and datasets using convolutional neural networks (ResNets), follow-up work (Chen et al., 2021) showed how SAM is particularly effective on Vision transformers (ViTs) (Dosovitskiy et al., 2020) and MLP-Mixers (Tolstikhin et al., 2021), architectures that are more prone than convolutional ones to land in sharp minima. Crucially, they show that when equipped with SAM, ViTs outperform ResNets of similar size and throughput *without the need for large-scale pre-training*.

Encouraged by wins in the vision domain, we ask whether SAM can deliver similar gains in the language domain. Our contributions are as follows:

1. We show that blithely applying SAM when fine-tuning public pre-trained checkpoints of the text-to-text transformer (T5) (Raffel et al., 2019) and its multilingual counterpart, mT5 (Xue et al., 2020) on SuperGLUE (Wang et al., 2019), GLUE (Wang et al., 2018), TyDiQA-GoldP (Clark et al., 2020) and the Closed-Book Question Answering (CBQA) tasks from Roberts et al. (2020) – Web Questions (Berant et al., 2013), Natural Questions (Kwiatkowski et al., 2019), and Trivia QA (Joshi et al., 2017) – improves test perfor-

mance quite markedly. Furthermore, by employing an approximation suggested by Brock et al. (2021), these gains come only at the cost of about 25% extra compute.

2. The improvement brought by SAM often increases with less labeled training data, making SAM indispensable for data-limited tasks. We test this by subsampling the training splits of CBQA and SuperGLUE datasets at rates ranging from 2% to 80%.

2 Related Works

Better Generalization. In light of flatter minima generalizing better, Smith and Le (2017) showed that the inherent noise in SGD serves as a form of *implicit regularization*, preventing the optimization from ever entering sharp valleys. Like SAM, entropy SGD (Chaudhari et al., 2019) *explicitly* encourages flatter minima. Smith et al. (2021); Barrett and Dherin (2020) analyzed SGD’s generalization formally by way of continuous-time gradient flow. Optimization routines based on adversarial risk (Zhu et al., 2019; He et al., 2020) and trust regions (Jiang et al., 2019; Aghajanyan et al., 2020) have been proposed and shown to improve generalization across settings.

While the number of methods which provide implicit or explicit regularization is overwhelmingly large, methods like early stopping, weight decay (or ℓ_2 -regularization), dropout (Srivastava et al., 2014), teacher-student or self-distillation (Hinton et al., 2015; Mobahi et al., 2020), label smoothing (Müller et al., 2019), batch normalization (Ioffe and Szegedy, 2015), mixup (Zhang et al., 2017), and data-augmentation more broadly are among the most widely used in practice. Marginalization of Bayesian neural networks, though challenging, has been shown to result in superior generalization in some settings (Wilson and Izmailov, 2020; MacKay, 1995).

While first-order optimization via SGD has been the prevailing way of training neural networks due to its efficiency and effectiveness, alternative second-order methods like K-FAC (Martens and Grosse, 2015) and Shampoo (Gupta et al., 2018) have slowly gained traction, often enabled by clever engineering to make them feasible at scale. Notably, Anil et al. (2020) presents a scalable implementation of Shampoo that provides significant convergence and wall-clock time improvements compared to first-order methods. They demonstrate

superior performance on machine translation and language modeling.

SAM. While this is, to the best of our knowledge, the first work detailing the benefits of SAM for language tasks, there have been successful applications of SAM in the vision domain. Notably, Chen et al. (2021) showed that convolution-free vision models like vision transformers (ViTs) (Dosovitskiy et al., 2020) and MLP-Mixers (Tolstikhin et al., 2021) suffer from sharp minima and that SAM indeed smooths their loss landscapes. They crucially show that ViTs and MLP-Mixers outperform ResNets of similar and greater size on ImageNet *without* the use of pre-training or data augmentations that would otherwise be necessary to achieve reasonable performance. They show that SAM induces sparsity in both architectures and leads to more perceptive attention maps in ViTs. They observe empirically that data augmentation and SAM are alike in that they both smooth the landscape on average, but the latter does so by explicitly controlling the worst-case curvature, whereas the former smooths over the directions induced by the augmentations. Furthermore, they observe that SAM encourages linearity with respect to the input, exhibiting an effect similar to that of mixup (Zhang et al., 2017). Lastly, they show that SAM helps contrastive learning and that it enables better robustness on corrupted examples from ImageNet-C (Hendrycks and Dietterich, 2019) and ImageNet-R (Hendrycks et al., 2021).

In a similar spirit, Brock et al. (2021) proposed speeding up SAM significantly by using fewer examples when computing the ascent step, a strategy which we employ in this work, and they were able to apply it to ResNet model variants to advance the state of the art on ImageNet without extra data.

Meanwhile, in an attempt to make SAM’s radius ρ invariant to the scale of the model parameters, Kwon et al. (2021) proposed an adaptive version named Adaptive Sharpness-Aware Minimization (ASAM), which they then show empirically to outperform normal SAM on a set of benchmark vision tasks.

3 Review of Sharpness-Aware Minimization (SAM)

We begin by briefly reviewing the SAM algorithm; interested readers can see the original paper for a thorough treatment. In our presentation, we use the ℓ_2 norm ($p = 2$ using notation from the origi-

nal paper), assume a general optimizer (instead of vanilla SGD), and use the approximation proposed by Brock et al. (2021) to compute the ascent gradient (adversarial point) efficiently. Given a loss function $L : \mathcal{W} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$, SAM seeks to find the parameter w whose neighborhood has low training loss by optimizing the minimax objective:

$$\min_w \max_{\|\epsilon\|_2 \leq \rho} L_{\text{train}}(w + \epsilon).$$

Finding the exact optima ϵ^* of the inner-maximization is challenging, so Foret et al. (2020) employ a first-order approximation, resulting in:

$$\begin{aligned} \hat{\epsilon}(w) &= \underset{\|\epsilon\|_2 \leq \rho}{\operatorname{argmin}} L_{\text{train}}(w) + \epsilon^T \nabla_w L_{\text{train}}(w) \\ &= \rho \nabla_w L_{\text{train}}(w) / \|\nabla_w L_{\text{train}}(w)\|_2. \end{aligned}$$

That is, $\hat{\epsilon}$ is just a scaling of the loss gradient at the current parameters. After computing $\hat{\epsilon}(w)$, SAM performs gradient descent using the gradient $\nabla_w L_{\text{train}}(w)|_{w_{\text{adv}}}$ at the nearby ‘‘adversarial’’ point $w_{\text{adv}}(w) \triangleq w + \hat{\epsilon}(w)$.

Put another way, SAM plugs-and-plays with any first-order optimizer by simply replacing the gradient of the mini-batch \mathcal{B} at the current model weights $w_t \in \mathcal{W}$ with the gradient computed at w_{adv} . w_{adv} itself is computed by taking a gradient ascent step of size ρ along the unit gradient vector $\nabla_w L_{\mathcal{M}}(w) / \|\nabla_w L_{\mathcal{M}}(w)\|_2|_{w_t}$, where \mathcal{M} can be the mini-batch \mathcal{B} , or a subset of it for enhanced efficiency. We found that setting \mathcal{M} to be 1/4-th of \mathcal{B} sped up the method significantly with little loss in quality, in line with the recommendation of Brock et al. (2021). The end-to-end algorithm is outlined in Algorithm 1.

4 Experiments

With SAM reviewed, we now discuss our experiments. We evaluate SAM on a range of natural language understanding tasks using the T5 (text-to-text Transformer) framework (Raffel et al., 2019). T5 casts NLU tasks as sequence-to-sequence ones that are learned using an encoder-decoder Transformer (Vaswani et al., 2017) architecture setup. These Transformer models are typically pre-trained on large corpora, like the Colossal Clean Crawled Corpus (C4) (Raffel et al., 2019), with, for example, the objective of predicting a short contiguous span of text that was intentionally corrupted in a snippet of input text. The pre-trained model is typically fine-tuned on a single task or a mixture of

Algorithm 1 Efficient SAM Algorithm.

- 1: **input:** training set $\mathcal{S} \triangleq \cup_{i=1}^n \{(x_i, y_i)\}$, loss function $L : \mathcal{W} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$, batch size b , neighborhood size $\rho > 0$ (default 0.15), ascent micro-batch size $a \leq b$ (default $b/4$), first-order optimizer update $\text{opt} : \mathcal{W} \times \mathcal{W} \rightarrow \mathcal{W}$.
 - 2: initialize parameters $w_0, t = 0$.
 - 3: **while** not converged **do**
 - 4: sample batch $\mathcal{B} = \{(x_1, y_1), \dots, (x_b, y_b)\}$.
 - 5: sample ascent micro-batch $\mathcal{M} = \{(x_1, y_1), \dots, (x_a, y_a)\}$.
 - 6: compute adversarial (ascent) point: $w_{\text{adv}} = w_t + \rho \frac{\nabla_w L_{\mathcal{M}}(w)}{\|\nabla_w L_{\mathcal{M}}(w)\|_2}|_{w_t}$.
 - 7: compute gradient approximation for the SAM objective: $g_{\text{adv}} = \nabla_w L_{\mathcal{B}}(w)|_{w_{\text{adv}}}$.
 - 8: update parameters: $w_{t+1} = \text{opt}(w_t, g_{\text{adv}})$.
 - 9: $t = t + 1$.
 - 10: **end while**
 - 11: **return** w_t
-

multiple tasks, the latter enabled by the fact that the framework treats all tasks as simple input-to-target sequence predictions.

To this end, we evaluate SAM in two ways:

1. When publicly available pre-trained checkpoints of the T5.1.1 model variant are fine-tuned with and without SAM, on SuperGLUE, GLUE, TyDiQA, and the Closed-Book Question Answering benchmarks: Web Questions, Natural Questions, TriviaQA. We show SAM improves generalization across benchmarks and four model sizes: Small (77M parameters), Base (250M), Large (800M), and XL (3B).
2. To show how it helps when task data is limited, we report results when the training splits of these benchmarks at various rates, ranging from 2% to 80%.

4.1 Setup

Framework. For all experiments, we train using Jax (Bradbury et al., 2018) and Google Cloud TPUs. To ensure fair comparisons, eliminate the impact of exogenous factors, and reduce the possibility of software bugs, we train both standard and SAM-enabled models using the same codebase and settings, so that the code paths are identical except for the gradient calculation at each step, wherein

Model	SGLue	BoolQ	CB	CoPA	MultiRC	ReCoRD	RTE	WiC	WSC
Small	67.7	72.6	89.4 / 89.3	67.0	68.5 / 21.4	61.7 / 60.8	69.3	65.4	72.1
Small + SAM (0.05)	68.4	73.5	92.1 / 89.3	61.0	68.5 / 22.8	62.1 / 61.0	69.7	65.7	79.8
Base	75.3	80.0	91.7 / 94.6	71.0	75.4 / 35.4	76.2 / 75.4	80.9	69.3	76.9
Base + SAM (0.15)	78.5	82.2	93.7 / 94.6	78.0	77.5 / 39.1	78.2 / 77.2	85.9	70.4	81.7
Large	84.3	86.6	99.4 / 98.2	89.0	83.7 / 51.0	86.5 / 85.6	89.2	72.9	84.6
Large + SAM (0.15)	84.6	88.0	95.0 / 96.4	86.0	84.0 / 53.7	87.3 / 86.4	89.2	75.2	86.5
XL	87.2	88.6	93.7 / 96.4	95.0	86.9 / 61.1	89.5 / 88.4	91.3	74.9	89.4
XL + SAM (0.15)	89.1	89.4	100.0 / 100.0	95.0	87.9 / 63.7	90.9 / 90.0	92.1	75.5	94.2

Table 1: Experimental results (dev scores) on the (full) SuperGLUE benchmark. Public checkpoints of various sizes are fine-tuned with and without SAM for 250k steps. We see that SAM improves performance across *all* model sizes.

Model	Glue	CoLA	SST	MRPC	STSB	QQP	MNLI	QNLI	RTE
Small	79.8	39.9	92.3	90.1 / 85.8	85.4 / 84.9	87.3 / 90.5	80.9 / 81.4	87.6	74.0
Small + SAM (0.05)	79.9	42.9	92.1	90.9 / 87.3	85.5 / 85.3	87.6 / 90.7	81.0 / 81.4	87.7	70.4
Base	84.7	50.9	94.2	91.5 / 88.2	88.5 / 88.3	88.5 / 91.4	87.3 / 87.6	92.1	81.6
Base + SAM (0.15)	85.1	49.8	94.2	93.4 / 90.7	90.0 / 89.7	89.0 / 91.7	87.3 / 87.5	92.5	82.7
Large	86.2	57.8	95.3	89.5 / 85.3	88.3 / 88.5	87.8 / 90.9	88.9 / 89.0	93.4	85.9
Large + SAM (0.15)	88.3	65.6	95.6	93.6 / 91.2	90.2 / 89.8	89.7 / 92.2	89.9 / 89.7	94.5	85.9
XL	89.9	70.7	96.3	92.8 / 90.0	91.2 / 91.0	89.6 / 92.1	91.1 / 91.3	95.5	90.6
XL + SAM (0.15)	90.2	69.7	96.9	93.2 / 90.7	91.7 / 91.8	90.2 / 92.7	91.4 / 91.5	95.8	91.0

Table 2: Experimental results (dev scores) on the (full) GLUE benchmark. Public checkpoints of various sizes are fine-tuned with and without SAM on the mixture of tasks for 250k steps. We see that SAM improves performance across *all* model sizes.

Model	avg. (F1/EM)
Small	73.4 / 62.1
Small + SAM (0.02)	74.3 / 63.0
Base	81.6 / 71.0
Base + SAM (0.02)	82.0 / 71.3
Large	85.6 / 75.3
Large + SAM (0.02)	85.9 / 76.1
XL	87.0 / 77.4
XL + SAM (0.05)	87.3 / 77.7

Table 3: Average results on TyDiQA-GoldP. Public checkpoints of mT5 of various sizes are fine-tuned on TyDiQA-GoldP for 20k steps. SAM boosts performance across model sizes here as well. We found that smaller values of ρ than those used for the English-only T5 model were necessary to achieve good performance here. Full, per-language results are shown in the Appendix.

SAM behaves differently. Our implementation of SAM is an adaptation of an existing open-source

implementation¹ to fit our framework for training language models.

Efficient SAM. In Foret et al. (2020), the idea of partitioning the ascent mini-batch into m disjoint micro-batches and computing a distinct adversarial point for each micro-batch and then averaging the SAM-gradients at each of these points was proposed under the name m -sharpness. It was noted there and in follow-up work (Chen et al., 2021) that $m > 1$ can result in better performance. This modification incurs m -times more compute under a naive sequential implementation (though it can be parallelized well if multiple devices are available).

Meanwhile, Brock et al. (2021) suggests (in the Appendix) using roughly 20% of the examples from the mini-batch for computing the adversarial point, observing little loss in model quality. With $m = 1$, this approximation roughly reduces SAM’s relative runtime from 2x to 1.2x. Since we understand how a $2m$ -x slow-down of model training may be prohibitive or significantly deter SAM’s widespread adoption, we, at the possible

¹<https://github.com/google-research/sam>

Model	Natural Q.	Web Q.	TriviaQA
Small	16.7 / 12.4	22.8 / 16.5	10.2 / 7.3
Small + SAM (0.05)	17.5 / 13.1	23.5 / 16.9	11.0 / 7.8
Base	23.2 / 18.1	29.7 / 22.5	19.3 / 15.3
Base + SAM (0.15)	25.7 / 20.6	31.0 / 24.5	21.5 / 17.4
Large	27.4 / 22.3	34.3 / 27.6	25.2 / 20.9
Large + SAM (0.15)	30.6 / 25.0	36.4 / 29.6	28.5 / 24.2
XL	33.5 / 27.5	39.3 / 31.6	36.5 / 31.1
XL + SAM (0.15)	34.7 / 28.8	40.7 / 33.3	38.0 / 32.6

Table 4: Experimental results (F1/EM) (test scores) on the (full) CBQA tasks. Public checkpoints of various sizes are fine-tuned with and without SAM on the mixture of tasks for 20k steps. We see that SAM improves performance across *all* model sizes.

Model	Natural Q.	Web Q.	TriviaQA
Small	19.2 / 15.0	23.8 / 17.7	10.9 / 8.1
Small + SAM (0.05)	20.8 / 16.5	25.9 / 20.4	12.3 / 9.4
Base	26.3 / 21.1	31.6 / 26.0	20.6 / 16.8
Base + SAM (0.15)	27.8 / 22.6	33.6 / 27.5	23.7 / 19.5
Large	28.1 / 23.0	32.7 / 25.8	25.1 / 20.8
Large + SAM (0.15)	30.8 / 25.3	34.4 / 28.0	28.8 / 24.2
XL	33.4 / 27.3	37.1 / 30.6	35.5 / 30.2
XL + SAM (0.15)	34.2 / 28.3	39.4 / 32.3	37.6 / 32.0

Table 5: Experimental results (F1/EM) (test scores) on the (full) CBQA tasks, where the model is trained on each of the three tasks separately, rather than on a mixture. We see that SAM improves performance across *all* model sizes, as we observed when training on the mixtures. This suggests that SAM’s gains are not solely due to some ability to better leverage multi-task learning.

loss of larger improvements, set $m = 1$ and use 1/4-th (25%) of the mini-batch, or the number of available training devices (TPU cores in our case), whichever is larger, to compute SAM’s adversarial point. This is necessary because the mini-batch gradient computation is parallelized over devices and each device must receive at least one example. We’ve observed from wall-clock times that with these settings, SAM is all in all about 25% slower than standard training.

Hyper-parameters. SAM has a single hyper-parameter ρ , which is size of the step taken along the unit adversarial gradient vector. We search the range $[0.02, 0.05, 0.1, 0.15, 0.2, 0.3]$ a *single* time *only* when fine-tuning on SuperGLUE. We found that 0.05 is a reasonable choice for T5.1.1 small models, and 0.15 for the Base, Large, and XL variants, and so for all subsequent experiments except for TyDiQA, we use these choices with-

out additional tuning. For the mT5 model on TyDiQA, we found that a smaller ρ was necessary for good performance. For this, we searched the range $[0.01, 0.02, 0.05]$.

For all fine-tuning, we use the AdaFactor optimizer with learning rate 1e-3, 128 batch size, and the T5.1.1 settings. For SuperGLUE, we use 10% dropout rate, 512 input sequence length, 62 target sequence length, and fine-tune for 250k steps. For Natural Questions, Web Questions, and TriviaQA, we use 5% dropout, 38 input sequence length, 18 target sequence length, and fine-tune for 20k steps. For TyDiQA, we use the official, public mT5 checkpoints, 10% dropout, 1024 input sequence length, 512 target sequence length, and fine-tune for 20k steps. We run each experiment once, due to resource constraints, and we take the best checkpoint (stored every 1k steps for SuperGLUE and GLUE and every 200 steps for all other datasets) across

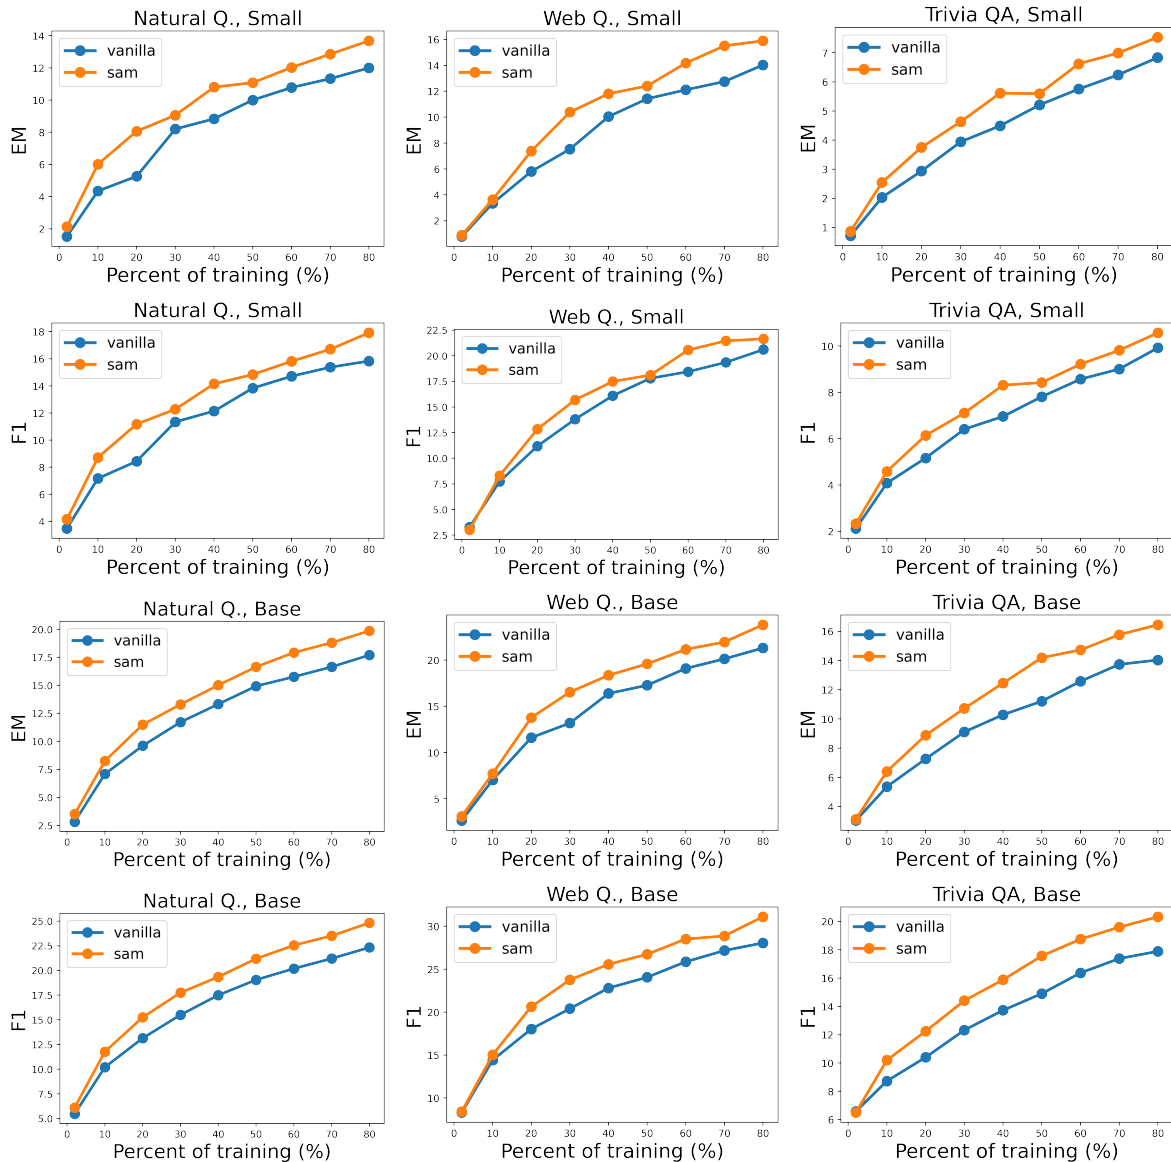


Figure 1: CBQA results at various training data sampling rates, for the Small (top half) and Base (bottom half) models. We see that SAM’s improvement is consistent across data size regimes and that the relative improvement is often largest in the ballpark of 20%.

training steps. Following standard practice, we report the best checkpoint for each task-metric pair (e.g. SuperGLUE CB F1) individually.

4.2 Full Data Results

Results for SuperGLUE and GLUE are shown in Table 1 and Table 2 respectively. We observe that SAM improves the overall scores for both benchmarks across all T5 model sizes. For Base and XL sizes on SuperGLUE, SAM brings 4.2% and 2.1% relative gains in overall score respectively, while the gain for Large on GLUE is 2.4%. As shown in Table 4, on Natural Questions, Web Questions, and Trivia QA tasks, we observe improvements

for each task, metric (F1 and EM), and model size. For Base, we see a 13.8%, 8.8%, and 13.7% gain on the exact match metric for Natural Questions, Web Questions, and Trivia QA respectively. For Large, these figures are 12.1%, 7.2%, and 15.7%. Table 3 shows the results for TyDiQA-GoldP. Here, we observe more modest improvements in the 1-2% range.

SAM improves performance on all model sizes.

In light of the conventional wisdom that “larger models generalize better,” we suspected, a priori, that SAM would be more helpful for the smaller models we consider, like Small and Base, and that we should expect substantial diminishing returns

Model	SGLue	BoolQ	CB	CoPA	MultiRC	ReCoRD	RTE	WiC	WSC
Small	50.2	60.8	37.0 / 55.4	52.0	60.1 / 11.0	33.9 / 32.5	54.5	54.4	65.4
Small + SAM (0.05)	51.9	60.5	45.6 / 66.1	53.0	61.1 / 12.5	36.7 / 34.5	52.3	55.2	66.3
Base	52.9	59.6	32.3 / 55.4	53.0	60.2 / 11.8	47.8 / 46.5	58.5	57.2	68.3
Base + SAM (0.15)	56.7	61.4	41.8 / 64.3	55.0	62.4 / 15.7	59.7 / 57.9	62.5	55.3	68.3
Large	62.8	65.3	40.1 / 62.5	62.0	71.6 / 24.0	80.4 / 78.9	69.7	57.7	69.2
Large + SAM (0.15)	64.3	77.3	47.9 / 69.6	59.0	69.0 / 20.4	81.5 / 80.0	65.0	59.6	69.2
XL	75.9	84.5	57.0 / 82.1	86.0	82.4 / 48.7	83.3 / 81.7	78.7	66.0	74.0
XL + SAM (0.15)	77.0	82.5	58.9 / 83.9	85.0	79.9 / 45.3	86.8 / 85.6	80.5	64.4	83.7

Table 6: SuperGLUE results when only 5% of the training data is available. We see again that SAM boosts performance across the board, adding a whopping 7.2% relative improvement on the Base model.

Model	Natural Q.	Web Q.	TriviaQA
Small	4.9 / 2.9	4.8 / 1.8	3.0 / 1.3
Small + SAM (0.05)	6.0 / 3.7	7.1 / 2.2	3.3 / 1.6
Base	7.9 / 4.8	13.7 / 4.6	7.6 / 4.2
Base + SAM (0.15)	8.6 / 5.6	12.2 / 5.7	7.7 / 4.4
Large	8.7 / 5.2	14.0 / 7.0	9.8 / 6.0
Large + SAM (0.15)	10.5 / 6.6	14.9 / 7.7	10.6 / 7.1
XL	13.1 / 8.0	20.6 / 11.9	19.6 / 15.3
XL + SAM (0.15)	13.4 / 8.1	22.9 / 13.6	19.1 / 14.5

Table 7: CBQA results when only 5% of the training data is available. We see SAM helps here as it did for sub-sampled SuperGLUE. For Natural Questions, SAM improves Base model performance by a relative 8.86%/16.6% (F1/EM).

as we scale up the model size. Surprisingly, we did not observe any clear pattern with regards to size: indeed, sometimes the gains on XL were larger than those on Small. Thus, we lean to recommend SAM to all practitioners regardless of the regime in model capacity they are working in.

SAM improves single-task and multi-task learning alike. Thus far, SAM has been trained on a mixture of tasks, where the influence of a particular task is proportional to the number of examples in its training split (i.e. no artificial up or down-weighting). To rule out the possibility that the gains observed are *solely* due to some ability of SAM’s to leverage multi-task learning and improve cross-task transfer, we conduct the following ablation. For each of the three CBQA tasks, we train only on a single task and report the performance on that task’s test set. Results are shown in Table 5. Indeed, we see similar gains when training and testing on each single task individually. We conclude that the mechanism driving SAM’s improvements affect single-task and multi-task learning alike.

4.3 When training data is limited

We now switch gears and evaluate whether or not SAM helps when training data is scarce. Prior work (Chen et al., 2021) showed that for vision models and tasks, SAM helps more when there is less training data to learn from. To test whether this holds for language, we do as follows: we sub-sample the training splits for both SuperGLUE and CBQA datasets at rates ranging from 2% to 80%, and observe test performance when the public checkpoint is fine-tuned with and without SAM. SuperGLUE and CBQA results at a 5% sampling rate are shown in Tables 6 and 7 respectively. In both cases we see again that SAM boosts performance across the board, adding, for example, a whopping 7.2% relative improvement on the Base model on 5% SuperGLUE and a relative 8.86%/16.6% to F1/EM on Natural Questions.

Figure 1 plots the performance on the three CBQA tasks as a function of the sampling rate. We observe consistent gains from SAM across the size of the subsampled training set, with the relative improvement appearing largest when the subsampling

rate is around 20%.

4.4 Sensitivity to hyper-parameters

Figure 2 shows the impact of SAM’s hyper-parameters ρ , the ascent micro-batch size a , and the sharpness factor m on the (full) SuperGLUE benchmark for the Base model. For ρ , we see that all tested values perform better than fine-tuning without SAM. However, 0.15 is a “sweet spot,” performing better than values below or above it. Thus, practitioners with little computational budget for hyper-parameter tuning may still see large gains by using a non-optimal ρ , while those with a generous budget should consider tuning. For the ascent micro-batch size a , we see that when the normal (descent) batch size is 128, there is improvement as a is increased to 32 but little past this point. Thus, setting a to be 1/4-th the descent batch size, as we do throughout our experiments, provides a good trade-off between performance and computational overhead. Increasing the sharpness m , where each of the m ascent micro-batches has size $32/m$, does not improve performance here. We thus recommend a default of 1, which is the setting used across our experiments. Full results are shown in the Appendix.

5 Conclusion

To the best of our knowledge, this paper is the first to demonstrate how the recently-proposed Sharpness-Aware Minimization can be applied for fine-tuning the ubiquitous text-to-text Transformer (T5) and its multilingual counterpart mT5 on language tasks of broad interest. We thereby corroborate the already-documented success the method has had in the vision domain. Furthermore, we reveal SAM’s benefits when data is limited by fine-tuning on subsamples of the original task training split. By approximating the ascent step of the algorithm via fewer samples, we show how large gains can be had across benchmarks and model sizes while adding only around 25% additional compute and wall-clock training time. Our hope is that this work will spur SAM’s adoption in the natural language processing community the way it is starting to in the vision one.

References

Armen Aghajanyan, Akshat Shrivastava, Ancht Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta.




Figure 2: Impact of SAM’s hyper-parameters ρ , the ascent micro-batch size a , and the sharpness m on the SuperGLUE benchmark for the Base model. Except for the ablated hyper-parameter, we use $\rho = 0.15$, $a = 32$, $m = 1$. For ρ , we see that all tested values perform better than fine-tuning without SAM. However, 0.15 is a “sweet spot,” performing better than values below or above it. For the ascent micro-batch size a , we see that when the descent batch size is 128, there is improvement as a is increased to 32 but little past this point. Thus, setting a to be 1/4-th the descent batch size provides a good trade-off between performance and computational overhead. Increasing the sharpness m does not help (in fact it hurts) here. Full results are shown in the Appendix.

2020. Better fine-tuning by reducing representational collapse. *arXiv preprint arXiv:2008.03156*.
- Rohan Anil, Vineet Gupta, Tomer Koren, Kevin Regan, and Yoram Singer. 2020. Scalable second order optimization for deep learning. *arXiv preprint arXiv:2002.09018*.
- David GT Barrett and Benoit Dherin. 2020. Implicit gradient regularization. *arXiv preprint arXiv:2009.11162*.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1533–1544.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. **JAX: composable transformations of Python+NumPy programs**.
- Andrew Brock, Soham De, Samuel L Smith, and Karen Simonyan. 2021. High-performance large-scale image recognition without normalization. *arXiv preprint arXiv:2102.06171*.
- Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. 2019. Entropy-sgd: Biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124018.
- Xiangning Chen, Cho-Jui Hsieh, and Boqing Gong. 2021. When vision transformers outperform resnets without pretraining or strong data augmentations. *arXiv preprint arXiv:2106.01548*.
- Jonathan H Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. Tydi qa: A benchmark for information-seeking question answering in typologically diverse languages. *Transactions of the Association for Computational Linguistics*, 8:454–470.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. 2020. Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*.
- Vineet Gupta, Tomer Koren, and Yoram Singer. 2018. Shampoo: Preconditioned stochastic tensor optimization. In *International Conference on Machine Learning*, pages 1842–1850. PMLR.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. 2021. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349.
- Dan Hendrycks and Thomas Dietterich. 2019. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.
- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2019. Smart: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization. *arXiv preprint arXiv:1911.03437*.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*.
- Bobby Kleinberg, Yuanzhi Li, and Yang Yuan. 2018. An alternative view: When does sgd escape local minima? In *International Conference on Machine Learning*, pages 2698–2707. PMLR.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. 2021. Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. *arXiv preprint arXiv:2102.11600*.
- David JC MacKay. 1995. Probable networks and plausible predictions—a review of practical bayesian methods for supervised neural networks. *Network: computation in neural systems*, 6(3):469.

- James Martens and Roger Grosse. 2015. Optimizing neural networks with kronecker-factored approximate curvature. In *International conference on machine learning*, pages 2408–2417. PMLR.
- Hossein Mobahi, Mehrdad Farajtabar, and Peter L Bartlett. 2020. Self-distillation amplifies regularization in hilbert space. *arXiv preprint arXiv:2002.05715*.
- Rafael Müller, Simon Kornblith, and Geoffrey Hinton. 2019. When does label smoothing help? *arXiv preprint arXiv:1906.02629*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam 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 arXiv:1910.10683*.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? *arXiv preprint arXiv:2002.08910*.
- Noam Shazeer and Mitchell Stern. 2018. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR.
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. 2016. [On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima](#). *arXiv e-prints*, page arXiv:1609.04836.
- Samuel L Smith, Benoit Dherin, David GT Barrett, and Soham De. 2021. On the origin of implicit regularization in stochastic gradient descent. *arXiv preprint arXiv:2101.12176*.
- Samuel L Smith and Quoc V Le. 2017. A bayesian perspective on generalization and stochastic gradient descent. *arXiv preprint arXiv:1710.06451*.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, et al. 2021. Mlp-mixer: An all-mlp architecture for vision. *arXiv preprint arXiv:2105.01601*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019. Super-glue: A stickier benchmark for general-purpose language understanding systems. *arXiv preprint arXiv:1905.00537*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Andrew Gordon Wilson and Pavel Izmailov. 2020. Bayesian deep learning and a probabilistic perspective of generalization. *arXiv preprint arXiv:2002.08791*.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2020. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*.
- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. 2017. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. 2019. FreeLB: Enhanced adversarial training for natural language understanding. *arXiv preprint arXiv:1909.11764*.

6 Appendix

6.1 TyDiQA-GoldP Results

Table 8 shows the per-language TyDiQA-GoldP scores. We found that the multilingual mT5 model benefited from a smaller ρ than the vanilla T5 model.

6.2 Impact of hyper-parameters on SuperGLUE

Table 9 shows the full (no subsampling) SuperGLUE results for the Base model for different hyper-parameter choices.

Model	avg (F1/EM)	en	ar	bn	fi	id	ko	ru	sw	te
Small	73.4 / 62.1	66.3 / 54.5	78.0 / 62.8	69.4 / 60.2	73.7 / 60.9	77.8 / 65.0	64.1 / 55.8	69.5 / 56.7	78.0 / 68.9	84.0 / 74.4
Small + SAM (0.02)	74.3 / 63.0	67.1 / 55.7	79.2 / 64.9	68.4 / 57.5	75.0 / 61.8	78.8 / 67.6	64.0 / 54.3	72.0 / 58.0	79.9 / 72.1	84.2 / 75.0
Base	81.6 / 71.0	76.7 / 65.7	84.2 / 70.4	81.7 / 71.7	81.6 / 69.3	85.1 / 74.2	73.6 / 66.3	78.8 / 64.7	84.4 / 77.2	87.9 / 79.4
Base + SAM (0.02)	82.0 / 71.3	76.1 / 65.0	83.9 / 70.2	84.4 / 75.2	81.1 / 69.6	85.4 / 74.0	74.5 / 66.7	79.5 / 65.3	84.8 / 76.4	88.8 / 79.1
Large	85.6 / 75.3	81.2 / 70.0	86.6 / 73.1	86.3 / 77.9	84.6 / 71.6	87.7 / 77.9	81.1 / 72.1	84.0 / 72.8	88.7 / 81.2	90.3 / 81.5
Large + SAM (0.02)	85.9 / 76.1	82.3 / 71.6	87.3 / 74.4	86.6 / 79.6	84.5 / 72.5	88.1 / 80.0	81.3 / 73.2	84.2 / 71.9	88.3 / 79.8	90.9 / 82.1
XL	87.0 / 77.4	82.1 / 72.0	87.3 / 74.4	88.9 / 82.3	85.8 / 74.6	89.7 / 80.4	81.7 / 73.6	85.1 / 73.5	90.7 / 83.2	91.5 / 82.4
XL + SAM (0.05)	87.3 / 77.7	82.8 / 73.4	87.6 / 75.0	89.2 / 81.4	86.4 / 74.2	89.6 / 80.2	82.6 / 75.0	85.5 / 74.8	90.7 / 83.0	91.4 / 82.5

Table 8: Full Results for TyDiQA-GoldP.

Model	SGlue	BoolQ	CB	CoPA	MultiRC	ReCoRD	RTE	WiC	WSC
Base + SAM (0.02)	76.6	80.5	92.4 / 92.9	73.0	76.2 / 36.8	77.3 / 76.4	81.9	70.8	80.8
Base + SAM (0.05)	77.2	80.3	97.4 / 96.4	73.0	76.4 / 37.9	77.8 / 76.9	83.8	71.9	76.9
Base + SAM (0.1)	77.4	81.7	94.8 / 94.6	72.0	76.8 / 38.3	79.3 / 78.3	83.8	72.7	77.9
Base + SAM (0.15)	78.5	82.2	93.7 / 94.6	78.0	77.5 / 39.1	78.2 / 77.2	85.9	70.4	81.7
Base + SAM (0.2)	77.7	82.3	95.0 / 96.4	75.0	77.2 / 39.5	77.9 / 76.8	84.1	71.8	76.9
Base + SAM (0.3)	77.8	81.1	93.6 / 94.6	74.0	77.4 / 40.0	79.8 / 78.7	85.9	70.8	78.8
Base + SAM (0.4)	77.6	80.6	94.3 / 96.4	79.0	76.3 / 37.5	77.9 / 76.9	81.9	71.0	78.8
Base + SAM (8)	77.8	81.9	97.4 / 96.4	72.0	77.5 / 40.8	78.6 / 77.5	84.5	71.0	78.8
Base + SAM (24)	78.1	82.5	97.4 / 96.4	74.0	77.2 / 40.2	79.2 / 78.4	83.8	69.4	80.8
Base + SAM (32)	78.5	81.9	96.1 / 94.6	77.0	77.6 / 40.3	78.9 / 77.8	85.6	71.9	78.8
Base + SAM (64)	78.5	82.3	96.1 / 94.6	79.0	78.4 / 40.8	79.3 / 78.2	84.5	71.3	76.9
Base + SAM (128)	78.8	82.4	97.4 / 96.4	74.0	77.6 / 40.6	79.0 / 78.0	87.0	72.9	79.8
Base + SAM (1)	78.6	81.5	96.1 / 96.4	73.0	77.2 / 40.0	79.0 / 77.9	87.0	72.6	81.7
Base + SAM (2)	78.1	81.8	96.1 / 94.6	77.0	76.5 / 37.6	78.5 / 77.6	84.8	71.8	78.8
Base + SAM (4)	78.3	82.3	97.4 / 96.4	75.0	77.3 / 38.5	78.2 / 77.1	85.2	72.1	79.8

Table 9: SuperGLUE results for Base for different values of ρ (**top**), ascent micro-batch size a (**middle**), and sharpness m (**bottom**). Except for the ablated hyper-parameter, we use $\rho = 0.15$, $a = 32$, $m = 1$.