Memory-Efficient Adaptive Optimization

Rohan Anil Vineet Gupta Google Brain {rohananil,vineet}@google.com Tomer Koren Google Brain and Tel Aviv Univ. tkoren@google.com Yoram Singer Princeton Univ. y.s@cs.princeton.edu

Abstract

Adaptive gradient-based optimizers such as Adagrad and Adam are crucial for achieving state-of-the-art performance in machine translation and language modeling. However, these methods maintain second-order statistics for each parameter, thus introducing significant memory overheads that restrict the size of the model being used as well as the number of examples in a mini-batch. We describe an effective and flexible adaptive optimization method with greatly reduced memory overhead. Our method retains the benefits of per-parameter adaptivity while allowing significantly larger models and batch sizes. We give convergence guarantees for our method, and demonstrate its effectiveness in training very large translation and language models with up to 2-fold speedups compared to the state-of-the-art.

1 Introduction

Adaptive gradient-based optimizers such as Adagrad [11] and Adam [15] are among the de facto methods of choice in modern machine learning. These methods adaptively tune the learning rate for each parameter during the optimization process using cumulative second-order statistics. Often offering superior convergence properties, these methods are very attractive in large scale applications due to their moderate time and space requirements, which are linear in the number of parameters. However, when training extremely large models even the modest memory overhead imposes grave limitations on the quality of the trained model. For example, recent advances in natural language processing [26, 17] show that models with hundreds of millions to billions of parameters, trained with adaptive optimization methods, achieve state-of-the-art results. In such cases, the memory overhead of the optimizer severely restricts the size of the model that can be used as well as the number of examples in each mini-batch, both of which have a dramatic effect on the accuracy of the model.

Motivated by these challenges, we describe an adaptive optimization method that retains the benefits of standard per-parameter adaptivity while significantly reducing memory overhead. Our construction is general and flexible, and very simple to implement. We give convergence guarantees for our method in the convex (online or stochastic) optimization setting, and demonstrate experimentally that it is particularly effective when the gradients exhibit natural *activation patterns*; namely, when the parameters can be subdivided into (not necessarily disjoint) sets where gradient entries within sets are correlated and of a similar order of magnitude. For example, we often observe in deep networks that the incoming (outgoing) edges into (from) a neuron are jointly activated and, loosely speaking, their associated gradients exhibit similar statistical characteristics. That said, our analysis of the algorithm makes no statistical assumptions on the gradients and is applicable for general stochastic convex optimization. Further, we do not assume that the activation pattern is fully prescribed a-priori.

Large scale experiments show that our algorithm achieves comparable, and at times superior, rates of convergence compared to standard linear-space adaptive methods. Focusing primarily on language modeling tasks where state-of-the-art models are extremely large, we further demonstrate that the reduction in memory footprint can be utilized for a substantial increase in the batch size, which greatly speeds up convergence in a distributed environment. For a fixed budget of computational resource our method is able to shorten the end-to-end walltime for convergence by up to 50%. Our

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

method exhibits slightly improved per-step time. The latter could be attributed to reduction in the frequency of memory accesses.

1.1 Related work

Adaptive learning rates in online and stochastic optimization date back at least to [5] and were popularized in [11], [16], the former of which introduced the well-known Adagrad algorithm. Several variants of Adagrad have now been proposed in the optimization and machine learning literature (see [19] and the references therein), the most notable of which is Adam [15]. All of these methods require (at least) linear space for maintaining various per-parameter statistics during their execution. One notable exception, which is directly related to our work, is the Adafactor algorithm [23] that was proposed as a way to reduce the memory costs of Adam, primarily for training large language models. While the memory requirements of our construction are similar to Adafactor's, the application scope and the convergence properties of the two algorithms are quite different. We discuss the relationship in more detail in Section [4] and give an empirical comparison between the algorithms in Section [5].

Spring et al. [25] provide an alternative way to reduce memory costs, making use of the Count-Sketch data structure [7] to maintain a compressed approximation to the auxiliary variables. One key difference between SM3 and Count-Sketch is that SM3 uses specific hash functions instead of random hash functions. Our hash functions are compatible with slices of parameter tensors and are geared towards exploiting empirically observed correlations between the auxiliary parameters, as we discuss below (see Section 4). As a result, our sketches can be 100x-1000x smaller than the original tensors—compared to the 5x reduction reported in [25]—while showing significantly smaller approximation error (we provide details in the full version of the paper [3]). In addition, randomized sketching is extremely inefficient to implement on GPUs and TPUs, since it involves sparse look-ups and is not cache-efficient. These differences allow us to show significant improvements for a large variety of tasks and models, as compared to the results in [25].

Also related to our work is the Shampoo algorithm for optimization over tensor structures [12]. The goal of Shampoo is very different from ours: going beyond entry-wise learning rates and employing *full-matrix* regularization in a computationally efficient way. Nonetheless, Shampoo can also be seen as a method to substantially reduce the memory footprint of full-matrix preconditioned algorithms (specifically, full-matrix Adagrad). In a sense, our algorithms are analogous to a diagonalized version of the Shampoo algorithm. Yet another recent adaptive optimization method is the GGT algorithm [2]. Similarly to Shampoo, the goal of the latter is to reduce the computation cost of full-matrix preconditioning in order to make it practical in large scale settings. However, GGT stores multiple copies of the gradient over the course of its execution, and as a result, its space requirements restricts it from being applied at large scale.

2 Preliminaries

2.1 Online optimization

We henceforth assume the general *online optimization* setting (see [22, [3]). Online optimization consists of rounds t = 1, ..., T, where in each round the algorithm chooses a parameter vector $w_t \in \mathbb{R}^d$. After making a choice on round t, the algorithm receives a loss function $\ell_t : \mathbb{R}^d \to \mathbb{R}$ which is used to form an update of the parameters. In our analysis, we focus on online *convex* optimization in which $\ell_1, ..., \ell_T$ are convex. Often, as is the case in this paper, the update is determined by the gradient $g_t = \nabla \ell_t(w_t)$ of the instantaneous loss ℓ_t at the current iterate w_t . The algorithm is measured by its *T*-round regret with respect to a given comparator $w^* \in \mathbb{R}^d$, defined as the quantity $\sum_{t=1}^{T} \ell_t(w_t) - \sum_{t=1}^{T} \ell_t(w^*)$. An online optimization algorithm is convergent if its regret is o(T), i.e., its average regret approaches zero as *T* grows.

The above setting includes stochastic (possibly mini-batched) optimization as a special case. In stochastic optimization the underlying goal is to minimize a population loss $L(w) = \mathbb{E}_{z \sim D}[\ell(w, z)]$ based on samples of *z*. Here $\ell(w, z)$ defines the loss of parameters *w* w.r.t a batch *z*. The online loss function $\ell_t(w) = \ell(w, z_t)$ is the average loss over a mini-batch z_t received on iteration *t*. The stochastic gradient g_t is a conditionally unbiased estimate of the gradient of *L* at the current parameter vector w_t . Under convexity assumptions, an online algorithm with vanishing average regret can be converted to a stochastic optimization algorithm for minimizing the population loss L [6].

2.2 Adaptive methods

For the sake of self-containment, we give a brief description of adaptive gradient methods, focusing on Adagrad [11]. Adagrad maintains at each step *t* parameter-wise accumulated statistics which are computed from the previously obtained gradients g_1, \ldots, g_t :

$$\gamma_t(i) = \sum_{s=1}^t g_s^2(i) , \qquad \forall i \in [d] .$$
⁽¹⁾

Based on these statistics, the update rule of the algorithm on step t takes the form:

$$w_{t+1}(i) = w_t(i) - \eta \frac{g_t(i)}{\sqrt{\gamma_t(i)}}, \qquad \forall i \in [d],$$

where $\eta > 0$ is an external learning rate parameter. Duchi et al. [11] proved the following regret bound for Adagrad with respect to a given w^* (with properly tuned η):

$$\sum_{t=1}^{T} \ell_t(w_t) - \sum_{t=1}^{T} \ell_t(w^*) = O\left(D\sum_{i=1}^{d} \sqrt{\sum_{t=1}^{T} g_t^2(j)}\right),$$
(2)

where $D \ge \max_t ||w_t - w^*||_{\infty}$. Adagrad has proved to be particularly useful in training sparse models, where the effective learning rates $\eta/\sqrt{\gamma_t(i)}$ decay in a moderate way for rare, yet potentially informative, features. In these settings, Adagrad can potentially lead to substantial improvements in convergence time; see for instance the discussion in [11]. Crucially, however, Adagrad must maintain auxiliary sequence of accumulators γ_t and thus needs $\Omega(d)$ additional space. The goal of this paper is to provide memory-efficient methods with comparable convergence characteristics that refrain from maintaining the full vectors γ_t .

3 The SM3 Algorithm

We now present our memory-efficient adaptive optimization algorithm. As an abstraction, the algorithm employs a *cover* of the parameters: a collection of k nonempty sets $\{S_r\}_{r=1}^k$, such that $S_r \subseteq [d]$ and $\bigcup_r S_r = [d]$. In particular, each index $i \in [d]$ may be contained in multiple sets S_r . The algorithm maintains a single variable for each set S_r in the cover. Thus, the additional space it requires is O(k) rather than the O(d) required by standard adaptive methods. In large scale applications, k will be chosen to be negligible in comparison to d, which would translates to substantial savings in memory; see Section 4 for a discussion on the covers used in practice.

Concretely, for each set S_r in the cover, the algorithm maintains a running sum, $\mu_t(r)$, of the *maximal* variance over all gradient entries $j \in S_r$. Next, for each parameter *i*, we take the *minimum* over all variables $\mu_t(r)$ associated with sets which cover *i*, denoted $S_r \ni i$. Thereafter, the learning rate corresponding to the *i*'th gradient entry is determined by taking the square-root of this minimum, denoted by $\nu_t(i)$. Accordingly, we name our algorithm the Square-root of Minima of Sums of Maxima of Squared-gradients Method, or in short, *SM3*. See Algorithm SM3-II for its pseudocode.

As noted above, SM3-**I** requires only O(k) space in addition to the space required for storing the parameters w_t themselves. The time per iteration of SM3-**I** is $O(\sum_{r=1}^{k} |S_r|)$. To see this, consider a bipartite graph defined over d+k vertices. Nodes on one side of the graph correspond to indices $i \in [d]$, while nodes on the other side correspond to indices $j \in [k]$. The edges of the graphs are all pairs (i, j) such that $i \in S_j$. The complexity of each inner for-loop of the algorithm scales with the number of edges in this graph, which is equal to $O(\sum_{r=1}^{k} |S_r|)$. Note that updating the weights w_t takes O(d) time, which is always dominated by the former quantity.

SM3-I

1: **parameters:** learning rate η 2: initialize $w_1 = 0$; $\forall r \in [k]$: $\mu_0(r) = 0$ 3: **for** t = 1, ..., T **do** 4: receive gradient $g_t = \nabla \ell_t(w_t)$ 5: **for** r = 1, ..., k **do** 6: set $\mu_t(r) \leftarrow \mu_{t-1}(r) + \max_{j \in S_r} g_t^2(j)$ 7: **for** i = 1, ..., d **do** 8: set $\nu_t(i) \leftarrow \min_{r:S_r \ni i} \mu_t(r)$ 9: update $w_{t+1}(i) \leftarrow w_t(i) - \eta g_t(i) / \sqrt{\nu_t(i)}$

 \triangleright with the convention that 0/0 = 0

The following provides convergence guarantees for SM3-I.

Proposition 1. Assume that the loss functions ℓ_1, ℓ_2, \ldots are convex, and let w_1, w_2, \ldots be the iterates generated by SM3 [I]. Then, for any $w^* \in \mathbb{R}^d$,

$$\sum_{t=1}^{T} \left(\ell_t(w_t) - \ell_t(w^*) \right) \le 2D \sum_{i=1}^{d} \sqrt{\min_{r:S_r \ni i} \sum_{t=1}^{T} \max_{j \in S_r} g_t^2(j)} ,$$

where $\max_t ||w_t - w^*||_{\infty} \leq D$ and choosing $\eta = D$.

For stochastic optimization, i.e., when the functions ℓ_t correspond to i.i.d. samples with $\mathbb{E}[\ell_t(w)] = L(w)$, the above bound translates via standard arguments to a $O(1/\sqrt{T})$ -type convergence guarantee for the average iterate $\overline{w}_T = \frac{1}{T} \sum_{t=1}^T w_t$ of the form

$$\mathbb{E}[L(\overline{w}_T)] - L(w^*) = O\left(\frac{1}{T}\sum_{i=1}^d \mathbb{E}_{\sqrt{\sum_{r:S_r \ni i}^T \max_{r:S_r \ni i} g_t^2(j)}}\right).$$

Note that adding more sets S_r to the cover used by SM3 always improves its convergence bound, but results in a worse space complexity and a higher runtime per step. When k = d and $S_i = \{i\}$ for all $i \in [d]$, SM3-[] reduces to the Adagrad algorithm, and the regret bound in Proposition [] then precisely recovers the bound attained by Adagrad (recall Eq. (2)). In general, the right-hand side of Proposition [] is never smaller than Adagrad's regret bound, as expected from a space-restricted scheme (this is a consequence of Claim 2 below). Nevertheless, the two bounds can be of similar order of magnitude in practical scenarios; see Section 4 below for a detailed discussion.

We now give a proof of Proposition []. First, we state two elementary properties of the step sizes the algorithm computes. For a proof, see the full version of the paper [3].

Claim 2. For any *i*, the sequence $\nu_1(i), \nu_2(i), \ldots$ is monotonically increasing, and $\nu_t(i) \ge \sum_{s=1}^t g_s^2(i)$.

Proof of Proposition [] Let us first assume that $g_1(i) > 0$ for all *i*, so that $\nu_t(i) > 0$ for all *i* and $t \ge 1$ due to Claim [2]. We start by observing that SM3-[] performs Online Mirror Descent updates, where the step on round *t* uses the positive definite diagonal matrix $H_t = \text{diag}(\nu_t^{1/2})$ for regularization. Then, employing a standard regret bound for the Online Mirror Descent algorithm with time-dependent regularization (see for instance []]]. Proposition 3]), the regret of the algorithm is bounded by

$$\frac{1}{2\eta} \sum_{t=1}^{T} \left(\|w_t - w^{\star}\|_{H_t}^2 - \|w_{t+1} - w^{\star}\|_{H_t}^2 \right) + \frac{\eta}{2} \sum_{t=1}^{T} \left(\|g_t\|_{H_t}^* \right)^2.$$

Here, $||x||_H = \sqrt{x^T H x}$ and $||\cdot||^*$ is the corresponding dual norm, $||x||_H^* = \sqrt{x^T H^{-1} x}$. Henceforth, for notational convenience we set $\nu_0 = 0$. Simplifying the first sum above using the fact that H_t are diagonal matrices, we have

$$\sum_{t=1}^{T} \left(\|w_t - w^{\star}\|_{H_t}^2 - \|w_{t+1} - w^{\star}\|_{H_t}^2 \right) \le \sum_{t=1}^{T} \left(\nu_t^{1/2} - \nu_{t-1}^{1/2} \right) \cdot \left(w_t - w^{\star} \right)^2$$
$$\le \sum_{t=1}^{T} \left(\nu_t^{1/2} - \nu_{t-1}^{1/2} \right) \cdot \left(\|w_t - w^{\star}\|_{\infty}^2 \mathbf{1}_d \right)$$
$$\le D^2 \left(\nu_T^{1/2} \cdot \mathbf{1}_d \right) = D^2 \operatorname{Tr}(H_T) .$$

Now, let $\gamma_t(i) = \sum_{s=1}^t g_s^2(i)$ and consider the positive definite diagonal matrix $G_t = \text{diag}(\gamma_t^{1/2})$. From [12] Lemma 2] with $\Phi(G) = \text{Tr}(G)$, we have

$$\sum_{t=1}^{T} \left(\|g_t\|_{G_t}^* \right)^2 \le \sum_{t=1}^{T} \left(\|g_t\|_{G_T}^* \right)^2 + \operatorname{Tr}(G_T) = \gamma_T^{-1/2} \cdot \gamma_T + \operatorname{Tr}(G_T) = 2 \operatorname{Tr}(G_T).$$

¹Here we implicitly assume that the iterates of the algorithm remain bounded and D is a constant. This can be enforced by projecting the iterates to a bounded set of choice; we avoid introducing projections explicitly as they are rarely used in practice.

Also, from Claim 2 we know that for all $t, H_t \succeq G_t$, thus

$$\sum_{t=1}^{T} \left(\|g_t\|_{H_t}^* \right)^2 \leq \sum_{t=1}^{T} \left(\|g_t\|_{G_t}^* \right)^2 \leq 2 \operatorname{Tr}(G_T) \leq 2 \operatorname{Tr}(H_T) \,.$$

In summary, we have established that

$$\sum_{t=1}^{I} \ell_t(w_t) - \ell_t(w^*) \le \left(\frac{D^2}{2\eta} + \eta\right) \operatorname{Tr}(H_T) \,.$$

Plugging in $\eta = D$ and the expression for the diagonal elements of H_T , we obtain the claim.

For the degenerate case where the matrices H_t may not be strictly positive definite, a careful yet technical inspection of the proof above reveals that our arguments apply to this case as well by replacing inverses with pseudo-inverses. The rest of the proof remains intact as the algorithm does not update parameter *i* on step *t* if the corresponding diagonal entry in H_t is zero.

3.1 SM3-II

We now discuss a slightly more efficient variant of SM3, which we describe in SM3-III It is similar to SM3-III and improves on the latter in the following sense.

Proposition 3. For any $i \in [d]$, the sequence $\nu'_1(i), \ldots, \nu'_T(i)$ is monotonically increasing. Further, fixing a sequence of gradients g_1, \ldots, g_T , we have for all t, i that $\sum_{s=1}^t g_s^2(i) \le \nu'_t(i) \le \nu_t(i)$, where $\nu_1(i), \ldots, \nu_T(i)$ is the sequence SM3-1 emits upon receiving the gradients g_1, \ldots, g_T .

(See the full version of the paper [3] for a proof.) In other words, SM3-[1] provides a tighter upper bound on the cumulative gradient squares than SM3-[]. Consequently, we can show, along similar lines to the proof of Proposition [], a slightly better bound for SM3-[1] that scales with the quantity $\sum_{i=1}^{d} \sqrt{\nu'_{i}(i)}$, which is always smaller than the one appearing in the bound of SM3-[].

4 Discussion

SM3-II 1: **parameters:** learning rate η 2: initialize $w_1 = 0$; $\forall r \in [k]$: $\mu'_0(r) = 0$ 3: for t = 1, ..., T do 4: receive gradient $g_t = \nabla \ell_t(w_t)$ 5: initialize $\mu'_t(r) = 0$ for all $r \in [k]$ 6: for i = 1, ..., d do $\nu'_{t}(i) \leftarrow \min_{r:S_{r} \ni i} \mu'_{t-1}(r) + g_{t}^{2}(i)$ $w_{t+1}(i) \leftarrow w_{t}(i) - \eta g_{t}(i) / \sqrt{\nu'_{t}(i)}$ 7: 8: \triangleright with the convention that 0/0 = 09: for all $r : S_r \ni i$ do $\mu'_t(r) \leftarrow \max\{\mu'_t(r), \nu'_t(i)\}$ 10:

Thus far, we gave an analysis of SM3 in a worstcase (convex) setting without placing any further

assumptions on the statistical characteristics of the underlying stochastic gradients. Further, we did not attempt to relate the cover used by SM3 to properties of the underlying stochastic optimization problem. It should not come as a surprise that in this general setting, the convergence of SM3 might be much worse, at least in theory, than its linear-memory counterpart Adagrad.

Activation patterns. Often in our experiments, we observe common statistical attributes that could be exploited by SM3. Specifically, we see that certain entries of the stochastic gradients have (on average) similar values, and exhibit what we refer to as an *activation pattern*. For example, in gradients of embedding layers of deep networks, an entire row (or column) is either zero or non-zero. Similarly, in intermediate layers we often observe that gradients associated with the same unit are of similar order of magnitude. In these cases, a similar phenomenon is observed in the second-order statistics maintained by adaptive methods. In Figure 1 we visualize this phenomenon for different layers of a Transformer network. In the full version of the paper 3 we give additional illustrations of similar phenomena in convolutional layers of image classification models.

Choice of covers. The intuitive notion of an activation pattern motivates a natural and generic choice for the cover used by SM3 in practice. For the parameters of deep networks, that are organized as a collection of tensors, we form a cover consisting of slices of co-dimension 1 for each tensor. Thus, for an $m \times n$ parameter matrix, the cover consists of rows and columns of the matrix. The memory requirements therefore drop from $\Theta(mn)$ to merely $\Theta(m + n)$. For a parameter tensor of dimension $n_1 \times \cdots \times n_p$, the reduction in memory consumption is even more pronounced, dropping



Figure 1: Visualization of Adagrad's statistics (cf. Eq. (1)) for different weight matrices in Transformer-Big model trained with Adagrad on WMT'14 en \rightarrow fr (color intensities are in log scale).

from $\Theta(\prod_{i=1}^{p} n_i)$ to $\Theta(\sum_{i=1}^{p} n_i)$. This virtually eliminates the memory overhead associated with maintaining the adaptive learning rates.

We argue, though only informally, that when choice of cover used by SM3 is compatible with the observed activation patterns, we expect the convergence of SM3 to be significantly better, and closely match Adagrad. Quantitatively, if each parameter $i \in [d]$ is covered by a set S_r such that $g_s(j) \approx g_s(i)$ for all $j \in S_r$, then $\max_{j \in S_r} g_s^2(j) \approx g_s^2(i)$, and thus $\min_{r:S_r \ni i} \sum_s \max_{j \in S_r} g_s^2(j) \approx \sum_s g_s^2(i)$. Thus, the bounds in Proposition 1 and Eq. (2) are of similar order of magnitude. In other words, in such scenarios we inherit the convergence properties of Adagrad while using a negligible amount of memory. We remark that the activation pattern need not be fully specified in advance; in particular, SM3 is robust to whether a certain parameter is "row tied" or "column tied", as long as both rows and columns are included in the cover.

Comparison with Adafactor. Adafactor [23] is a very effective method for space-efficient adaptive optimization. SM3 and Adafactor differ in a number of important ways. First, Adafactor is only defined for matrix-shaped parameters while SM3 applies to tensors of arbitrary dimensions, and even more generally, to any predefined cover of the parameters. Second, Adafactor is in essence a fixed learning-rate algorithm, being a memory-constrained variation of Adam, and often requires a manually devised learning-rate schedule to ensure convergence. In contrast, SM3 adapts its learning rates in an adaptive, data-driven manner similar to Adagrad. Finally, SM3 comes with rigorous convergence guarantees in stochastic convex optimization settings.

5 Experiments

We demonstrate the practical efficacy of SM3 on several machine learning tasks using published state-of-the-art architectures. We focus on three domains: machine translation, language modeling, and image classification. We implemented SM3 as an optimizer in TensorFlow [1]; source code is publicly available at [4]. Our implementation follows the pseudocode of SM3-[1] as it performed slightly yet consistently better than SM3-[1] in our experiments (as predicted by our bounds). We use covers induced by rows and columns of matrices, and more generally, by slices of higher-order tensors (e.g., in convolutional layers represented by 4-dimensional tensors), as described in Section [4]. In addition to being compatible with the natural activation patterns, these covers facilitates efficient tensor operations available on GPUs and TPUs for computing *max* and *min* over the sets. In all experiments, we used the Cloud TPU-v2 device [14] where each core has 8GiB of memory. For more details on all of our experiments, including the precise hyperparameters used in each of them, refer to the full version of the paper [3].

5.1 Machine translation

We experimented with machine translation tasks on two standard datasets from WMT'14: English to French (en \rightarrow fr) with 36.3M sentence pairs, and English to German (en \rightarrow de) with 4.5M sentence pairs. We used the state-of-the-art Transformer architecture Vaswani et al. [26]. The basic version of this

model has 93.3M parameters and consumes 0.36GiB memory. The larger variant (Transformer-Big) has 375.4M parameters (1.432GiB) and consists of 6 layers for its encoder and decoder, where each layer has 1024 model dimensions, 8192 hidden dimensions, and 16 attention heads.

Here we report our results on the larger Transformer-Big, and defer results on the basic Transformer to the full version of the paper [3]. We trained Transformer-Big on the en \rightarrow fr dataset with batches of size 384, and compared SM3 with several standard optimizers in each of the tasks. In all cases, we used momentum (including for Adagrad) and extensively tuned all hyperparameters. We also ran SGD with momentum (with various exponential decay schedules), but it performed poorly and hence it is omitted from the figures. The results are provided in Figure 2 and Table 1, and demonstrate that SM3 performed substantially better and provided a large improvement in BLEU score compared to Adam and Adafactor. In addition, the small memory requirements of SM3 and Adafactor allowed us to *double the number of examples in a batch* to a total of 768, with minimal additional computation resources. In this setting, we found that SM3 outperformed Adafactor in terms of the number of steps as well as the wall-time to convergence by roughly a factor of 2. We further observed that SM3 approximated the 2nd-order statistics tightly. For more details, see the full version of the paper [3].

Both models were trained on a 4×4 Cloud TPU-v2 using the Lingvo [24] sequence modeling framework, with 32K word-pieces [21] for each language pair. BLEU scores were computed on the Newstest 2014 for evaluation, on tokenized, true-case outputs, and without manual post-processing of the text, similar to [28]. Our BLEU scores are not directly comparable to those of [26]. We instead followed the experimental protocol described in a later work [8].



Figure 2: Test log-perplexity of a Transformer-Big model on WMT'14 en \rightarrow fr, when training with batch sizes of 384 (left) and 768 (right). For batch size of 768, Adam and Adagrad were infeasible as they exceeded the available memory.

Optimizer	BATCH SIZE PER CORE (TOTAL)	Memory Usage per core	BLEU
Adam Adagrad Adafactor SM3	12 (384) 12 (384) 12 (384) 12 (384) 12 (384)	6.88 GiB 6.85 GiB 5.43 GiB 5.36 GiB	$\begin{array}{c} 38.96 \pm 0.002 \\ 39.90 \pm 0.003 \\ 37.89 \pm 0.002 \\ 39.81 \pm 0.002 \end{array}$
Adafactor SM3	24 (768) 24 (768)	7.04 GiB 7.02 GiB	$\begin{array}{c} 39.65 \pm 0.002 \\ \textbf{40.50} \pm \textbf{0.001} \end{array}$

Table 1: BLEU scores and memory usage for various batch sizes on the WMT'14 en \rightarrow fr dataset.

5.2 Language modeling

Next, we considered a language modeling task on the concatenation of Wikipedia and BooksCorpus [29], with 2.5B and 800M words respectively. We used the recent Bidrectional Encoder Representation (BERT) architecture of Devlin et al. [10], focusing on its larger variant, coined BERT-Large. BERT-Large is a large bidirectional transformer model containing 24 transformer blocks with 1024 hidden dimensions and 16 self attention heads. It has 340M parameters (1.297 GiB), and is set up to jointly optimize two objectives: (a) masked language model (Masked-LM) loss where the task is to

Optimizer	BATCH SIZE PER CORE (TOTAL)	Memory Usage per core
Adam SM3	8 (1024) 8 (1024)	6.15 GiB 4.90 GiB
SM3	16 (2048)	6.02 GiB

Table 2: Training memory consumption at different batch sizes for BERT-Large on 8x8 TPUs.

predict masked tokens based on surrounding context, and (b) next sentence prediction (NSP) loss where the task is to predict whether two given sentences are consecutive in the text.



Figure 3: Masked LM test accuracy (left), and number of steps to get 70% test accuracy as a function of the batch size (right), of the BERT-Large language model trained on Wikipedia+BooksCorpus. SM3 with batch size 2048 uses about the same amount of memory as Adam/Adagrad with batch size 1024, and scales linearly up to a batch size of 2¹⁶, at which point we hit the hardware memory limits.

As before, we compared SM3 with Adagrad, Adam and Adafactor. Our results are presented in Figure 3. We see that SM3 worked as well as Adam and Adagrad for a fixed batch size. However, the savings in memory allowed us to train SM3 with double the batch size, resulting in a substantial increase in accuracy. The experiments were run using the open sourced code from [10] on a 8×8 Cloud TPU-V2 configuration.

To underscore the importance of our memory savings in the context of very large models, we report additional results on the number of steps required for reaching a given solution quality for various batch sizes. We chose a solution quality of 70% Masked-LM accuracy on the holdout set, which Adam and AdaGrad reached at 500k steps. We use Cloud TPU-v3 device which has 16Gib per core for this experiment. We measured the number of steps SM3 needed to reach this accuracy as a function of the batch size. Our results are presented in Figure 3. SM3 scaled almost linearly with the batch size, up to a size of 2^{16} , at which point the training program reached the limits of memory available on hardware. We also found that SM3 came out ahead in terms of wall-time: with the same batch size, a step of SM3 was faster than Adam's by 3%, and doubling the batch size allowed it to reach the same solution quality in almost 35% less wall-time for the same computational budget.

5.3 AmoebaNet-D on ImageNet

Finally, we report results from a different domain: image classification on ImageNet [20] with the state-of-the-art AmoebaNet-D architecture [18], that has recently won the Stanford DAWNBench competition [9]. We compared SM3 with SGD with momentum (Adam performed poorly on this task). SM3 performed very well in this task and achieved improved convergence to state-of-the-art performance, reaching 78.71% top-1 and 94.31% top-5 test accuracies. The fully detailed convergence plots are provided in the full version of the paper [3].

6 Summary

Motivated by the large increase in models sizes and the huge amounts of memory required for training them, we have presented a new memory-efficient adaptive optimization algorithm for stochastic optimization called SM3. We demonstrated empirically that SM3 can be used effectively in training

modern mammoth-sized models and dramatically decrease memory overhead. Utilizing the freed memory for increasing the batch size, our experiments indicate that this saving can also lead to significant improvements in performance. Our theoretical investigation focused on convex objectives. As with many other optimization scenarios, we believe the analysis of convex memory-efficient adaptive optimization could serve as a basis for understanding non-convex settings.

Our memory savings virtually eliminate the overhead coming from the second-order statistics γ_t with little and often no impact on convergence. Additional and potentially substantial improvements in memory consumption could come from compressing or sketching the momentum terms employed by virtually all first-order optimizers used in practice. We leave the exploration of this promising direction for future work.

Acknowledgements

We would like to thank Luke Metz, Kunal Talwar and Yonghui Wu for numerous helpful discussions and suggestions. Special thanks go to Samy Bengio who made it possible for us to conduct large scale experiments on a tight schedule. We would also like to thank Zhifeng Chen for coming up with the shorthand 'SM3'.

References

- [1] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. Tensorflow: A system for largescale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, 2016.
- [2] N. Agarwal, B. Bullins, X. Chen, E. Hazan, K. Singh, C. Zhang, and Y. Zhang. The case for full-matrix adaptive regularization. *CoRR*, abs/1806.02958, 2018.
- [3] R. Anil, V. Gupta, T. Koren, and Y. Singer. Memory-efficient adaptive optimization for largescale learning. arXiv preprint arXiv:1901.11150, 2019.
- [4] R. Anil, V. Gupta, T. Koren, and Y. Singer. SM3 tensorflow optimizer. <u>https://github.com/google-research/google-research/tree/master/sm3</u>, 2019.
- [5] P. Auer, N. Cesa-Bianchi, and C. Gentile. Adaptive and self-confident on-line learning algorithms. *Journal of Computer and System Sciences*, 64(1):48–75, 2002.
- [6] N. Cesa-Bianchi, A. Conconi, and C. Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.
- [7] M. Charikar, K. Chen, and M. Farach-Colton. Finding frequent items in data streams. In Proceedings of the 29th International Colloquium on Automata, Languages and Programming, ICALP '02, pages 693–703, Berlin, Heidelberg, 2002. Springer-Verlag. ISBN 3-540-43864-5. URL http://dl.acm.org/citation.cfm?id=646255.684566.
- [8] M. X. Chen, O. Firat, A. Bapna, M. Johnson, W. Macherey, G. Foster, L. Jones, M. Schuster, N. Shazeer, N. Parmar, A. Vaswani, J. Uszkoreit, L. Kaiser, Z. Chen, Y. Wu, and M. Hughes. The best of both worlds: Combining recent advances in neural machine translation. In *Proceedings* of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, pages 76–86, 2018.
- [9] C. Coleman, D. Kang, D. Narayanan, L. Nardi, T. Zhao, J. Zhang, P. Bailis, K. Olukotun, C. Re, and M. Zaharia. Analysis of dawnbench, a time-to-accuracy machine learning performance benchmark. arXiv preprint arXiv:1806.01427, 2018.
- [10] J. Devlin, M. Chang, K. Lee, and K. Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018.
- [11] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.

- [12] V. Gupta, T. Koren, and Y. Singer. Shampoo: Preconditioned stochastic tensor optimization. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 1842–1850, 2018.
- [13] E. Hazan. Introduction to online convex optimization. Foundations and Trends in Optimization, 2(3-4):157–325, 2016.
- [14] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, et al. In-datacenter performance analysis of a tensor processing unit. In *Computer Architecture (ISCA), 2017 ACM/IEEE 44th Annual International Symposium on*, pages 1–12. IEEE, 2017.
- [15] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] H. B. McMahan and M. Streeter. Adaptive bound optimization for online convex optimization. COLT 2010, page 244, 2010.
- [17] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language models are unsupervised multitask learners. 2019.
- [18] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le. Regularized evolution for image classifier architecture search. arXiv preprint arXiv:1802.01548, 2018.
- [19] S. J. Reddi, S. Kale, and S. Kumar. On the convergence of adam and beyond. 2018.
- [20] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [21] M. Schuster and K. Nakajima. Japanese and korean voice search. In *ICASSP*, pages 5149–5152. IEEE, 2012.
- [22] S. Shalev-Shwartz. Online learning and online convex optimization. Foundations and Trends in Machine Learning, 4(2):107–194, 2012.
- [23] N. Shazeer and M. Stern. Adafactor: Adaptive learning rates with sublinear memory cost. In Proceedings of the 35th International Conference on Machine Learning, ICML 2018, pages 4603–4611, 2018.
- [24] J. Shen, P. Nguyen, Y. Wu, Z. Chen, et al. Lingvo. https://github.com/tensorflow/ lingvo.
- [25] R. Spring, A. Kyrillidis, V. Mohan, and A. Shrivastava. Compressing gradient optimizers via count-sketches. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5946–5955, Long Beach, California, USA, 09–15 Jun 2019. PMLR. URL http://proceedings.mlr.press/v97/spring19a.html.
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [27] A. Vaswani, S. Bengio, E. Brevdo, F. Chollet, A. N. Gomez, S. Gouws, L. Jones, L. Kaiser, N. Kalchbrenner, N. Parmar, R. Sepassi, N. Shazeer, and J. Uszkoreit. Tensor2tensor for neural machine translation. *CoRR*, abs/1803.07416, 2018. URL http://arxiv.org/abs/1803. 07416.
- [28] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, et al. Google's neural machine translation system: Bridging the gap between human and machine translation. arXiv preprint arXiv:1609.08144, 2016.
- [29] Y. Zhu, R. Kiros, R. Zemel, R. Salakhutdinov, R. Urtasun, A. Torralba, and S. Fidler. 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, 2015.