

Adaptivity without Compromise: A Momentumized, Adaptive, Dual Averaged Gradient Method for Stochastic Optimization

Aaron Defazio

Facebook AI Research, New York

ADEFAZIO@FB.COM

Samy Jelassi

Princeton University, Princeton

SJELASSI@PRINCETON.EDU

Editor: Animashree Anandkumar

Abstract

We introduce MADGRAD, a novel optimization method in the family of AdaGrad adaptive gradient methods. MADGRAD shows excellent performance on deep learning optimization problems from multiple fields, including classification and image-to-image tasks in vision, and recurrent and bidirectionally-masked models in natural language processing. For each of these tasks, MADGRAD matches or outperforms both SGD and ADAM in test set performance, even on problems for which adaptive methods normally perform poorly. An implementation is available at <https://github.com/facebookresearch/madgrad>¹

Keywords: Stochastic Optimization, AdaGrad, Optimization, Optimization for machine learning

1. Introduction

Optimization for deep learning forms a relatively new and growing sub-field in the optimization community. Compared to classical first order optimization, deep learning problems introduce additional concerns which require new tools to overcome. Deep learning problems are characterized by very large parameter vector sizes D , making it computationally infeasible to store matrices of size $D \times D$, and even “limited memory” approaches can be impractical for problems such as the 100+ billion parameter models currently being explored (Rajbhandari et al., 2019; Brown et al., 2020). The practical limit on these problems is storage that is fixed at a small multiple of the parameter vector size.

For this reason, diagonal scaling approaches have become the industry standard for deep learning. In this class of methods, adaptivity is performed independently for each coordinate, so that memory usage scales as $O(D)$. We consider Adam (Kingma and Ba, 2014) the benchmark method in this class; it has seen widespread adoption, and there are no alternative adaptive methods that consistently out-perform it (Choi et al., 2020; Schmidt et al., 2020).

1. This work revises and extends the technical report Jelassi and Defazio (2020) with a new algorithm and further theory and experiments.

Adam builds upon a rich history of diagonal adaptive methods. The AdaGrad method (Duchi et al., 2011) introduced a principled approach to diagonal adaptivity, that arises naturally as a simplification of a full-matrix adaptivity scheme. This approach is clearly motivated and yields natural convergence rate bounds for convex losses. Also within this family, the RMSProp method (Tieleman and Hinton, 2012) arose as a well-performing empirical method in this class, albeit with little theoretical motivation. The development of the Adam method can be seen as a natural extension of the scaling used in RMSProp to include a form of momentum, as well as a stabilizing “bias-correction” that significantly dampens the adaptivity and step-size during the early stages of optimization.

Despite its widespread success, Adam is far from a panacea for deep learning optimization. Wilson et al. (2017) show that Adam as well as other common adaptive optimizers converge to bad local minima on some important problems, such as the widely studied problem of image classification. This has led to the general claim that adaptive methods generalize poorly. As we will show, this is not necessarily the case. The method we develop in this work combines adaptivity with strong generalization performance.

Our MADGRAD (Momentumized, Adaptive, Dual averaged GRADient) method performs consistently at a state-of-the-art level across a varied set of realistic large-scale deep learning problems, without requiring any more tuning than Adam. MADGRAD is constructed from the lesser-used dual averaging form of AdaGrad, through a series of direct and systematic changes that adapt the method to deep learning optimization.

2. Problem Setup

We consider the stochastic optimization framework, where the goal is to minimize a parameterized function

$$f(x) = \mathbb{E}_{\xi} [f(x, \xi)],$$

where $x \in \mathbb{R}^D$, and each ξ is a random variable drawn from a fixed known distribution. In the case of empirical risk minimization, ξ is a data-point drawn from the data distribution, typically further processed by a stochastic data-augmentation procedure. At each step k , a stochastic optimization algorithm is given ξ_k and has access to $f(x_k, \xi_k)$ and $\nabla f(x_k, \xi_k)$ for a pre-specified iterate x_k .

3. Related Work

The theory of adaptive methods for non-convex optimization is still in its infancy. The current best known convergence theory for Adam due to Défossez et al. (2020) greatly improves over earlier theory (Zou et al., 2019b), but has the important caveat that it requires momentum values of the order $\beta = 1 - 1/N$ for N iterations, which is far from the values used in practice, which are of the order $\beta = 0.9$ to $\beta = 0.99$. Results for these settings may not be possible, as Reddi et al. (2018) show via a counter-example that Adam may fail to converge under common parameter settings, even in the convex case. When β_1 & β_2 are small, the Adam update is close to sign-sgd (i.e. $x_{k+1} = x_k - \gamma \text{sign}(\nabla f(x_k, \xi_k))$), a method that also fails to converge in the general stochastic case (Balles and Hennig, 2018), although some theory is possible under a large batch assumption Bernstein et al. (2018) where the behavior is closer to the non-stochastic case.

AdaGrad’s convergence in the non-convex case has also been studied. Ward et al. (2019) establish convergence for a restricted variant where only a global step size is adaptively updated. Li and Orabona (2019) establish almost sure convergence for a variant of AdaGrad where the most recently seen gradient is omitted from the denominator. Convergence with high probability is also established for a variant with global rather than coordinate-wise step size. More recently Zhou et al. (2020) and Zou et al. (2019a) establish convergence of non-momentum and momentum variants respectively, although with bounds that are much worse than established by Défossez et al. (2020), who also cover AdaGrad in their analysis.

Weighted AdaGrad as we use in this work has been explored to varying degrees before, including the non-convex case in the aforementioned work by Zou et al. (2019a), and the convex case by Levy et al. (2018). Weighting is particularly interesting in the strongly convex case, where weights such as $\lambda_k \propto k^2$ can be used to achieve accelerated convergence. Neither of these works cover the dual averaged form of AdaGrad which we explore.

4. Adaptivity in deep learning beyond Adam

To understand the motivation and design of the MADGRAD method, a clear understanding of the short-comings of existing methods is needed. Consider Adam, the most heavily used adaptive method in practice. Although it works remarkably well on some important problems, it also suffers from the following issues:

- It greatly under-performs the non-adaptive SGD-M method in a number of important situations including the widely studied ImageNet training problem.
- Problems can be constructed on which it will fail to converge entirely, even in the convex setting.
- The exponential moving average updates used are non-sparse when given sparse gradients, which makes the method poorly suited to sparse problems.

Due to these issues, Adam doesn’t quite reach the goal of being a general-purpose deep learning optimizer. The MADGRAD method is directly designed to address these issues. MADGRAD:

- Achieves state-of-the-art performance across problems traditionally tackled by Adam, while simultaneously achieving state-of-the-art on problems where Adam normally under-performs.
- Has provable and strong convergence theory on convex problems.
- Is directly applicable to sparse problems when momentum is not used.

5. Design

The MADGRAD method is the combination of a number of techniques that individually address separate short-comings in the AdaGrad method when applied to deep learning optimization problems. By building upon a method with known convergence theory, we are able to construct a method that is still provably convergent (under convexity assumptions)

without sacrificing the practical performance characteristics of Adam. We will detail each of these techniques in turn, to build up MADGRAD from its foundations.

5.1 Dual averaging for deep learning

MADGRAD is based upon the dual averaging formulation of AdaGrad, rather than the mirror descent formulation. Although the original seminal work on AdaGrad (Duchi et al., 2011) presents the dual averaging formulation with equal weight as the mirror descent (MD) form, the dual averaging form has seen virtually no use for deep learning optimization. The AdaGrad implementations available in major deep learning frameworks (PyTorch, Tensorflow) contain the mirror descent form only. This is despite the theory presented for the dual averaging formulation being arguably more elegant than the mirror descent theory. The dual averaging form of AdaGrad satisfies the following bound:

$$\sum_{i=1}^k f(x_i) - f(x_*) \leq \frac{1}{\gamma} \psi_k(x_*) + \frac{\gamma}{2} \sum_{i=1}^k \|\nabla f_i(x_i)\|_{\psi_{i-1}^*}^2$$

Whereas the mirror descent form satisfies the following more complex bound, involving the Bregman divergence of ψ :

$$\sum_{i=1}^k f(x_i) - f(x_*) \leq \frac{1}{\gamma} B_{\psi_1}(x_*, x_1) + \frac{1}{\gamma} \sum_{i=1}^{k-1} [B_{\psi_{i+1}}(x_*, x_{i+1}) - B_{\psi_i}(x_*, x_{i+1})] + \frac{\gamma}{2} \sum_{i=1}^k \|\nabla f_i(x_i)\|_{\psi_i^*}^2.$$

Given the clear advantage in terms of theoretical simplicity, why then are dual averaging approaches not used more widely? We believe this is due to a number of misconceptions. The first misconception is that dual averaging is only interesting in the composite optimization setting, where sophisticated regularizers are used to encourage sparsity or induce other properties of the solution. It is true that for smooth non-stochastic optimization, gradient descent and mirror descent coincide (under optimal hyper-parameters). However, when the objective is stochastic or non-smooth, the methods become distinct, and actually behave quite differently.

Dual averaging has the general form, given a proximal function ψ :

$$\begin{aligned} g_k &= \nabla f(x_k, \xi_k), \\ s_{k+1} &= s_k + \lambda_k g_k, \\ x_{k+1} &= \arg \min_x \{ \langle s_{k+1}, x \rangle + \beta_{k+1} \psi(x) \}. \end{aligned} \tag{1}$$

The gradient buffer s_0 is initialized as the zero vector. The simplest form of dual averaging occurs when the standard Euclidean squared norm is used: $\psi(x) = \frac{1}{2} \|x - x_0\|^2$, and $\lambda_k = 1$ in which case the method takes the form:

$$x_{k+1} = x_0 - \frac{1}{\beta_{k+1}} \sum_{i=0}^k g_i. \tag{2}$$

If the objective is either non-smooth or stochastic (or both), β sequences of the form $\beta_{k+1} = \sqrt{k+1}$ give a convergent method. Although Equation 2 has little resemblance to SGD as written, SGD's update:

$$x_{k+1} = x_k - \gamma_k \nabla f(x_k, \xi_k),$$

can be written in the more comparable form:

$$x_{k+1} = x_0 - \sum_{i=0}^k \gamma_i g_i. \tag{3}$$

where to achieve convergence without a fixed stopping time, a step size of the form $\gamma_i \propto 1/\sqrt{i+1}$ is standard. Comparing SGD and DA at a step k , it's clear that the weighting sequence used by SGD places a smaller weight on newer g_i in the summation compared to earlier g_i , whereas the sequence used by DA places equal weight on all g_i . This difference is key to understanding why methods in the DA family behaves differently from SGD in practice, even without additional regularization or non-Euclidean proximal functions.

The second misconception arises from implementing the dual averaging form of AdaGrad without considering what modifications need to be made for the deep learning setting. The algorithm as originally stated, uses an initial point of the origin $x_0 = 0$, and a proximity function $\psi_t(x) = \frac{1}{2} \langle x, H_t x \rangle$ that is quadratic, but centered around the origin. It is well known that neural network training exhibits pathological behavior when initialized at the origin, and so naive use of this algorithm does not perform well. When centering around 0, we have observed severely degraded empirical performance and a high risk of divergence. Instead, a proximity function centered about x_0 needs to be used:

$$\psi_t(x) = \frac{1}{2} \langle x - x_0, H_t (x - x_0) \rangle,$$

with initialization of x_0 following standard conventions for the network being trained.

5.2 Dual averaging generalizes well

In addition to the theoretical advantages of dual averaging methods, we have also observed that they also enjoy a strong practical advantage in the form of better generalization performance. Dual averaging based methods include a form of implicit regularization, which we believe is a crucial factor contributing to their good generalization performance. To see this, consider the classical dual averaging update:

$$x_{k+1} = x_0 - \frac{1}{\sqrt{k+1}} \sum_{i=0}^k g_i,$$

This update can be written in a form closer to the SGD update by substituting for x_0 :

$$\begin{aligned} x_{k+1} &= \left(x_k + \frac{1}{\sqrt{k}} \sum_{i=0}^{k-1} g_i \right) - \frac{1}{\sqrt{k+1}} \sum_{i=0}^k g_i, \\ &= x_k - \frac{1}{\sqrt{k+1}} \left[g_k - \left(\frac{\sqrt{k+1}}{\sqrt{k}} - 1 \right) \sum_{i=0}^{k-1} g_i \right], \\ &= x_k - \frac{1}{\sqrt{k+1}} \left[g_k + \left(\sqrt{k+1} - \sqrt{k} \right) (x_k - x_0) \right]. \end{aligned}$$

Since $\sqrt{k+1} - \sqrt{k} \approx 1/(2\sqrt{k+1})$, the behavior of dual averaging resembles a SGD step with a step-dependent regularizer:

$$\frac{1}{4\sqrt{k}} \|x_k - x_0\|^2,$$

which decays in strength during the course of optimization. We speculate that the indirect decaying regularization inherent in dual averaging methods may explain why MADGRAD also requires less decay than other methods to match their performance. The strong initial regularization may have a positive effect during early iterations, while not negatively affecting the ability of the model to fit to the data during the later "fine-tuning" epochs. Given the practical advantages we observe in our experiments, we believe further research into the effect of using stronger regularization at the early stages of optimization may be interesting more generally.

In the online learning setting, where dual averaging is known as follow-the-regularized-leader, the differences between the two approaches are more heavily explored. There are situations where the DA variant of a method is clearly superior to the mirror descent variant of the same method. For instance, when the stopping time is not known in advance, scale-free mirror descent may suffer from linear regret, whereas DA does not (Orabona and Pál, 2015).

Another advantage of DA approaches are their dependence on the sub-optimality of the initial point. They have regret bounds that depend multiplicatively on the initial distance to the solution $\|x_0 - x_*\|^2$ compared to $\max_{k=0\dots n} \|x_k - x_0\|^2$ for mirror descent. The MD bound may be much worse in cases where the iterate sequence is poorly behaved.

5.3 λ sequences for deep learning

Even with this modification, dual averaging both with and without adaptivity is not competitive with SGD on standard benchmark problems such as CIFAR10, as shown in Figure 1. The top row shows AdaGrad and DA methods using a flat learning rate schedule, and the bottom row shows a stage-wise schedule. SGD is shown as a baseline. For the DA family methods, λ_k is decreased for the stage-wise schedules. Both AdaGrad, DA and AdaGrad-DA under-perform SGD with either learning rate schedule. Part of this performance gap can be attributed to the fact that each of these methods either implicitly or explicitly use a $1/\sqrt{i+1}$ learning rate sequence. This sequence is actually **harmful**, as we can confirm by testing SGD using a schedule of the form:

$$\gamma_i = \frac{a}{\sqrt{i+b}},$$

Figure 2 illustrates the learning curves achievable for varying b values on CIFAR-10. Full description of our experimental setup is in Section 7. We performed a hyper-parameter search over a separately for each b , with test accuracy as the target quantity. All sqrt-decay sequences are significantly worse than the baseline stage-wise schedule, where the learning rate is decreased 10 fold at epochs 150 and 225. We speculate that the sqrt-decay sequences result in convergence that is too rapid, skipping over the initial annealing stage of learning, resulting in convergence to a poor local minima. The AdaGrad and AdaGrad-DA methods also use an implicitly decreasing sequence, although the rate of decrease depends on the

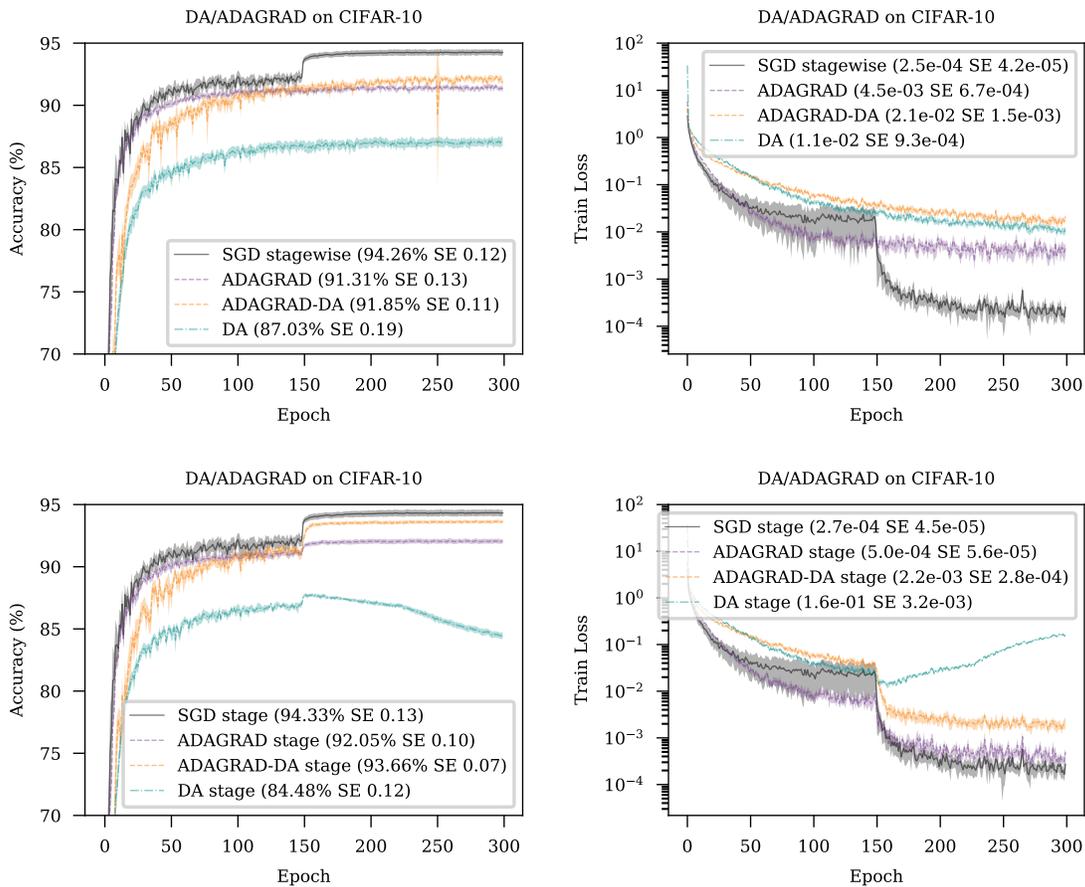


Figure 1: Comparison of SGD without momentum to DA and DA-AdaGrad and AdaGrad on CIFAR-10. Left column is test classification performance, right column is training loss. The "stage" learning rate scheme involves a 10 fold decrease in the learning rate at epochs 150 and 225. See Section 7 for a full description of the experimental setup.

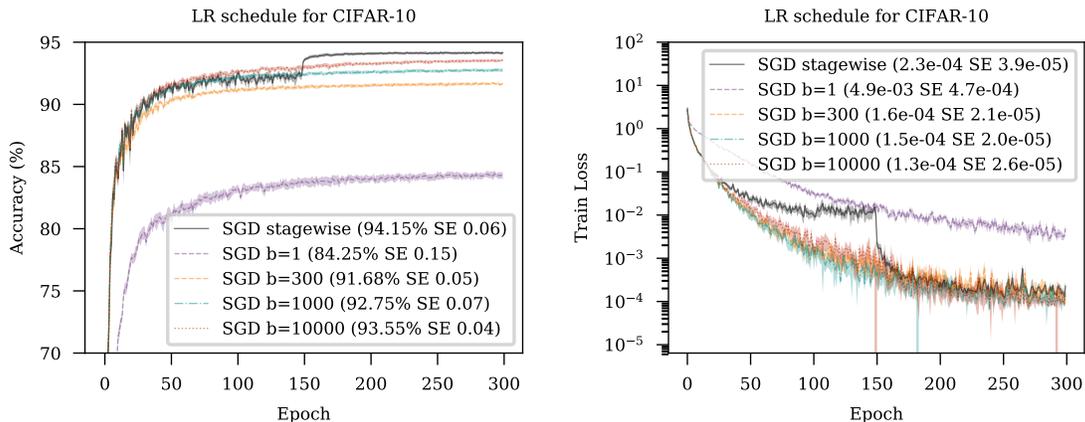


Figure 2: Sqrt-decay learning rate schedules under-perform stage-wise schedules. With batch-size 128 on CIFAR-10. No momentum is used in this comparison. A range of offsets b in the rate $a/\sqrt{i+b}$ were tried with values up to 10,000 shown. Larger values of b up to 100,000 were also tested, they also failed to match the performance of the stage-wise schedule. Left column is test classification performance, right column is training loss.

magnitude of the gradients, which is very problem dependent. If gradients stay of similar magnitude over a particular time-scale, then the rate of decrease will also be a $1/\sqrt{k}$ rate for step k . This step size scheme is also undesirable as prevents the use of standard SGD & Adam step size sequences for choosing the explicit step size constants λ_i .

Since in practice the same learning rate scheme is commonly used when comparing different optimization methods, this schedule contributes to the commonly held perception that AdaGrad is not as effective as other adaptive methods such as Adam.

For the DA method, we propose to remedy this issue by introducing a scaling of the λ values to counter-act the step size sequence. In particular we propose the choice:

$$\lambda_i = (i+1)^{1/2} \gamma_i,$$

where γ is a conventional (SGD/Adam) step size sequence. The advantage of this choice is that the leading term in the sum in Equation 2 has constant weight across k :

$$\begin{aligned} x_{k+1} &= x_0 - \frac{1}{\sqrt{k+1}} \sum_{i=0}^k \lambda_i g_i, \\ &= x_0 - \gamma_k g_k - \frac{1}{\sqrt{k+1}} \sum_{i=0}^{k-1} \lambda_i g_i, \end{aligned}$$

mirroring the behavior of SGD during a constant step size phase, but retaining the $\sqrt{k+1}$ decay of past gradients. This simple change is sufficient to greatly improve the test-set performance of DA when using the same learning rate schedule as SGD.

Another advantage of this sequence is that it will place higher weights on latter gradients in the final convergence rate bound. This makes no difference if we expect gradients to be of similar magnitude at all stages of optimization (which can happen for non-smooth problems in the worse case), but in practice even for non-smooth objectives the gradient typically shrinks to some degree during optimization, leading to tighter bounds when using a forward weighted lambda sequence. We discuss this difference further in Section 6.

5.4 Momentum

The use of momentum on top of SGD is known to be highly beneficial, if not crucial, for deep learning optimization across a wide variety of architectures and problem settings (Sutskever et al., 2013). Given how crucial it can be to maintaining competitive performance, we now examine how we can add a form of momentum to the dual averaging updates, and latter the AdaGrad updates.

We will consider an update of the following form, which was first explored in this general form by Nesterov and Shikhman (2015) under the name Dual Averaging with Double Averaging:

$$\begin{aligned}
 g_k &= \nabla f(x_k, \xi_k), \\
 s_{k+1} &= s_k + \lambda_k g_k, \\
 z_{k+1} &= \arg \min_x \{ \langle s_{k+1}, x \rangle + \beta_{k+1} \psi(x) \}, \\
 x_{k+1} &= (1 - c_{k+1}) x_k + c_{k+1} z_{k+1}.
 \end{aligned}
 \tag{4}$$

The essential idea behind this algorithm is simple. Instead of evaluating the gradient at each step at the value of the argmin operation as with regular DA, instead it's evaluated at a moving average point instead. This serves to smooth the iterate sequence. This technique has the advantage in the convex setting of making it possible to prove convergence properties of the last iterate x_{k+1} rather than the average iterate $\bar{x}_{k+1} = \frac{1}{k+1} \sum_{i=0}^k x_i$. Essentially the averaging operation is incorporated into the algorithm itself.

Momentum is normally thought of as performing more than just a smoothing of the iterate sequence, although a line of recent research has shown that inline averaging of the above form is actually exactly equivalent to momentum (Sebbouh et al., 2020; Defazio, 2020). This is clearly illustrated when momentum is added on top of SGD, where inline averaging:

$$\begin{aligned}
 z_{k+1} &= z_k - \eta_k \nabla f(x_k, \xi_k), \\
 x_{k+1} &= (1 - c_{k+1}) x_k + c_{k+1} z_{k+1},
 \end{aligned}$$

is actually exactly equivalent to more common equational forms of momentum:

$$\begin{aligned}
 m_{k+1} &= \beta_k m_k + \nabla f(x_k, \xi_k), \\
 x_{k+1} &= x_k - \alpha_k m_{k+1},
 \end{aligned}$$

for appropriate choices of the hyper-parameters. In the convex setting the advantage of this form arises when $c_{k+1} = \frac{1}{k+1}$, which corresponds to an equal weighted moving average $x_{k+1} = \frac{1}{k+1} \sum_{i=0}^k z_i$. Under this setting convergence of the last iterate can be shown just

as when this kind of averaging is used with dual averaging (Defazio and Gower, 2020). In the non-convex setting, constant c_{k+1} values, which correspond to an exponential moving average, appear to be the best choice (Defazio, 2020).

5.5 Adaptivity

Our goal is to combine these ideas together with the adaptivity technique from the AdaGrad method. The dual averaging form of coordinate-wise AdaGrad has the following form:

$$x_{k+1} = x_0 - \frac{1}{\sqrt{\sum_{i=0}^k \gamma_i g_i^2}} \circ \sum_{i=0}^k \gamma_i g_i,$$

where \circ represents the element-wise (Hadamard) product, and γ is a fixed step size hyperparameter. There are many different ways of combining this kind of coordinate-wise adaptivity with the weighted gradient sequence $\lambda_i = \sqrt{i+1}$ that we have proposed. Due to the flexibility of the dual averaging framework, it's possible to prove a convergence rate of some form for practically any choice of denominator sequence. However, we must take into consideration that we also want to maintain the magnitude of the “effective” step size, as discussed in Section 5.3.

We also need to ensure that the weighted dominator includes γ_i not just $\sqrt{i+1}$, as this mitigates a problem illustrated for DA in Figure 1: when λ is decreased 10 fold at epoch 150, the method starts to diverge. At this point, the β sequence continues to decrease at a square-root rate, while the sum-of-gradients starts growing ten times slower. This results in the method shrinking the iterates towards x_0 far to strongly.

We review a number of possible alternatives below and discuss their practicality.

5.5.1 UNWEIGHTED DENOMINATOR

One possibility is keep the denominator the same but just weight the gradients in the sum:

$$x_{k+1} = x_0 - \frac{1}{\sqrt{\sum_{i=0}^k \gamma_i g_i^2}} \circ \sum_{i=0}^k (i+1)^{1/2} \gamma_i g_i,$$

This is appealing as it maintains the constant effective step size property, however the resulting convergence rate bound derivable from this form depends on $\sqrt{\sum_{i=0}^k \gamma_i g_i^2}$ rather than $\sqrt{\sum_{i=0}^k (i+1)^{1/2} \gamma_i g_i^2}$, which defeats the purpose of using a front-weighted gradient sequence.

5.5.2 WEIGHTED DENOMINATOR

We can weight the gradient sequence in the denominator by λ also:

$$x_{k+1} = x_0 - \frac{1}{\sqrt{\sum_{i=0}^k (i+1)^{1/2} \gamma_i g_i^2}} \circ \sum_{i=0}^k (i+1)^{1/2} \gamma_i g_i.$$

This form does not maintain a constant effective step size, which results in poor empirical performance. We experimented with mitigations such as adding additional terms to the numerator that would counteract this growth, however this still resulted in unsatisfactory empirical results.

5.5.3 WEIGHTED NUMERATOR

The AdaGrad variant proposed by Zou et al. (2019a) uses a weighting scheme where the weights λ_k are included in the numerator as well as the denominator:

$$x_{k+1} = x_0 - \frac{\gamma_i}{\sqrt{t}} \frac{\sqrt{\sum_{i=0}^k \lambda_i}}{\sqrt{\sum_{i=0}^k \lambda_i g_i^2}} \circ g_i = x_0 - \frac{\gamma_i}{\sqrt{t}} \frac{\sqrt{\sum_{i=0}^k (i+1)^{1/2}}}{\sqrt{\sum_{i=0}^k (i+1)^{1/2} g_i^2}} \circ g_i.$$

This numerator is proportional to $t^{1/4}$. To adapt this sequence to dual averaging, we must include a step size parameter in the weights. It's unclear exactly how to do this in a way that maintains the effective step size property, since if $\lambda_i \propto \gamma_i$ then the step size will cancel between the numerator and denominator.

5.5.4 MADGRAD'S CUBE-ROOT DENOMINATOR

To maintain the correct effective step size we propose the use of a cube root instead:

$$x_{k+1} = x_0 - \frac{1}{\sqrt[3]{\sum_{i=0}^k (i+1)^{1/2} \gamma_i g_i^2}} \circ \sum_{i=0}^k (i+1)^{1/2} \gamma_i g_i. \quad (5)$$

Although this modification appears ad-hoc, the use of a cube root here can actually be motivated by a similar argument used to motivate the standard square-root formulation. Duchi et al. (2011) consider the following minimization problem over a D dimensional vector s :

$$\min_s \sum_{i=0}^k \sum_{d=0}^D \frac{g_{id}^2}{s_d}, \quad \langle 1, s \rangle \leq c, \quad \forall d : s_d > 0,$$

which is solved by $s_d \propto \sqrt{\sum_{i=0}^k g_{id}^2}$. The motivation for this surrogate problem is to minimize weighted square norm of the gradients in hind-sight. Rather than a linear penalty on the size of s , which when combined with the positivity constraint is just a L1 norm penalty $\|s\|_1 \leq c$, if we instead use a L2 norm penalty:

$$\min_s \sum_{i=0}^k \sum_{d=0}^D \frac{g_{id}^2}{s_d}, \quad \|s\|_2^2 \leq c, \quad \forall d : s_d > 0$$

then we recover a cube-root solution $s_d \propto \sqrt[3]{\sum_{i=0}^k g_{id}^2}$. We show this in the Appendix. The cube root maintains the effective step size as can be seen by considering that $\sum_i^k (i+1)^{1/2} \propto (k+1)^{3/2}$ which after the cube root operation gives the necessary $\sqrt{k+1}$ scaled denominator required to cancel against λ 's square-root growth.

Algorithm 1 MADGRAD

Require: γ_k stepsize sequence, c_k momentum sequence, initial point x_0 , epsilon ϵ

- 1: $s_0 : d = 0, \nu_0 : d = 0$
- 2: **for** $k = 0, \dots, T$ **do**
- 3: Sample ξ_k and set $g_k = \nabla f(x_k, \xi_k)$
- 4: $\lambda_k = \gamma_k \sqrt{k+1}$
- 5: $s_{k+1} = s_k + \lambda_k g_k$
- 6: $\nu_{k+1} = \nu_k + \lambda_k (g_k \circ g_k)$
- 7:

$$z_{k+1} = x_0 - \frac{1}{\sqrt[3]{\nu_{k+1}} + \epsilon} \circ s_{k+1}$$

- 8: $x_{k+1} = (1 - c_{k+1})x_k + c_{k+1}z_{k+1}$.
 - 9: **end for**
 - 10: **return** x_T
-

One disadvantage of this weighting is that it results in a final convergence rate bound that is not fully adaptive in the sense that the choices of global step size will depend on an expression involving the gradient norms. We don't believe this is a significant problem given that the choice of step size still depends on other unknown quantities even when using a fully adaptive sequence such as the function sub-optimality gap and gradient bound G .

6. Convergence Theory

The MADGRAD algorithm, combining the discussed ideas, is listed in Algorithm 1. In order to establish convergence results for potentially non-smooth functions, we rely on a bounded gradient assumption:

$$\|\nabla f(x, \xi)\|_\infty \leq G \text{ for all } x, \xi.$$

We also assume each $f(\cdot, \cdot)$ is proper and convex in x over all \mathbb{R}^D . Our analysis uses a slight variant of Algorithm 1, where the denominator includes an extra term $\lambda_k G^2$:

$$z_{k+1} = x_0 - \frac{1}{\sqrt[3]{\lambda_{k+1} G^2 + \nu_{k+1}}} \circ s_{k+1}, \quad (6)$$

A similar term is also needed by the original DA-AdaGrad method in Duchi et al. (2011). We don't believe this term plays an important role in practice as its magnitude quickly diminishes, and so we have not included this term in Algorithm 1. Orabona and Pál (2015) describe a way to remove this term from AdaGrad DA methods, at the expense of a more complicated analysis. A per-coordinate upper bound G_d may be used instead of G to further tighten the theory.

Theorem 1. *After k steps of MADGRAD using the update in Equation 6,*

$$\mathbb{E}[f(x_k) - f(x_*)] \leq \frac{6}{k^{1/2}} \|x_0 - x_*\|_2 G D^{1/2},$$

if $c_k = \frac{3/2}{k+3/2}$ and

$$\gamma = \frac{1}{k^{3/4} D^{3/4} G^{1/2}} \|x_0 - x_*\|_2^{3/2}.$$

This bound is very loose. It results from the application of $\nabla f(x, \xi)_i \leq G$ to bound each index of the gradient at each time-step separately, which does not capture any of the adaptivity of the convergence rate. We discuss more precise bounds below. Note that $\|g\|_2 \leq D^{1/2} \|g\|_\infty = GD^{1/2}$, so the dependence on dimensionality here is comparable to bounds established for non-adaptive stochastic methods which have bounds on the 2-norm of the gradient on the right instead. Note also that we recommend using a flat $c_k = c$ momentum for non-convex problems, this decaying rate is only optimal in the convex case. A value of $c = 0.1$ corresponds to the $\beta = 0.9$ momentum commonly used with SGD and Adam.

6.1 Adaptivity

To understand the adaptivity of the method at a more granular level, we can express the convergence rate as:

$$\begin{aligned} \mathbb{E} [f(x_k) - f(x_*)] &\leq \frac{3}{\gamma} \frac{1}{(k+1)^{3/2}} \sum_{d=0}^D \left(\mathbb{E} \left[\lambda_k \left(\sum_{i=0}^k \lambda_i g_{id}^2 \right)^{2/3} \right] \right) \\ &\quad + \frac{3}{\gamma} \frac{1}{(k+1)^{3/2}} \sum_{d=0}^D (x_{0x} - x_{*d})^2 \mathbb{E} \left(\lambda_{k+1} G^2 + \sum_{i=0}^k \lambda_i g_{id}^2 \right)^{1/3} \end{aligned}$$

The convergence rate heavily depends on a weighted sequence:

$$\sum_{d=0}^D \sum_{i=0}^k \lambda_i g_{id}^2 = \gamma \sum_{d=0}^D \sum_{i=0}^k (i+1)^{1/2} g_{id}^2,$$

rather than an unweighted sum $\sum_{d=0}^D \sum_{i=0}^k g_{id}^2$ used in AdaGrad. This is key to understanding the performance characteristics of MADGRAD over traditional AdaGrad. In particular, large gradients at the early stages have a smaller effect on the overall bound than they do for AdaGrad. This can be quantified by considering the behavior when the gradient norm bound is time dependent, i.e. $\|\nabla f(x_i, \xi)\|_\infty \leq G_i$. Then as we show in the appendix, for MADGRAD, when using optimal step-sizes:

$$\mathbb{E} [f(x_k) - f(x_*)] \leq \frac{6}{(k+1)^{5/4}} \|x_0 - x_*\|_2 D^{1/2} \left(\sum_{i=0}^k (i+1)^{1/2} G_i^2 \right)^{1/2},$$

whereas for AdaGrad with the use of momentum:

$$\mathbb{E} [f(x_k) - f(x_*)] \leq \frac{6}{(k+1)} \|x_0 - x_*\|_2 D^{1/2} \left(\sum_{i=0}^k G_i^2 \right)^{1/2}.$$

In MADGRAD the effect of an ‘‘outlier’’ G_i that is particularly large at time-step i decays at a faster rate, with a power 5/4 compared to linearly for AdaGrad. Using λ_i with larger power than 1/2 is also possible within our momentumized-dual averaged gradient framework,

which would result in a faster decay. We have found that the $1/2$ factor is a "Sweet-spot", as larger values result in empirically slower convergence. Similar convergence rate bounds can be derived using the same proof technique, although they are prefixed by progressively larger constants (growing factorially in the power) as the power used is increased. In general, the advantage of MADGRAD over AdaGrad manifests in the common situation where the gradients are largest at the early stages of optimization.

6.2 Comparison to Adam

Although Adam is known to potentially diverge, we can consider the theoretical properties of the AMSGrad variant of Adam, which is perhaps the smallest modification to Adam that results in provable convergence. For AMSGrad, parameterized by momentum $\beta_1 \lambda^{i-1}$ at step i , assuming a bounded domain with $R = \max_{x,y} \|x - y\|_\infty$, defining $\gamma = \beta_1 / \sqrt{\beta_2}$, and using step size $\alpha_i = \alpha / \sqrt{i}$ (Reddi et al., 2018):

$$\mathbb{E} \sum_{i=1}^k f(x_i) - f(x_*) \leq \frac{\beta_1 R G}{(1 - \beta_1)^2 (1 - \lambda)^2} + \frac{R \sqrt{T}}{\alpha (1 - \beta_1)} \sum_{d=1}^D (\hat{v}_{k,d})^{1/2} + \frac{\alpha \sqrt{1 + \log k}}{(1 - \beta_1)^2 (1 - \gamma) \sqrt{1 - \beta_2}} \sum_d^D \left(\sum_{i=1}^k g_{id}^2 \right)^{1/2}$$

\hat{v} is the maximum of the exponential moving average of the squared gradients, see Reddi et al. (2018) for further details. This result has a number of shortcomings compared to the MADGRAD. Firstly, note that the momentum term $1 - \beta_1$, comparable to c in MADGRAD divides each term in the bound. This means that momentum hurts rather than improves performance. The dependence on a bounded domain is also an undesirable property compared to MADGRAD, and the convergence theory of MADGRAD avoids log factors.

7. Experimental Results

In our experiments we compared MADGRAD against SGD, Adam and AdaGrad. SGD is known to perform well on computer vision classification problems due to its ability to produce solutions that generalize better than adaptive methods. In contrast, Adam is the method of choice in other domains with structured output where overfitting is less of an issue. We present results across a large number of problems across both categories to validate the general purpose utility of the MADGRAD approach.

In our experiments we use the most common step-size reduction scheme used in the literature for each respective problem. For all algorithms, we performed a learning rate and decay sweep on a grid on intervals of $[1 \times 10^i, 2.5 \times 10^i, 5 \times 10^i]$ for a range of i large enough to ensure the best parameters for each problem and method were considered. We present the results from the best learning rate and decay for each method when considering test set performance. For other hyper-parameters, we used commonly accepted defaults for each respective problem. Full parameter settings used for each method are listed in the appendix. All presented results are averaged over a number of seeds with error bars indicating 2 standard errors. Ten seeds were used for CIFAR-10 and IWSLT14, whereas only five seeds were used for the remaining larger scale problems.

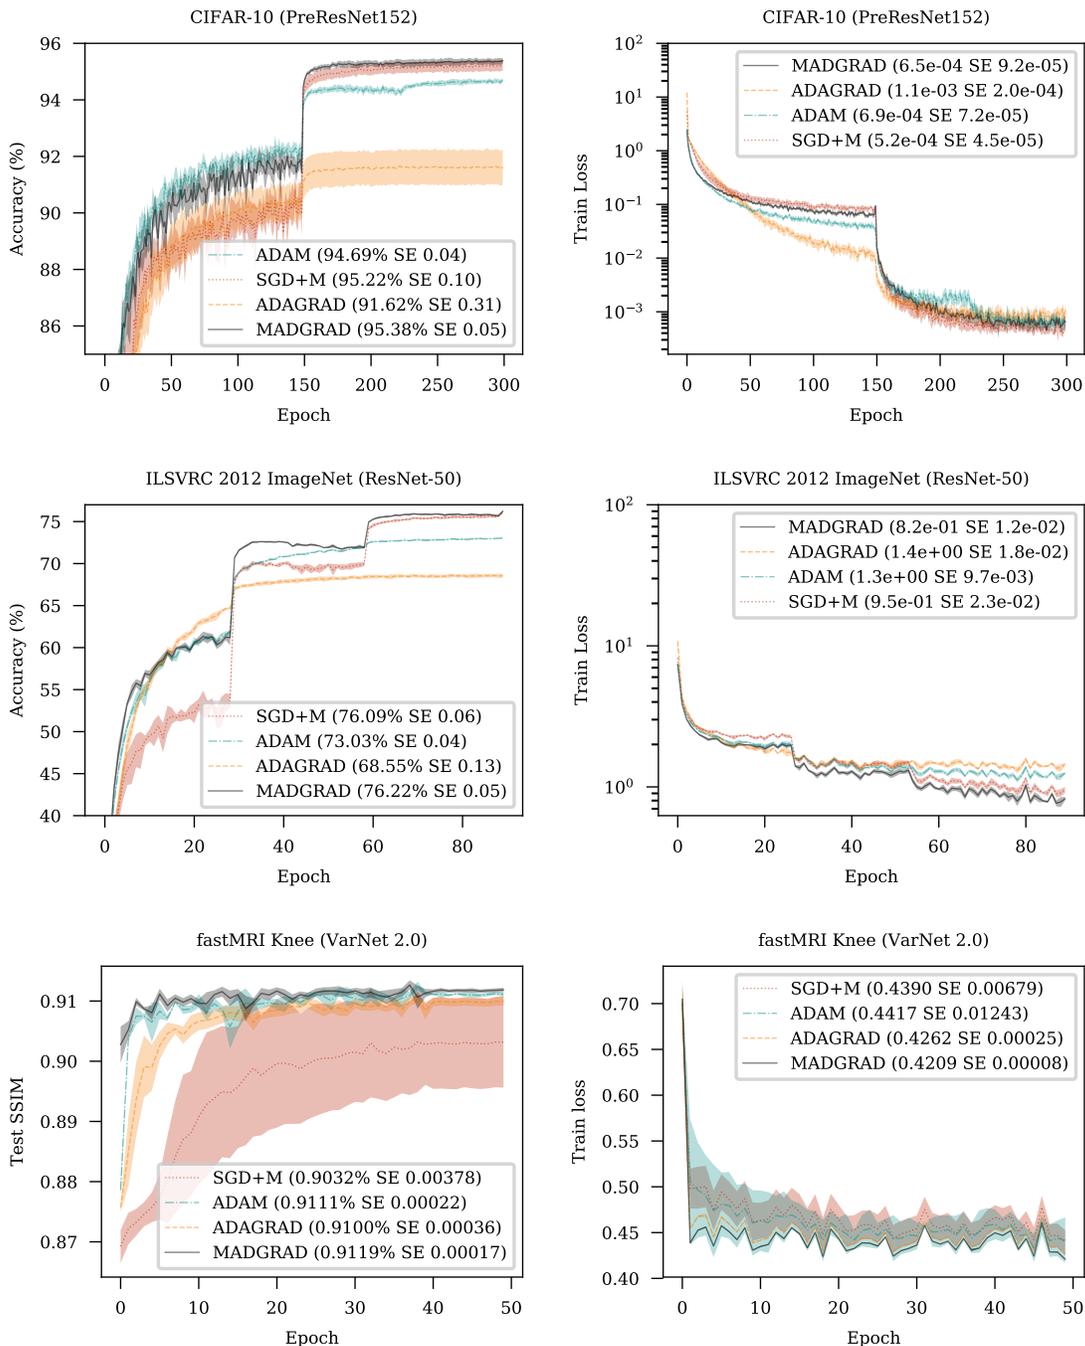


Figure 3: Experimental results for the CIFAR-10, ImageNet and fastMRI Knee problems. Left column shows test set performance and the right column shows training set performance.

7.1 CIFAR10 image classification

CIFAR10 (Krizhevsky, 2009) is an established baseline method within the deep learning community due to its manageable size and representative performance within the class of data-limited supervised image classification problems. It is particularly notable for showing clear differences between adaptive and non-adaptive methods, as the former tend to overfit considerably on this problem. Following standard practice, we apply a data-augmentation step consisting of random horizontal flipping, 4px padding followed by random cropping to 32px at training time only. We used a high-performance pre-activation ResNet architecture (He et al., 2016b) which is known to work well on this problem, consisting of 58,144,842 parameters across 152 layers. The depth of this network is representative of the typical point of diminishing returns for network depth on computer vision problems. As this network is greatly over-parameterized, each method can be expected to fit the training data exactly, achieving near zero loss, even with this data augmentation. For this reason, this task is particularly sensitive to difference in generalization performance of each method.

As illustrated in Figure 3, both Adam and AdaGrad perform poorly on this problem in terms of test accuracy. The under-performance of Adam on this problem is well known (Wilson et al., 2017), and is typically attributed to convergence to poor local minima, as the training set convergence is very rapid initially.

MADGRAD exhibits excellent test accuracy results on this problem, achieving the highest test accuracy among the methods considered. This demonstrates that unlike Adam and AdaGrad, MADGRAD’s adaptivity does not come at the cost of inferior generalization performance.

7.2 ILSVRC 2012 ImageNet image classification

The ImageNet problem (Krizhevsky et al., 2012) is a larger problem more representative of image classification problems encountered in industrial applications where a large number of classes and higher resolution input images are encountered. Like CIFAR10, overfitting can be an issue on this problem for adaptive methods. We ran experiments using the ResNet-50 architecture, which is considered the standard baseline for this problem. This combination of data set and architecture are one of the most studied in all of machine learning, which makes it an ideal testing ground for optimization algorithms.

Our setup used data preprocessing consisting of a mean [0.485, 0.456, 0.406] and std [0.229, 0.224, 0.225] normalization of the three respective color channels, followed by a RandomResizedCrop PyTorch operation to reduce the resolution to 224 pixels followed by a random 50% chance of horizontal flipping. For test set evaluation a resize to 256 pixels followed by a center crop to 224 pixels is used instead. This setup was used as it is standard within the PyTorch community, however it differs from the setup in He et al. (2016a), meaning that test accuracy is close but not directly comparable.

On this problem both Adam and AdaGrad show similar convergence properties as were seen on the CIFAR-10 problem. They both greatly under-perform SGD with momentum. MADGRAD shows strong performance here as well, achieving higher test accuracy than any other method for the majority of the training time, and yielding the best final test accuracy. The accuracy of MADGRAD at epoch 70 is 75.87, a level only reached by SGD+M after

the learning rate reduction at epoch 90, more than 28% longer. MADGRAD also performs the best on training loss on this problem.

7.3 fastMRI challenge MRI reconstruction

The fastMRI Knee challenge (Zbontar et al., 2018) is a recently proposed large-scale image-2-image problem. Unlike the previously explored classification problems, the scale of this problem makes overfitting a non-concern given the number of weights in the largest models currently trainable on contemporary hardware, meaning that adaptive methods are not prone to overfitting. This problem is also particularly notable for being poorly conditioned among image processing problems. Part of the reason for the poor conditioning is the high depth of current SOTA models, such as the VarNet 2.0 Sriram et al. (2020) model that we used. This model has 12,931,532 parameters over 273 layers. Our implementation uses 16 auto-calibration lines and an offset equispaced sampling pattern (Defazio, 2019), which is much closer to a realistic clinical configuration than the challenge’s random sampling mask.

Figure 3 shows a number of interesting properties of the methods. SGD+M exhibits extremely variable performance on this problem, and under-performs other methods by a large margin. AdaGrad also has a clear performance gap compared to the top performing methods, MADGRAD and Adam. MADGRAD is the best performer, with a small but statistically significant improvement over Adam, which is the standard method for this problem. Training set performance shows a much higher degree of variability, making comparisons difficult, however MADGRAD appears to also be the best performing method on training loss as well.

7.4 Machine translation with a recurrent neural network

For a machine translation baseline we trained our model on the IWSLT14 Germain-to-English dataset (Cettolo et al., 2014), using a popular LSTM variant introduced by Wiseman and Rush (2016).

Figure 4 shows that all of the adaptive methods out-perform SGD on this problem by a significant margin. The results are close but MADGRAD has a small performance lead, yielding 4.33 test loss compared to 4.38 for AdaGrad and 4.35 for Adam. In training loss AdaGrad’s lead over the other methods can be attributed to a slight degree of overfitting; these is a slight increase in test loss near the end of optimization for AdaGrad which indicates this.

7.5 Masked language modeling with a Transformer

Bidirectional training objectives, as used in the BERT approach (Devlin et al., 2019), have quickly established themselves as the new standard for large-scale pre-training of natural language models. We performed our experiments using the RoBERTa variant of BERT_BASE (Liu et al., 2019), a 110M parameter transformer model. This model is large enough to provide a realistic optimization test-bed for large-scale Transformer models while still being trainable in in time comparable to a ResNet-50 model on ImageNet.

Similar to the LSTM problem, SGD+M performs poorly here. It exhibits some spikes where training loss rapidly degrades then recovers quickly. both Adam and MADGRAD

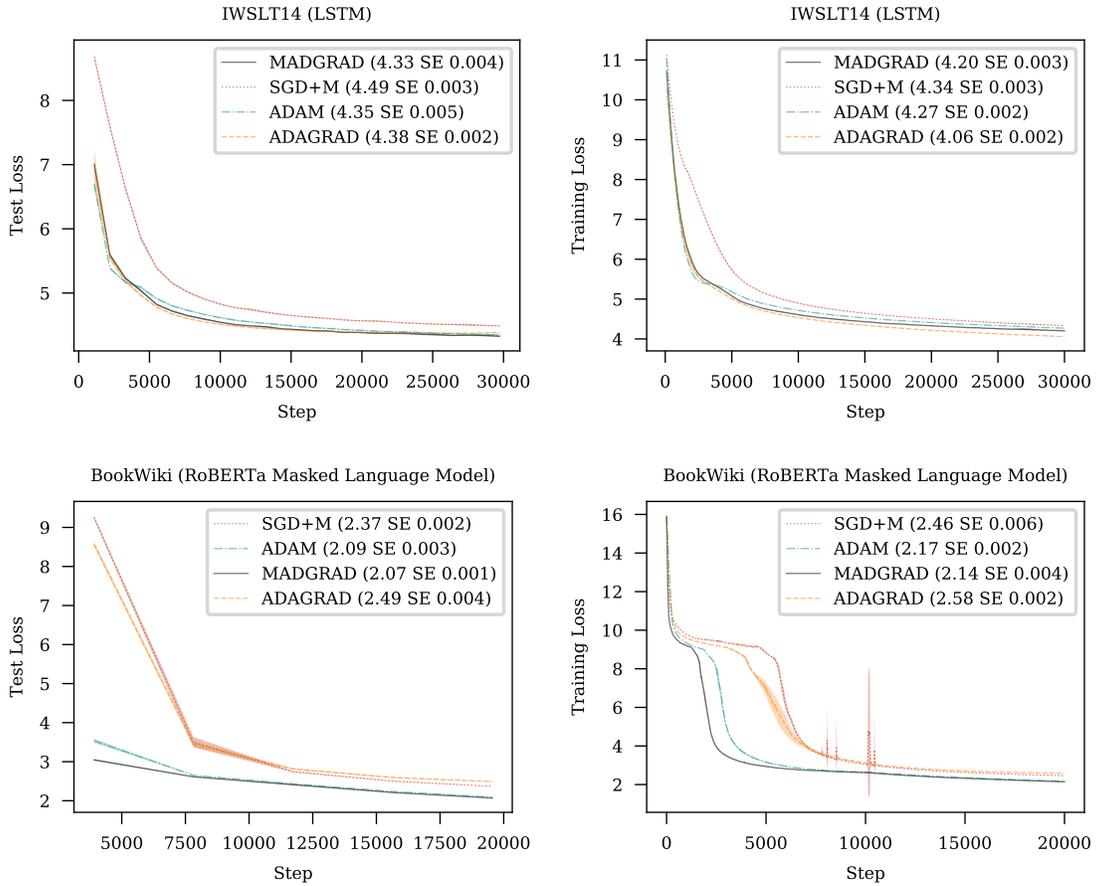


Figure 4: Experimental results for the IWSLT14 and BookWiki problems. Left column shows test set performance and the right column shows training set performance.

perform well, however MADGRAD is significantly faster initially, and also achieves a better final test loss of 2.07 compared to 2.09 achieved by Adam.

8. Discussion

8.1 Hyper-parameter settings

We have made the following observations during our experimentation:

- Typically, using the default weight decay from previous SGD/Adam training runs will result in poor generalization performance. Weight decay will need to be much less, potentially even 0, for good performance. We recommend reducing the weight-decay before any learning rate tuning.
- Learning rate values are not directly comparable to SGD/Adam, a full learning rate sweep is necessary to find the optimal value. In the appendix we list the best LR values for each of our test problems, which should form a good starting point. Sweeping across a power-of-2 grid is recommended as the value several of orders of magnitude different from SGD/Adam.
- Momentum values used for SGD/Adam should work without issue, by setting $c = 1 - \beta$ for momentum β .

8.2 Empirical results in deep learning

We believe our experimental validation is one of the most comprehensive performed for any newly proposed deep learning optimization method. More than 20,000 hours of GPU time were needed to perform the grid search and final evaluation mentioned above, as we performed the search for each of the methods considered, rather than just the MADGRAD method. This prevents our method looking better than it would otherwise look due to hyper-parameter optimization rather than an actual performance advantage. Our comparison also includes a number of large and realistic problems, which are better representative of modern deep learning compared to small scale problems. Finally, our final results are averaged over a sufficiently large number of seeds for each problem to ensure that run-to-run variation is not mistaken for actual performance differences. This is particularly a problem with CIFAR-10, yet many published results still use only a single seed for comparisons on that problem. For these reasons, we believe our experimental results for MADGRAD are representative of the performance of the method across modern large-scale empirical risk minimization problems.

8.3 Sparsity

The reliance on a slowly updating moving average for the squared gradient within the Adam method greatly hinders its application to sparse models. In contrast, MADGRAD maintains a simple sum of the squared gradient entries which may be updated in a sparse fashion. One potential problem in the sparse case is that the buffer of iterates (rather than gradients) is maintained with a moving average. To support sparse applications, the iterate buffer may be removed, effectively equivalent to setting $c = 1$.

9. Conclusion

We have proposed the MADGRAD (Momentumized, Adaptive, Dual averaged GRADient) method as a general purpose optimizer for deep learning. Given MADGRAD’s state-of-the-art empirical performance, together with its strong theoretical foundations, it is an excellent first choice of optimizer across many sub-fields of machine learning.

References

- Lukas Balles and Philipp Hennig. Dissecting adam: The sign, magnitude and variance of stochastic gradients. In *Proceedings of the 35th International Conference on Machine Learning (ICML2018)*, 2018.
- Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, and Animashree Anandkumar. signSGD: Compressed optimisation for non-convex problems. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *Advances in Neural Information Processing Systems 33 pre-proceedings (NeurIPS 2020)*, 2020.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, and Marcello Federico. Report on the 11th IWSLT evaluation campaign, IWSLT 2014. 2014.
- Dami Choi, Christopher J. Shallue, Zachary Nado, Jaehoon Lee, Chris J. Maddison, and George E. Dahl. On empirical comparisons of optimizers for deep learning, 2020.
- Aaron Defazio. Offset sampling improves deep learning based accelerated mri reconstructions by exploiting symmetry. *arXiv preprint arXiv:1912.01101*, 2019.
- Aaron Defazio. Understanding the role of momentum in non-convex optimization: Practical insights from a lyapunov analysis. *arXiv preprint arXiv:2010.00406*, 2020.
- Aaron Defazio and Robert M. Gower. The power of factorial powers: New parameter settings for (stochastic) optimization, 2020.
- Alexandre Défossez, Léon Bottou, Francis Bach, and Nicolas Usunier. A simple convergence proof of adam and adagrad, 2020.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Tout. Bert: Pre-training of deep-bidirectional transformers for language understand. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, 2019.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016a.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Computer Vision – ECCV 2016*, 2016b.
- Samy Jelassi and Aaron Defazio. Dual averaging is surprisingly effective for deep learning optimization, 2020. URL <https://arxiv.org/abs/2010.10502>.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- Kfir Y. Levy, Alp Yurtsever, and Volkan Cevher. Online adaptive methods, universality and acceleration. In *Advances in Neural Information Processing Systems*, 2018.
- Xiaoyu Li and Francesco Orabona. On the convergence of stochastic gradient descent with adaptive stepsizes. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 983–992. PMLR, 16–18 Apr 2019. URL <http://proceedings.mlr.press/v89/li19c.html>.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Yu Nesterov and Vladimir Shikhman. Quasi-monotone subgradient methods for nonsmooth convex minimization. *Journal of Optimization Theory and Applications*, 165(3):917–940, 2015.
- Yurii Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical programming*, 120(1):221–259, 2009.
- Francesco Orabona and Dávid Pál. Scale-free algorithms for online linear optimization. In Kamalika Chaudhuri, CLAUDIO GENTILE, and Sandra Zilles, editors, *Algorithmic Learning Theory*, pages 287–301, Cham, 2015. Springer International Publishing. ISBN 978-3-319-24486-0.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. ArXiv, 2019.
- Sashank J. Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. In *International Conference on Learning Representations*, 2018.

- Robin M. Schmidt, Frank Schneider, and Philipp Hennig. Descending through a crowded valley – benchmarking deep learning optimizers, 2020.
- Othmane Sebbouh, Robert M Gower, and Aaron Defazio. On the convergence of the stochastic heavy ball method. *arXiv preprint arXiv:2006.07867*, 2020.
- Anuroop Sriram, Jure Zbontar, Tullie Murrell, Aaron Defazio, C Lawrence Zitnick, Nafissa Yakubova, Florian Knoll, and Patricia Johnson. End-to-end variational networks for accelerated mri reconstruction. *arXiv preprint arXiv:2004.06688*, 2020.
- Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pages 1139–1147, 2013.
- Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5 - rmsprop, coursera: Neural networks for machine learning. technical report, 2012. URL https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf.
- Rachel Ward, Xiaoxia Wu, and Leon Bottou. AdaGrad stepsizes: Sharp convergence over nonconvex landscapes. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. In *Advances in Neural Information Processing Systems*, 2017.
- Sam Wiseman and Alexander M. Rush. Sequence-to-sequence learning as beam-search optimization. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2016.
- Jure Zbontar, Florian Knoll, Anuroop Sriram, Matthew J Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J Geras, Joe Katsnelson, Hersh Chandarana, et al. fastMRI: An open dataset and benchmarks for accelerated MRI. *arXiv preprint arXiv:1811.08839*, 2018.
- Dongruo Zhou, Jinghui Chen, Yuan Cao, Yiqi Tang, Ziyang Yang, and Quanquan Gu. On the convergence of adaptive gradient methods for nonconvex optimization, 2020.
- Fangyu Zou, Li Shen, Zequn Jie, Ju Sun, and Wei Liu. Weighted adagrad with unified momentum, 2019a.
- Fangyu Zou, Li Shen, Zequn Jie, Weizhong Zhang, and Wei Liu. A sufficient condition for convergences of adam and rmsprop. *CVPR*, 2019b.

Appendix A. Parameter settings

CIFAR10

Our data augmentation pipeline followed standard practice: random horizontal flipping, then random cropping to 32x32, then normalization by centering around (0.5, 0.5, 0.5).

Hyper-parameter	Value
Architecture	PreAct ResNet152
Epochs	300
GPUs	1xV100
Batch Size per GPU	128
LR schedule	150-225 tenthing
Seeds	10

Method	LR	Decay
MADGRAD	2.5e-4	0.0001
AdaGrad	0.01	0.0001
Adam	0.00025	0.0001
SGD	0.1	0.0001

ImageNet

A standard LR schedule was used, where the learning rate is decreased 10 fold every 30 epochs. Interestingly, for this problem, a smaller decay constant improved the performance of MADGRAD, but didn't yield any improvement to the other methods considered.

Hyper-parameter	Value
Architecture	ResNet50
Epochs	90
GPUs	8xV100
Batch size per GPU	32
LR schedule	30-60-90 tenthing
Seeds	5

Method	LR	Decay
MADGRAD	0.001	2.5e-5
AdaGrad	0.01	0.0001
Adam	0.00025	0.0001
SGD	0.1	0.0001

fastMRI

For this task, the best learning rate schedule is a flat schedule, with a small number of fine-tuning epochs at the end to stabilize. To this end, we decreased the learning rate 10 fold at epoch 40.

Hyper-parameter	Value	
Architecture	12 layer VarNet 2	
Epochs	50	
GPUs	8xV100	
Batch size per GPU	1	
Acceleration factor	4	
Low frequency lines	16	
Mask type	Offset-1	
LR schedule	40 tenthing	
Seeds	5	
Method	LR	Decay
MADGRAD	0.01	0.0
AdaGrad	0.25	0.0
Adam	0.00025	0.0
SGD	0.01	0.0

IWSLT14

Our implementation used FairSeq defaults except for the parameters listed below.

Hyper-parameter	Value	
Architecture	lstm_wiseman_iwslt_de_en	
Max updates	60,000	
GPUs	1xV100	
Max tokens per batch	4096	
Warmup steps	4000	
Dropout	0.3	
Label smoothing	0.1	
Share decoder/input/output embed	True	
Float16	True	
Update Frequency	1	
LR schedule	Inverse square-root	
Seeds	10	
Method	LR	Decay
MADGRAD	0.025	5e-6
AdaGrad	0.25	1e-5
Adam	0.01	0.05
SGD	1.0	1e-5

BookWiki

Our implementation used FairSeq defaults except for the parameters listed below.

Hyper-parameter	Value
Architecture	roberta_base
Task	masked_lm
Max updates	20,000
GPUs	8xV100
Max tokens per sample	512
Dropout	0.1
Attention Dropout	0.1
Max sentences	16
Warmup	10,000
Sample Break Mode	Complete
Share decoder/input/output embed	True
Float16	True
Update Frequency	16
LR schedule	Polynomial decay
Seeds	5
Gradient clipping	0.5

Method	LR	Decay
MADGRAD	0.005	0.0
AdaGrad	0.01	0.0
Adam	0.001	0.0
SGD	1.0	0.0

Appendix B. Theory

B.1 Theoretical variant

We analyze a variant of the MADGRAD algorithm, using fixed step size γ , and $\lambda_k = \gamma\sqrt{k+1}$:

$$\begin{aligned}
 s_{k+1} &= s_k + \lambda_k g_k, \\
 v_{k+1} &= v_k + \lambda_k g_k^2, \\
 z_{k+1} &= x_0 - \frac{1}{\sqrt[3]{\lambda_{k+1} G^2 + v_{k+1}}} s_{k+1}, \\
 x_{k+1} &= (1 - c_{k+1}) x_k + c_{k+1} z_{k+1}.
 \end{aligned} \tag{7}$$

This variant differs from Algorithm 1 just with the addition of $\lambda_k G^2$ in the denominator, which is necessitated by our analysis method. Note that the AdaGrad DA formulation originally proposed by Duchi et al. (2011) also requires this extra term.

B.2 Support function

We define a matrix analogue of the support function from Nesterov (2009):

$$V_{A_k}(-s_k) = \max_x \left\{ -\langle s_k, x - x_0 \rangle - \frac{1}{2} \|x - x_0\|_{A_k}^2 \right\}. \tag{8}$$

In this work we only consider diagonal A_k , represented by a vector a_k :

$$A_k = \text{diag}(a_k).$$

In this notation, we have $\alpha_k = \sqrt[3]{\lambda_k G^2 + v_k}$. The maximizer of expression 8 is (using component-wise division):

$$z_k = x_0 - \frac{s_k}{\alpha_k}.$$

Since v_{k+1} is non-decreasing, it's clear that:

$$V_{A_{k+1}}(-s_k) \leq V_{A_k}(-s_k). \quad (9)$$

We will also use the following properties, which follow directly by modifying the argument in Nesterov (2009) to handle scaling matrices instead of constants:

$$\nabla V_{A_k}(-s_k) = z_k - x_0, \quad (10)$$

$$V_{A_k}(s + \delta) \leq V_{A_k}(s) + \langle \delta, \nabla V_{A_k}(s) \rangle + \frac{1}{2} \|\delta\|_{A_k^{-1}}^2. \quad (11)$$

B.3 Lemmas

Lemma 2. *For all natural k , assuming $\lambda_{k+1} \geq \lambda_k$:*

$$\sum_{t=0}^k \frac{\lambda_t^2 g_t^2}{\left(\lambda_t G^2 + \sum_{i=0}^{t-1} \lambda_i g_i^2\right)^{1/3}} \leq \frac{3}{2} \lambda_k \left(\sum_{i=0}^k \lambda_i g_i^2\right)^{2/3}.$$

Proof We prove by induction. For the base case:

$$\frac{g_0^2}{(G^2)^{1/3}} \leq g^{2(1-1/3)} = (g^2)^{2/3} \leq \frac{3}{2} (g^2)^{2/3}.$$

Now assume the lemma holds for $k-1$ then using the inductive hypothesis

$$\begin{aligned} \sum_{t=0}^k \frac{\lambda_t^2 g_t^2}{\left(\lambda_t G^2 + \sum_{i=0}^{t-1} \lambda_i g_i^2\right)^{1/3}} &\leq \frac{\lambda_k^2 g_k^2}{\left(\lambda_k G^2 + \sum_{i=0}^{k-1} \lambda_i g_i^2\right)^{1/3}} + \frac{3}{2} \lambda_{k-1} \left(\sum_{i=0}^{k-1} \lambda_i g_i^2\right)^{2/3}, \\ &\leq \frac{\lambda_k^2 g_k^2}{\left(\lambda_k G^2 + \sum_{i=0}^{k-1} \lambda_i g_i^2\right)^{1/3}} + \frac{3}{2} \lambda_k \left(\sum_{i=0}^{k-1} \lambda_i g_i^2\right)^{2/3}. \end{aligned}$$

Define $b_k = \sum_{i=0}^k \lambda_i g_i^2$ and $a_k = g_k^2$ then we have:

$$\sum_{t=0}^k \frac{\lambda_t^2 g_t^2}{\left(\lambda_t G^2 + \sum_{i=0}^{t-1} \lambda_i g_i^2\right)^{1/3}} \leq \lambda_k^2 a_k \left(\lambda_k G^2 + b_k - \lambda_k a_k\right)^{-1/3} + \frac{3}{2} \lambda_k (b_k - \lambda_k a_k)^{2/3}.$$

We have two terms on the right to consider. For the first term, note that since $a_k \leq G^2$,

$$\lambda_k^2 a_k (\lambda_k G^2 + b_k - \lambda_k a_k)^{-1/3} \leq \lambda_k^2 a_k (b_k)^{-1/3}.$$

For the 2nd term, we can use concavity to get:

$$\frac{3}{2} \lambda_k (b_k - \lambda_k a_k)^{2/3} \leq \frac{3}{2} \lambda_k (b_k)^{2/3} - \lambda_k^2 a_k (b_k)^{-1/3}.$$

Combining gives:

$$\sum_{t=0}^k \frac{\lambda_t^2 g_t^2}{\left(\lambda_t G^2 + \sum_{i=0}^{t-1} \lambda_i g_i^2\right)^{1/3}} \leq \frac{3}{2} \lambda_k (b_k)^{2/3},$$

and so the inductive case is proven. ■

Lemma 3. *Let $0 < r < 1$ and $j \geq 0$. Then define:*

$$c_k = \frac{r+1}{k+j+r},$$

for all $k \geq 0$ it then holds that:

$$\frac{1-c_k}{c_k} (k+j)^r \leq \frac{1}{c_{k-1}} (k+j-1)^r.$$

Proof We start by simplifying:

$$\begin{aligned} \frac{1-c_k}{c_k} (k+j)^r &= \frac{1 - \frac{r+1}{k+j+r}}{\frac{r+1}{k+j+r}} (k+j)^r, \\ &= \frac{k+j-1}{r+1} (k+j)^r, \\ &= \frac{k+j+r-1}{r+1} \frac{k+j-1}{k+j+r-1} (k+j)^r, \\ &= \frac{1}{c_k} \frac{k+j-1}{k+j+r-1} (k+j)^r. \end{aligned}$$

So we need:

$$(k+j)^r \leq \frac{k+j+r-1}{k+j-1} (k+j-1)^r.$$

Recall the concavity upper bound:

$$f(x) \leq f(y) + \langle \nabla f(y), x-y \rangle,$$

using $f(x) = (k+j)^r$ which is concave for $r \in (0, 1)$, and $x = k+j, y = k+j-1$, we have:

$$\begin{aligned} (k+j)^r &\leq (k+j-1)^r + r(k+j-1)^{r-1}, \\ &= (k+j-1)^r + \frac{r}{k+j-1} (k+j-1)^r, \\ &= \frac{k+j-1+r}{k+j-1} (k+j-1)^r. \end{aligned}$$

Which establishes the result. ■

Lemma 4. *The dual averaging iterates obey:*

$$z_k = x_k - \frac{1 - c_k}{c_k} (x_{k-1} - x_k). \quad (12)$$

Proof We rearrange the x update:

$$\begin{aligned} x_{k+1} &= (1 - c_{k+1}) x_k + c_{k+1} z_{k+1}. \\ \therefore x_k &= (1 - c_k) x_{k-1} + c_k z_k, \\ \therefore c_k z_k &= x_k - (1 - c_k) x_{k-1}, \\ \therefore z_k &= \frac{1}{c_k} x_k - \frac{1 - c_k}{c_k} x_{k-1}. \end{aligned}$$

■

Theorem 5. *Consider the MADGRAD method. We upper bound the quantity $V_{A_{k+1}}(-s_{k+1})$ as follows: For the first step $k = 0$:*

$$V_{A_1}(-s_1) \leq \frac{\lambda_0^2}{2} \|\nabla f(x_0, \xi_k)\|_{A_0^{-1}}^2.$$

For subsequent steps $k \geq 1$:

$$\begin{aligned} V_{A_{k+1}}(-s_{k+1}) &\leq V_{A_k}(-s_k) + \frac{\lambda_k^2}{2} \|\nabla f(x_k, \xi_k)\|_{A_k^{-1}}^2 + \lambda_k \langle \nabla f(x_k, \xi_k), x_0 - x_* \rangle \\ &\quad - \frac{1}{c_k} \lambda_k [f(x_k, \xi_k) - f(x_*, \xi_k)] + \frac{1 - c_k}{c_k} \lambda_k [f(x_{k-1}, \xi_k) - f(x_*, \xi_k)]. \end{aligned}$$

Proof Base case:

$$V_{A_1}(-s_1) \quad (13)$$

$$\leq -\lambda_0 \langle \nabla f(x_0, \xi_k), \nabla V_0(-s_0) \rangle + \frac{\lambda_0^2}{2} \|\nabla f(x_0, \xi_k)\|_{A_0^{-1}}^2 \quad (\text{Eq. 11}),$$

$$= \lambda_k \langle \nabla f(x_k, \xi_k), x_0 - x_0 \rangle + \frac{\lambda_0^2}{2\beta_0} \|\nabla f(x_0, \xi_k)\|_{A_0^{-1}}^2, \quad (\text{Eq. 10})$$

$$= \frac{\lambda_0^2}{2} \|\nabla f(x_0, \xi_k)\|_{A_0^{-1}}^2. \quad (14)$$

Inductive case:

$$\begin{aligned}
& V_{A_{k+1}}(-s_{k+1}) \\
& \leq V_{A_k}(-s_{k+1}) \\
& \leq V_{A_k}(-s_k) - \lambda_k \langle \nabla f(x_k, \xi_k), \nabla V_{A_k}(-s_k) \rangle + \frac{\lambda_k^2}{2} \|\nabla f(x_k, \xi_k)\|_{A_k^{-1}}^2, \quad (\text{Eq. 11}) \\
& = V_{A_k}(-s_k) + \lambda_k \langle \nabla f(x_k, \xi_k), x_0 - z_k \rangle + \frac{\lambda_k^2}{2} \|\nabla f(x_k, \xi_k)\|_{A_k^{-1}}^2, \quad (\text{Eq. 10}) \\
& = V_{A_k}(-s_k) + \frac{\lambda_k^2}{2} \|\nabla f(x_k, \xi_k)\|_{A_k^{-1}}^2 \\
& + \lambda_k \left\langle \nabla f(x_k, \xi_k), x_0 - x_k + \left(\frac{1-c_k}{c_k} \right) (x_{k-1} - x_k) \right\rangle, \quad (\text{Eq. 12}) \\
& = V_{A_{k+1}}(-s_k) + \frac{\lambda_k^2}{2} \|\nabla f(x_k, \xi_k)\|_{A_k^{-1}}^2 \\
& + \lambda_k \langle \nabla f(x_k, \xi_k), x_0 - x_k \rangle + \lambda_k \frac{1-c_k}{c_k} \langle \nabla f(x_k, \xi_k), x_{k-1} - x_k \rangle, \\
& = V_{A_{k+1}}(-s_k) + \frac{\lambda_k^2}{2} \|\nabla f(x_k, \xi_k)\|_{A_k^{-1}}^2 \\
& + \lambda_k \langle \nabla f(x_k, \xi_k), x_0 - x_* \rangle + \lambda_k \langle \nabla f(x_k, \xi_k), x_* - x_k \rangle \\
& + \lambda_k \frac{1-c_k}{c_k} \langle \nabla f(x_k, \xi_k), x_{k-1} - x_k \rangle.
\end{aligned}$$

Now we use:

$$\langle \nabla f(x_k, \xi_k), x_* - x_k \rangle \leq f(x_*, \xi_k) - f(x_k, \xi_k),$$

and:

$$\langle \nabla f(x_k, \xi_k), x_{k-1} - x_k \rangle \leq f(x_{k-1}, \xi_k) - f(x_k, \xi_k),$$

to give:

$$\begin{aligned}
V_{A_{k+1}}(-s_{k+1}) & \leq V_{A_k}(-s_k) + \frac{\lambda_k^2}{2} \|\nabla f(x_k, \xi_k)\|_{A_k^{-1}}^2 \\
& + \lambda_k \langle \nabla f(x_k, \xi_k), x_0 - x_* \rangle \\
& + \lambda_k [f(x_*, \xi_k) - f(x_k, \xi_k)] + \lambda_k \frac{1-c_k}{c_k} [f(x_{k-1}, \xi_k) - f(x_k, \xi_k)],
\end{aligned}$$

grouping function value terms gives the result. ■

B.4 Convergence rate

Theorem 6. *After k steps of MADGRAD,*

$$\mathbb{E} [f(x_k) - f(x_*)] \leq \frac{6}{k^{1/2}} \|x_0 - x_*\| GD^{1/2},$$

if $c_k = \frac{3/2}{k+3/2}$ and

$$\gamma = \frac{1}{k^{3/4} D^{3/4} G^{1/2}} \|x_0 - x_*\|^{3/2}.$$

We assume that $\gamma_k = \gamma$ is a constant. First note that for our choice of $\lambda_k = \gamma(k+1)^{1/2}$ and:

$$c_k = \frac{3/2}{k + 3/2},$$

applying Lemma 3 gives that:

$$\frac{1 - c_k}{c_k} \lambda_k \leq \frac{1}{c_{k-1}} \lambda_{k-1}.$$

Using this bound we can telescope the bound from Theorem 5 after taking expectations:

$$\begin{aligned} \frac{1}{c_k} \lambda_k [f(x_k, \xi_k) - f(x_*, \xi_k)] &\leq -\mathbb{E} [V_{A_{k+1}}(-s_{k+1})] + \frac{1}{2} \mathbb{E} \left[\sum_{t=0}^k \lambda_t^2 \|\nabla f(x_t, \xi_t)\|_{A_t^{-1}}^2 \right] \\ &\quad + \mathbb{E} \left\langle \sum_{i=0}^k \lambda_i \nabla f(x_i, \xi_i), x_0 - x_* \right\rangle. \end{aligned}$$

Now note that $s_{k+1} = \sum_{i=0}^k \lambda_i \nabla f(x_i, \xi_i)$, so:

$$\begin{aligned} \mathbb{E} [V_{A_{k+1}}(-s_{k+1})] &= \mathbb{E} \left[\max_x \left\{ \langle -s_{k+1}, x - x_0 \rangle - \frac{1}{2} \|x - x_0\|_{A_{k+1}}^2 \right\} \right], \\ &\geq \mathbb{E} \left[\langle -s_{k+1}, x_* - x_0 \rangle - \frac{1}{2} \|x_* - x_0\|_{A_{k+1}}^2 \right], \\ &= \mathbb{E} \left\langle \sum_{i=0}^k \lambda_i \nabla f(x_i, \xi_i), x_0 - x_* \right\rangle - \frac{1}{2} \|x_* - x_0\|_{A_{k+1}}^2. \end{aligned}$$

So combining this bound and further using the definition of c_k and λ_k :

$$\frac{k + 3/2}{3/2} \gamma (k + 1)^{1/2} \mathbb{E} [f(x_k) - f(x_*)] \leq \frac{1}{2} \mathbb{E} \left[\sum_{t=0}^k \lambda_t^2 \|\nabla f(x_t, \xi_t)\|_{A_t^{-1}}^2 \right] + \frac{1}{2} \|x_* - x_0\|_{A_{k+1}}^2.$$

To simplify further we need to start working in a coordinate wise fashion. Let D be the number of dimensions in x , then we can write the above bound using Lemma 2 applied coordinate wise as:

$$\begin{aligned} \frac{k + 3/2}{3/2} \gamma (k + 1)^{1/2} \mathbb{E} [f(x_k) - f(x_*)] &\leq \frac{1}{2} \sum_{d=0}^D \left(\mathbb{E} \left[\frac{3}{2} \lambda_k \left(\sum_{i=0}^k \lambda_i g_{id}^2 \right)^{2/3} \right] \right) \\ &\quad + \frac{1}{2} \sum_{d=0}^D (x_{0d} - x_{*d})^2 \mathbb{E} \left(\lambda_{k+1} G^2 + \sum_{i=0}^k \lambda_i g_{id}^2 \right)^{1/3}. \end{aligned}$$

We now apply the bound $g_{id} \leq G$:

$$\begin{aligned} \frac{k+3/2}{3/2} \gamma (k+1)^{1/2} \mathbb{E}[f(x_k) - f(x_*)] &\leq \frac{3}{4} \sum_{d=0}^D \left(\lambda_k \left(\sum_{i=0}^k \lambda_i G^2 \right)^{2/3} \right) \\ &\quad + \frac{1}{2} \sum_{d=0}^D (x_{0x} - x_{*d})^2 \left(\sum_{i=0}^{k+1} \lambda_i G^2 \right)^{1/3}. \end{aligned}$$

Since $\lambda_k = \gamma (k+1)^{1/2}$, we can further simplify using the summation property:

$$\sum_{i=0}^k (i+1)^{1/2} \leq \frac{2}{3} (k+2)^{3/2},$$

we apply on the two locations on the right to give:

$$\begin{aligned} \frac{k+3/2}{3/2} \gamma (k+1)^{1/2} \mathbb{E}[f(x_k) - f(x_*)] &\leq \frac{1}{2} \gamma^{5/3} \sum_{d=0}^D (k+1)^{1/2} (k+2) G^{4/3} \\ &\quad + \frac{1}{3} \gamma^{1/3} \sum_{d=0}^D (x_{0x} - x_{*d})^2 (k+3)^{1/2} G^{2/3}. \end{aligned}$$

Note that:

$$\begin{aligned} \frac{(k+3)^{1/2}}{(k+3/2)(k+1)} &\leq \frac{(k+3/2)^{1/2} + (3/2)^{1/2}}{(k+3/2)(k+1)} \\ &\leq \frac{1}{k+1} + \frac{1}{(k+1)} \\ &\leq \frac{2}{k+1} \end{aligned}$$

and likewise:

$$\frac{k+2}{k+3/2} \leq 2$$

so after rearranging:

$$\begin{aligned} \frac{2}{3} \mathbb{E}[f(x_k) - f(x_*)] &\leq 2\gamma^{2/3} G^{4/3} D \\ &\quad + \gamma^{-2/3} G^{2/3} \frac{2}{k+1} \sum_{d=0}^D (x_{0x} - x_{*d})^2, \end{aligned}$$

$$\mathbb{E}[f(x_k) - f(x_*)] \leq 3\gamma^{2/3} G^{4/3} D + \frac{3}{k+1} \gamma^{-2/3} G^{2/3} \|x_0 - x_*\|^2.$$

Taking the gradient with respect to γ to zero gives

$$0 = \frac{2}{3} \gamma^{-1/3} G^{4/3} D - \frac{2}{3(k+1)} \gamma^{-5/3} G^{2/3} \|x_0 - x_*\|^2,$$

$$\begin{aligned}\therefore \gamma^{-1} G^4 D^3 &= \frac{1}{(k+1)^3} \gamma^{-5} G^2 \|x_0 - x_*\|^6, \\ \therefore \gamma^4 &= \frac{1}{(k+1)^3 D^3 G^2} \|x_0 - x_*\|^6, \\ \therefore \gamma &= \frac{1}{(k+1)^{3/4} D^{3/4} G^{1/2}} \|x_0 - x_*\|^{3/2}.\end{aligned}$$

Using this optimal γ gives:

$$\gamma^{2/3} = \frac{1}{k^{1/2} D^{1/2} G^{1/3}} \|x_0 - x_*\|.$$

and so:

$$\mathbb{E}[f(x_k) - f(x_*)] \leq \frac{6}{k^{1/2}} \|x_0 - x_*\| G D^{1/2}.$$

Note that $\|g\|_2 \leq D^{1/2} \|g\|_\infty = D^{1/2} G$, so the dependence on dimensionality here is comparable to standard stochastic method proofs which have $\|g\|_2$ on the right instead.

B.5 Time varying case

Consider the situation where the bound on the gradient potentially varies over time.

$$\|\nabla f(x_i, \xi)\|_\infty \leq G_i \text{ for all } x, \xi.$$

Then using the same argument as in the previous section we arrive at:

$$\begin{aligned}\mathbb{E}[f(x_k) - f(x_*)] &\leq 3\gamma^{2/3} \frac{1}{(k+1)} D \left(\sum_{i=0}^{k+1} (i+1)^{1/2} G_i^2 \right)^{2/3} \\ &\quad + 3\gamma^{-2/3} \frac{1}{(k+1)^{3/2}} \|x_0 - x_*\|_2^2 \left(\sum_{i=0}^{k+1} (i+1)^{1/2} G_i^2 \right)^{1/3}.\end{aligned}$$

We may solve for the optimal step size, giving:

$$\begin{aligned}\gamma^{4/3} &= \frac{1}{(k+1)^{1/2}} \frac{\|x_0 - x_*\|_2^2 \left(\sum_{i=0}^{k+1} (i+1)^{1/2} G_i^2 \right)^{1/3}}{D \left(\sum_{i=0}^{k+1} (i+1)^{1/2} G_i^2 \right)^{2/3}}, \\ \therefore \gamma^{4/3} &= \frac{1}{(k+1)^{1/2}} \frac{\|x_0 - x_*\|_2^2}{D \left(\sum_{i=0}^{k+1} (i+1)^{1/2} G_i^2 \right)^{1/3}}, \\ \therefore \gamma^{2/3} &= \frac{1}{(k+1)^{1/4}} \frac{\|x_0 - x_*\|_2}{D^{1/2} \left(\sum_{i=0}^{k+1} (i+1)^{1/2} G_i^2 \right)^{1/6}}.\end{aligned}$$

Then substituting this in gives:

$$\mathbb{E}[f(x_k) - f(x_*)] \leq 6 \frac{1}{(k+1)^{5/4}} D^{1/2} \|x_0 - x_*\|_2 \left(\sum_{i=0}^{k+1} (i+1)^{1/2} G_i^2 \right)^{1/2}.$$

When applying $\lambda_i = \gamma$, as in AdaGrad, we instead get:

$$\begin{aligned} \mathbb{E}[f(x_k) - f(x_*)] &\leq 3 \frac{\gamma^{1/2}}{(k+1)} D \left(\sum_{i=0}^{k+1} G_i^2 \right)^{1/2} \\ &\quad + 3 \frac{1}{(k+1) \gamma^{1/2}} \|x_0 - x_*\|_2^2 \left(\sum_{i=0}^{k+1} G_i^2 \right)^{1/2}, \end{aligned}$$

solving for the optimal step size:

$$\begin{aligned} \frac{\gamma^{1/2}}{(k+1)} D \left(\sum_{i=0}^k G_i^2 \right)^{1/2} &= \frac{1}{(k+1) \gamma^{-3/2}} \|x_0 - x_*\|_2^2 \left(\sum_{i=0}^k G_i^2 \right)^{1/2}, \\ \therefore \gamma^2 &= \frac{\|x_0 - x_*\|_2^2}{D}. \end{aligned}$$

So:

$$\mathbb{E}[f(x_k) - f(x_*)] \leq \frac{6}{(k+1)} \|x_0 - x_*\|_2 D^{1/2} \left(\sum_{i=0}^k G_i^2 \right)^{1/2}.$$

Appendix C. Cube root formulation

Consider the minimization problem parameterized by $g : k \times D$ and a single vector $s : D$, where

$$\min_s \sum_{i=0}^k \sum_{d=0}^D \frac{g_{id}^2}{s_d}, \quad \|s\|_2^2 \leq c, \quad \forall d : s_d > 0$$

In this section we show that $s_d \propto \sqrt[3]{\sum_{i=0}^k g_{id}^2}$ is a solution. Without loss of generality we disregard the inequality constraint on s_d and consider only positive solutions to the equality constrained problem. We will apply the method of Lagrange multipliers.

Firstly we form the Lagrangian with multiplier μ :

$$L(s, \mu) = \sum_{i=0}^k \sum_{d=0}^D \frac{g_{id}^2}{s_d} + \frac{\mu}{2} \left(\sum_{d=0}^D s_d^2 - c \right)$$

Saddle-points of the Lagrangian can be found by equating the gradients to zero and solving.

$$\frac{\partial L}{\partial s_d} L(s, \mu) = -\frac{1}{s_d^2} \sum_{i=0}^k g_{id}^2 + \mu s_d,$$

$$\frac{\partial L}{\partial \mu} L(s, \mu) = \frac{1}{2} \left(\sum_{d=0}^D s_d^2 - c \right).$$

From the first equation:

$$\begin{aligned} \frac{1}{s_d^2} \sum_{i=0}^k g_{id}^2 &= \mu s_d, \\ \therefore s_d^3 &= \frac{1}{\mu} \sum_{i=0}^k g_{id}^2. \end{aligned}$$

Therefore $s_d = \mu^{-1/3} \left(\sum_{i=0}^k g_{id}^2 \right)^{1/3}$ since s_d is positive we take the positive root, The Lagrange multiplier μ is given by the requirement that

$$\begin{aligned} \sum_{d=0}^D s_d^2 &= c, \\ \therefore \mu^{2/3} &= c^{-1} \sum_{d=0}^D \left(\sum_{i=0}^k g_{id}^2 \right)^{2/3}, \\ \therefore \mu^{1/3} &= \sqrt{c^{-1} \sum_{d=0}^D \left(\sum_{i=0}^k g_{id}^2 \right)^{2/3}}. \end{aligned}$$

So

$$s_d = \frac{1}{\sqrt{c^{-1} \sum_{d=0}^D \left(\sum_{i=0}^k g_{id}^2 \right)^{2/3}}} \left(\sum_{i=0}^k g_{id}^2 \right)^{1/3}.$$

We can verify that this is an extreme point of the original problem by noting that the linear independence constraint qualification (LICQ) condition trivially holds when using one equality constraint. Since the objective is convex for $s_d > 0$, this point must be a minimizer.