
DoG is SGD’s Best Friend: A Parameter-Free Dynamic Step Size Schedule

Maor Ivgi¹ Oliver Hinder² Yair Carmon¹

Abstract

We propose a tuning-free dynamic SGD step size formula, which we call Distance over Gradients (DoG). The DoG step sizes depend on simple empirical quantities (distance from the initial point and norms of gradients) and have no “learning rate” parameter. Theoretically, we show that, for stochastic convex optimization, a slight variation of the DoG formula enjoys strong, high-probability parameter-free convergence guarantees and iterate movement bounds. Empirically, we consider a broad range of vision and language transfer learning tasks, and show that DoG’s performance is close to that of SGD with tuned learning rate. We also propose a per-layer variant of DoG that generally outperforms tuned SGD, approaching the performance of tuned Adam. A PyTorch implementation of our algorithms is available at <https://github.com/formll/dog>.

1. Introduction

While stochastic optimization methods drive continual improvements in machine learning, choosing the optimization parameters—and particularly the learning rate—remains a difficulty. Standard methodologies include searching over a set of learning rates, or simply picking the learning rate from prior work. The former incurs a substantial computational overhead, while the latter risks training a suboptimal model.

The rich literature on adaptive gradient methods (AdaGrad, Adam, and their many variants) offers optimization algorithms that better exploit problem structure (e.g., Duchi et al., 2011; Kingma & Ba, 2015; Gupta et al., 2018; Shazeer & Stern, 2018; Loshchilov & Hutter, 2019). However, these methods still have a learning rate parameter that requires tuning. The theoretically-optimal value of this parameter depends on unknown problem properties. For ex-

¹Tel Aviv University ²University of Pittsburgh. Correspondence to: Maor Ivgi <maor.ivgi@cs.tau.ac.il>.

Proceedings of the 40th International Conference on Machine Learning, Honolulu, Hawaii, USA. PMLR 202, 2023. Copyright 2023 by the author(s).

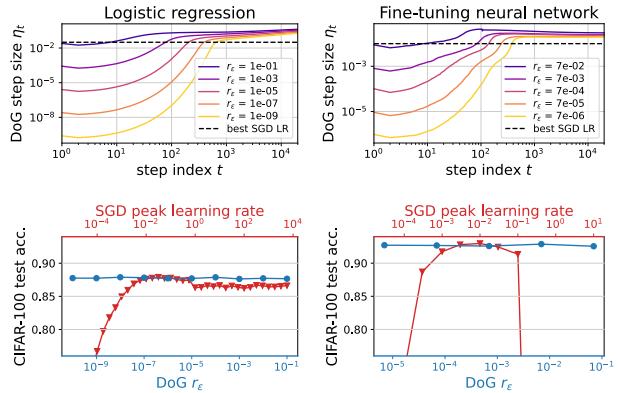


Figure 1. Illustration of DoG for CIFAR-100 classification using logistic regression on last-layer features of a pre-trained ViT-B/32 (left) or end-to-end fine-tuning of the model (right). The top row shows the DoG step size sequence η_t for different values of the initial movement r_ϵ , and the bottom row shows that DoG attains test error on par with carefully tuned SGD (with cosine annealing), even when varying r_ϵ by several orders of magnitude. See details in Appendix E.6.

ample, on convex problems the optimal learning rate of AdaGrad is related to the distance between the initial point and the optimal solution, while in non-convex settings it is related to the function’s smoothness and initial optimality gap (Gupta et al., 2017; Ward et al., 2019; Faw et al., 2022).

Parameter-free optimization aims to remove the need for such tuning by designing algorithms that achieve a near-optimal rate of convergence with almost no knowledge of the problem properties (Streeter & McMahan, 2012). Most works in this field (e.g., Luo & Schapire, 2015; Orabona & Pál, 2016; Cutkosky & Orabona, 2018; Mhammedi & Koolen, 2020; Bhaskara et al., 2020; Jacobsen & Cutkosky, 2022; Zhang et al., 2022) use advanced online learning techniques to construct algorithms that, for the fundamental setting of stochastic convex optimization (SCO) with bounded stochastic gradients, achieve optimal rates of convergence up to logarithmic factors. While practical parameter-free algorithms exist (e.g. Orabona, 2014; Orabona & Tommasi, 2017; Kempka et al., 2019; Chen et al., 2022), there is little research into practical parameter-free step size selection methods for SGD. Recently, Carmon & Hinder (2022) have shown that performing a careful bisection over the SGD step size yields a parameter-free optimization method that is optimal for SCO up to a

double-logarithmic factor. While theoretically novel, on a practical level the result leaves much to be desired, as it essentially prescribes the standard recipe of running SGD multiple times with different learning rates.

Proposed algorithm. In this work, we use key insights from [Carmon & Hinder \(2022\)](#) to go a step further and develop a parameter-free step size schedule. For SGD iterations of the form $x_{t+1} = x_t - \eta_t g_t$, where x_t denotes the model parameters at the t ’th iteration and g_t denotes the stochastic gradient of the loss function, our proposed dynamic steps size is (for all $t \geq 1$)

$$\eta_t = \frac{\max_{i \leq t} \|x_i - x_0\|}{\sqrt{\sum_{i \leq t} \|g_i\|^2}}. \quad (\text{DOG})$$

In words, the step size at iteration t is the maximum distance to between the initial point and observed iterates, divided by the sum of squared stochastic gradient norms, i.e., Distance over Gradients (DOG). At the first step, we set η_0 to be $r_\epsilon / \|g_0\|$, i.e., we take a normalized gradient step of size r_ϵ ; we show that, as long as r_ϵ is small, its precise setting has only mild effect.

Crucially, DOG has no multiplicative “learning rate” parameter: if one considers step sizes of the form $\eta_t = c \cdot \frac{\max_{i \leq t} \|x_i - x_0\|}{\sqrt{\sum_{i \leq t} \|g_i\|^2}}$ then $c = 1$ is a universally good setting (see Section 2 for a heuristic justification and Section 4.3 for empirical evidence for this claim).

Figure 1 highlights key aspects of DOG. The top row shows the DOG step size sequence for different values of r_ϵ in convex (left) and non-convex (right) stochastic optimization problems. The DOG step size increases rapidly (note the logarithmic x scale) and stabilizes around values close to the optimal SGD step size with little dependence on r_ϵ . The bottom row of the figure compares the test errors of DOG and SGD with various step sizes, showing that (for all choices of r_ϵ) DOG is on par with well-tuned SGD.

1.1. Summary of results

Theoretical guarantees. In Section 3 we analyze DOG for stochastic convex optimization with bounded stochastic gradients and a (potentially unbounded) closed convex domain. To present our results, let \mathcal{B} denote a ball around the initial point x_0 with radius $3d_0$, where d_0 is the distance between x_0 and an optimum.

First, we show that if the iterates of DOG remain in \mathcal{B} , then with high probability DOG achieves a convergence rate that is optimal up to a factor of $O(\log(1 + \frac{d_0}{r_\epsilon}))$. In practice, DOG appears to indeed be stable as long as r_ϵ is sufficiently small. However, DOG is not always stable: on pathological functions its iterates can move far from the optimum.

To address this, we consider a theoretical, tamed variant of

DOG, which we call T-DOG, whose step sizes are smaller by a logarithmic factor. We prove that, with high probability, the T-DOG iterates never leave \mathcal{B} . Thus, we obtain a high probability parameter-free convergence guarantee is optimal up logarithmic factors.

To our knowledge, T-DOG is the first parameter-free stochastic optimization method to attain such theoretical guarantee,¹ and only the third high probability parameter-free guarantee in the literature (following [Carmon & Hinder, 2022](#); [Zhang & Cutkosky, 2022](#)).

Empirical study. Our experiments in Section 4 focus on fine-tuning neural networks, because this is a practically important setting that still allows for thorough experiments at a reasonable computational budget. We also perform a small-scale experiment with training a neural network from scratch. Our experiments span 23 natural language understanding and image classification tasks and 8 popular model architectures.

Our results indicate that, compared to DOG, SGD with a cosine step size schedule and tuned base learning rarely attains a relative error improvement of more than 5% (e.g., the difference between accuracy 95% and 95.25%). For convex problems (linear probes), the relative difference in errors is below 1%. In our testbed, well-tuned Adam tends to outperform both SGD and DOG, but a layer-wise version of DOG (which we call L-DOG) closes some of this performance gap.

We also test the sensitivity of DOG to the value of r_ϵ . We find that for most model/task combinations, DOG performs consistently well across a wide range of r_ϵ values as our theory predicts. However, in certain cases, choosing r_ϵ to be too low results in poor performance. We provide some preliminary findings showing that this is due in part to batch normalization.

Put together, our theory and experiments suggest DOG has the potential to save significant computation currently spent on learning rate tuning at little or no cost in performance—especially if we reinvest some of the saved computation in training a larger model on more data.

2. Algorithm Derivation

Before providing rigorous theoretical guarantees for DOG, in this section we explain the origin of the algorithm. Our starting point is the following result by [Carmon & Hinder \(2022\)](#). Suppose we run T iterations of SGD with fixed

¹[Carmon & Hinder \(2022\)](#) guarantee boundedness of the point they output, but do not have guarantees on intermediate query points. [Orabona & Pál \(2021, Lemma 25\)](#) show that the algorithm iterates are bounded but not to within a constant factor of the initial distance to optimality.

step size η , i.e., the recursion $x_{t+1} = x_t - \eta g_t$, where x_t is the SGD iterate and g_t is the stochastic gradient at step t . If, for some $c \in (0, 1)$, it happens to hold that

$$\eta = c \cdot \frac{\max_{k \leq T} \|x_k - x_0\|}{\sqrt{\sum_{k \leq T} \|g_k\|^2}}, \quad (1)$$

then the averaged iterates satisfies an excess loss bound that is at most a factor $\frac{1}{c(1-c^2)}$ larger than the worst-case optimal bound achieved by perfectly tuned SGD.²

The condition (1) is an implicit equation: it allows us to check whether the choice of step size η is good only after running T steps of SGD using that η . Solving this implicit equation therefore requires multiple calls to SGD. We derive the DOG step size sequence by making the equation explicit: we choose η_t so that equation (1) holds at each step. For $c = 1$, this yields the step size formula (DOG). Our reason for choosing $c = 1$ is that it is the threshold under which a solution to the implicit equation yields an optimal rate of convergence. Therefore, in practice we expect 1 to be close to the highest stable value of c , and thus obtain the best performance; we verify this empirically in Section 4.3.

3. Theoretical Analysis

3.1. Preliminaries

Problem setting. Our goal is to minimize a loss function $f : \mathcal{X} \rightarrow \mathbb{R}$ where $\mathcal{X} \subseteq \mathbb{R}^m$ (including the unconstrained setting $\mathcal{X} = \mathbb{R}^m$ as an important special case). We perform our analysis under the following standard convexity assumption.

Assumption 3.1 (Convexity). The function f is convex, its domain \mathcal{X} is closed and convex, and its minimum is attained at some $x_\star \in \mathcal{X}$, i.e., $f_\star := \inf_{x \in \mathcal{X}} f(x) = f(x_\star)$.

In Appendix A we discuss a possible relaxation of convexity under which our results continue to hold.

To minimize f we assume access to a *stochastic gradient oracle* \mathcal{G} . When queried at a point $x \in \mathcal{X}$ the oracle returns a stochastic (sub)gradient estimator $\mathcal{G}(x)$ satisfying $\mathbb{E}[\mathcal{G}(x) \mid x] \in \partial f(x)$. With slight abuse of notation, we write $\nabla f(x) := \mathbb{E}[\mathcal{G}(x) \mid x]$. We make the following assumption, where $\|\cdot\|$ denotes the Euclidean norm.

Assumption 3.2 (Bounded stochastic gradients). There exists some $L > 0$ such that $\|\mathcal{G}(x)\| \leq L$ almost surely.

We can relax Assumption 3.2 to a *local* stochastic gradient norm; see Appendix B.

²This results holds in the non-stochastic case (Carmon & Hinder, 2022, Proposition 1), but a qualitatively similar results holds with high probability in the stochastic case as well (Carmon & Hinder, 2022, Proposition 3).

Algorithm statement. We study (projected) SGD with dynamic learning rate schedule $\{\eta_t\}$, i.e.,

$$x_{t+1} = \text{Proj}_{\mathcal{X}}(x_t - \eta_t g_t)$$

where x_0 is a given initialization, $g_k := \mathcal{G}(x_k)$, and $\text{Proj}_{\mathcal{X}}(\cdot)$ is the Euclidean projection onto \mathcal{X} . To succinctly state and analyze DOG, we define the following quantities:

$$r_t := \|x_t - x_0\|, \bar{r}_t = \max_{k \leq t} r_k \vee r_\epsilon \text{ and } G_t := \sum_{k=0}^t \|g_k\|^2,$$

where $a \vee b := \max\{a, b\}$ and r_ϵ is a small user-specified initial movement size parameter. With this notation, we define a family of DOG-like learning rate schedules.

Definition 3.3. A step size schedule is *DOG-like* if

$$\eta_t = \frac{\bar{r}_t}{\sqrt{G'_t}}$$

for a positive nondecreasing sequence G'_t that depends only on x_0, g_0, \dots, g_t and satisfies $G'_t \geq G_t$.

DOG corresponds to simply setting $G'_t = G_t$; in Section 3.3 we consider a theoretical (or tamed) DOG-like algorithm for which we guarantee bounded iterates by making G'_t larger than G_t by polylogarithmic factors. Throughout, we bound the error of the weighted average sequence

$$\bar{x}_t := \frac{1}{\sum_{k=0}^{t-1} \bar{r}_k} \sum_{k=0}^{t-1} \bar{r}_k x_k. \quad (2)$$

Finally, to streamline the analysis we define:

$$d_t := \|x_t - x_\star\|, \bar{d}_t := \max_{k \leq t} d_k,$$

and

$$\theta_{t,\delta} := \log \left(\frac{60 \log(6t)}{\delta} \right).$$

Logarithm conventions. Throughout the paper \log is base e and $\log_+(\cdot) := 1 + \log(\cdot)$.

3.2. Optimality gap bounds assuming bounded iterates

In this section, we bound the optimality gap attained by any DOG-like algorithm. Our bounds depend on the quantities \bar{r}_T and G_T , and are nearly optimal when $\bar{r}_T = O(d_0)$ (i.e., the DOG iterates don't move too far away from x_0) and G'_T is not much larger than G_T . In the next section we describe a specific DOG-like algorithm that is guaranteed to satisfy both requirements.

Convexity and Jensen's inequality imply that \bar{x}_t satisfies

$$f(\bar{x}_t) - f_\star \leq \frac{1}{\sum_{k=0}^{t-1} \bar{r}_k} \sum_{k=0}^{t-1} \bar{r}_k \langle \nabla f(x_k), x_k - x_\star \rangle. \quad (3)$$

The sum in the RHS decomposes to two components:

$$\underbrace{\sum_{k=0}^{t-1} \bar{r}_k \langle g_k, x_k - x_\star \rangle}_{\text{weighted regret}} - \underbrace{\sum_{k=0}^{t-1} \bar{r}_k \langle \Delta_k, x_k - x_\star \rangle}, \quad (4)$$

where $\Delta_k := g_k - \nabla f(x_k)$. We give probability 1 bounds for the weighted regret (Lemma 3.4) and high probability bounds for the noise term (Lemma 3.5). In each case, the key challenge is replacing a-priori bounds on d_0 (or the domain size) with the empirically observed \bar{r}_T . We present and discuss each lemma in turn.

Lemma 3.4 (Weighted regret bound). *If \mathcal{X} is a closed convex set then any DOG-like scheme (Definition 3.3) satisfies $\sum_{k=0}^{t-1} \bar{r}_k \langle g_k, x_k - x_\star \rangle \leq \bar{r}_t (2\bar{d}_t + \bar{r}_t) \sqrt{G'_{t-1}}, \forall t \geq 1$.*

The proof of Lemma 3.4 appears in Appendix D.1. While it is similar to the analysis of adaptive SGD (where $\eta_t = \frac{\rho}{\sqrt{G_t}}$ (Gupta et al., 2017)), there are a couple of key differences. First, the DOG step sizes can increase, which typically makes adaptive gradient methods difficult to analyze (Reddi et al., 2018). We bypass this difficulty by considering regret weighted by \bar{r}_k , which factors out the increasing portion of the step size. Second, the standard adaptive SGD analysis yields a bound proportional to \bar{d}_t^2 (typically further bounded using the domain diameter) rather than $\bar{r}_t \bar{d}_t$ as in our bound. This is a crucial difference, since—as we soon argue— \bar{r}_t “cancels” when dividing through by $\sum_{k < t} \bar{r}_k$, while \bar{d}_t does not. We obtain the improved result by keeping around the last term in a telescoping sum, a trick similar to Carmon & Hinder (2022, Lemma 1).

Next, we handle the noise term in (4), recalling the notation $\Delta_t := g_t - \nabla f(x_t)$ and $\theta_{t,\delta} := \log \frac{60 \log(6t)}{\delta}$.

Lemma 3.5 (Noise bound). *Under Assumption 3.2, for all $\delta \in (0, 1)$, $T \in \mathbb{N}$ and $L > 0$ we have*

$$\mathbb{P} \left(\exists t \leq T : \left| \sum_{k=0}^{t-1} \bar{r}_k \langle \Delta_k, x_k - x_\star \rangle \right| \geq b_t \right) \leq \delta$$

where $b_t = 8\bar{r}_{t-1}\bar{d}_{t-1}\sqrt{\theta_{t,\delta}G_{t-1} + \theta_{t,\delta}^2L^2}$.

The proof of Lemma 3.5 appears in Appendix D.2 and is based on a new concentration bound, Lemma D.2, which allows us to bound the noise term despite having no deterministic bound on the magnitude of the martingale difference sequence $\bar{r}_k \langle \Delta_k, x_k - x_\star \rangle$. The proof of Lemma D.2 involves combining time-uniform Bernstein bounds (Howard et al., 2021) and a general bound on the cumulative sums of sequence products (Lemma C.2), which may be of independent interest.

Combining the above results, we obtain the following.

Proposition 3.6. *For all $\delta \in (0, 1)$ and $L > 0$, if Assumption 3.1, Assumption 3.2, and Definition 3.3 hold then with probability at least $1 - \delta$, for every $t \leq T$ the optimality gap $f(\bar{x}_t) - f_\star$ is*

$$O \left(\frac{(d_0 + \bar{r}_t) \sqrt{G'_{t-1} + G_{t-1}\theta_{t,\delta} + L^2\theta_{t,\delta}^2}}{\sum_{i < t} \bar{r}_i / \bar{r}_t} \right).$$

Proof. Follows from Equations (3) and (4), Lemma 3.4, Lemma 3.5 and the fact that $\bar{d}_t \leq d_0 + \bar{r}_t$. \square

The following algebraic fact shows that there is always an iteration $\tau \leq T$ where the denominator $\sum_{i < \tau} \frac{\bar{r}_i}{\bar{r}_t} \geq \Omega(T / \log \frac{\bar{r}_T}{r_\epsilon})$; see Appendix C.3 for proof.

Lemma 3.7. *Let s_0, s_1, \dots, s_T be a positive nondecreasing sequence. Then*

$$\max_{t \leq T} \sum_{i < t} \frac{s_i}{s_t} \geq \frac{1}{e} \left(\frac{T}{\log_+ (s_T / s_0)} - 1 \right).$$

Combining Proposition 3.6 and Lemma 3.7 yields the following (see short proof in Appendix C.3).

Corollary 3.8. *Under the setting of Proposition 3.6, let $\tau \in \arg \max_{t \leq T} \sum_{i < \tau} \frac{\bar{r}_i}{\bar{r}_t}$. Then, with probability at least $1 - \delta$, the optimality gap $f(\bar{x}_\tau) - f_\star$ is*

$$O \left(\log_+ \left(\frac{\bar{r}_\tau}{r_\epsilon} \right) \frac{(d_0 + \bar{r}_\tau) \sqrt{G'_{\tau-1} + G_{\tau-1}\theta_{\tau,\delta} + L^2\theta_{\tau,\delta}^2}}{T} \right).$$

Corollary 3.8 is immediately useful when \mathcal{X} is bounded but its exact diameter is unknown, for example when \mathcal{X} is a polytope as is common in two-stage stochastic programming (Nemirovski et al., 2009).

Simplifying the bound for typical DOG trajectories. Suppose that $\bar{r}_T = O(d_0)$ and note that Assumption 3.2 DOG satisfies $G'_t = G_t \leq L^2T$. Substituting into Corollary 3.8 yields an optimality gap bound of $O\left(\frac{d_0 L}{\sqrt{T}} \theta_{T,\delta} \log \frac{\bar{r}_T}{r_\epsilon}\right)$, which is minimax optimal up to a term double-logarithmic in T and logarithmic in $\frac{1}{r_\epsilon}$ (Agarwal et al., 2012).

Furthermore, in realistic DOG trajectories, even the multiplicative term $\log \frac{\bar{r}_T}{r_\epsilon}$ is likely too pessimistic. This is because \bar{r}_t typically increases rapidly for $t_0 < 1000$ steps and then plateaus (see Figure 12 in the appendix). Consequently, $\bar{r}_i / \bar{r}_t \geq \frac{1}{10}$ for most of the optimization trajectory, and $\sum_{i < t} \frac{\bar{r}_i}{\bar{r}_t} \geq \frac{t}{10} - t_0$. Substituting back into Proposition 3.9, we get that \bar{x}_T is $O\left(\frac{d_0 L}{\sqrt{T-t_0}} \theta_{T,\delta}\right)$ suboptimal.

DOG can run wild. While DOG is empirically stable, there exist (non-stochastic) examples where \bar{r}_t grows much larger than d_0 : in Appendix D.4 we describe a variant of Nemirovski's function (Nemirovski & Yudin, 1983; Nemirovski, 1994) for which $\bar{r}_t = r_\epsilon \sqrt{t}$ and therefore \bar{r}_t/d_0 diverges as t grows. Next, we show that by slightly decreasing the DOG step sizes we can guarantee that $\bar{r}_T/d_0 \leq 3$ with high probability.

3.3. Iterate stability bound

This section introduces a new DOG-like step size scheme whose iterates are guaranteed to remain bounded with high probability. We call this scheme T-DOG, where the T stands for “theoretical” or “tamed.” The step sizes depend on the iteration budget T , the failure probability δ , and an upper bound L on the stochastic gradient norms (defined in Assumption 3.2), and are given by $\eta_t = \bar{r}_t/\sqrt{G'_t}$, where

$$G'_t = 8^4 \theta_{T,\delta} \log_+^2(t+1)(G_{t-1} + 16\theta_{T,\delta} L^2), \quad (\text{T-DOG})$$

using $G_{-1} := 0$. The dependencies on T , δ and L are weak, since $\theta_{t,\delta} := \log\left(\frac{\log(6t)}{\delta}\right)$ and the L -dependent term (that barely grows with t) will typically be smaller than the term proportional to G_{t-1} .

With the definition of T-DOG in hand, we are ready to state its key property: guaranteed iterate stability.

Proposition 3.9. *Suppose that Assumptions 3.1 and 3.2 hold and $r_\epsilon \leq 3d_0$. For any $\delta \in (0, 1)$, and $T \in \mathbb{N}$, the iterations of T-DOG satisfy $\mathbb{P}(\bar{r}_T > 3d_0) \leq \delta$.*

We defer the full proof to Appendix D.5 and proceed to highlight the key argument by proving the result in the noiseless case.

Proof of Proposition 3.9 in the noiseless case. In this case we have $g_k = \nabla f(x_k)$ and therefore $\langle g_k, x_k - x_* \rangle \geq f(x_k) - f_*$. Substituting into (5) and rearranging gives $d_{k+1}^2 - d_k^2 \leq \eta_k^2 \|g_k\|^2$. Assuming by induction that $\bar{r}_t \leq 3d_0$ and telescoping yields

$$\begin{aligned} d_{t+1}^2 - d_0^2 &\leq \sum_{k=0}^t \frac{\bar{r}_k^2 \|g_k\|^2}{G'_k} \stackrel{(i)}{\leq} \frac{\bar{r}_t^2}{8^4} \sum_{k=0}^t \frac{G_k - G_{k-1}}{(G_k + L^2) \log_+^2 \frac{G_k + L^2}{L^2}} \\ &\stackrel{(ii)}{\leq} \frac{\bar{r}_t^2}{8^4} \stackrel{(iii)}{\leq} \frac{9d_0^2}{8^4} \implies d_{t+1} \leq 2d_0, \end{aligned}$$

where (i) uses that $\|g_k\|^2 = G_k - G_{k-1}$ (with the shorthand $G_{-1} := 0$) and

$$G'_k \geq 8^4(G_k + L^2) \log_+^2 \frac{G_k + L^2}{L^2}$$

since $G_k \leq kL^2$ for all $k \leq t$, (ii) uses Lemma C.3 with $a_k = G_k + L^2$, and (iii) uses $\bar{r}_t \leq 3d_0$ again. Therefore, $r_{t+1} \leq d_{t+1} + d_0 \leq 3d_0$ by the triangle inequality, completing the induction step. \square

Finally, we state the main guarantee for T-DOG.

Theorem 3.10. *Suppose that Assumptions 3.1 and 3.2 hold. For any $\delta \in (0, \frac{1}{2})$, $T \in \mathbb{N}$, consider T iterations of T-DOG with $r_\epsilon \leq 3d_0$. Then for $\tau \in \arg \max_{t \leq T} \sum_{i < \tau} \bar{r}_i/\bar{r}_t$ we have, with probability at least $1 - 2\delta$, that $f(\bar{x}_\tau) - f_*$ is upper bounded by*

$$O\left(c_{\delta,r_\epsilon,T} \frac{d_0 \sqrt{G_\tau + L^2}}{T}\right) \leq O\left(c_{\delta,r_\epsilon,T} \frac{d_0 L}{\sqrt{T}}\right),$$

$$\text{where } c_{\delta,r_\epsilon,T} = \log_+(T) \log_+\left(\frac{d_0}{r_\epsilon}\right) \log\left(\frac{\log_+(T)}{\delta}\right).$$

Proof. Follows by Corollary 3.8, Proposition 3.9 and (T-DOG). \square

Theorem 3.10 yields the optimal convergence bound (Agarwal et al., 2012) up to logarithmic factors. Moreover, to the best of our knowledge our method is the first to produce an iterate bound of the form $\bar{r}_T = O(d_0)$. In Appendix D.6 we discuss how this iterate bound can lead to better convergence rates in least squares problems.

4. Experiments

To test DOG in practical scenarios, we perform extensive experiments over a diverse set of tasks and model architectures in both the vision and language domains. We construct a testbed that consists of over 20 tasks and 7 model architecture, covering natural language understanding and computer vision (Section 4.1). In this testbed we compare DOG to SGD and Adam (Section 4.2), showing that DOG performs on par with tuned SGD, but not as well as tuned Adam. Nevertheless, a per-layer version of DOG (defined below) closes much of this gap with Adam without requiring tuning. We also use our testbed to analyze the sensitivity of DOG to its fixed parameters (Section 4.3), demonstrate its effectiveness in convex logistic regression settings (Section 4.4). Finally, we apply DOG and L-DOG to fine-tuning a CLIP model on ImageNet (Section 4.5) and training a CIFAR10 model from scratch (Section 4.6), and provide preliminary comparison to previously-proposed tuning free methods (Section 4.7). A PyTorch implementation of DOG is available at <https://github.com/formll/dog>.

Layer-wise DOG. Neural models in general and transformer-based models in particular often benefit from using a per-parameter or per-layer step sizes (Kingma & Ba, 2015; You et al., 2020). With this in mind, we consider

a per-layer version of DOG, which we call L-DOG, where we apply the (DOG) formula separately for every layer. Namely, if we consider x_t^l to be the weights in layer³ l at step t , then we set the learning rate for that layer to be $\eta_t^l = \frac{\max_{i \leq t} \|x_i^l - x_0^l\|}{\sqrt{\sum_{i \leq t} \|g_i^l\|^2 + \epsilon}}$, where $\epsilon = 10^{-8}$ is added to the denominator for numerical stability. While we do not provide theoretical guarantees for L-DOG, we show below that it performs well in practice.

4.1. Fine-tuning testbed

Our main experiments focus on fine-tuning pre-trained models, which allows us to experiment with advanced models while also thoroughly tuning the learning rate for the baseline optimizers, using an academic computational budget.

Common hyperparameters. For each baseline algorithm, we use best-practice learning rate schedule (cosine annealing for all experiments, with a warm-up stage for language experiments) and sweep over the peak learning rate for each model/task pair. We give each pair a fixed step budget designed to suffice for convergence, performing evaluation throughout the training. In all cases, we use polynomial decay averaging⁴ as proposed by [Shamir & Zhang \(2013\)](#), and select the best checkpoint (either averaged or not) based on evaluation performance. We repeat relevant learning setups with 5 different seeds, and report the mean performance across the seeds. For simplicity, we do not use weight decay throughout. The complete set of hyper-parameters appears in Appendix E.

Natural language understanding (NLU). To test DOG’s efficacy in modern NLU, we use it to fine-tune transformer language models ([Vaswani et al., 2017](#)) on the well-studied GLUE benchmark ([Wang et al., 2019b](#)) which measures models’ performance on diverse text classification tasks (listed in Appendix E.3).

Additionally, we fine-tune models on SQuAD 1.1, a question answering dataset ([Rajpurkar et al., 2016](#)). We fine-tune a RoBERTa-base ([Liu et al., 2019](#)) checkpoint and T5-base ([Raffel et al., 2020](#)).⁵ For each task, we use the official evaluation metrics defined in by [Wang et al. \(2019b\)](#) and [Rajpurkar et al. \(2016\)](#) as well as their original proposed splits, and report the results over the evaluation set.

Computer vision. We also fine-tune 5 models architectures on 12 different computer vision tasks from the VTAB

³More precisely, our implementation treats each element in the PyTorch `.parameters()` list as a separate layer.

⁴We apply the weight averaging with a fixed parameter ($\gamma = 8$, following ([Levy et al., 2020](#))); we did not try any other parameter in our experiments.

⁵Throughout the paper we often use the shorthand names RoBERTa-b and T5-b, respectively.

benchmark ([Zhai et al., 2019](#)) (see Appendix E.3); of the other 7 tasks in VTAB, 5 are trivial (accuracy greater than 99%) and 2 have small validation splits leading to unreliable model selection. We follow the training, validation and test splits defined in VTAB, and report performance on the test split (using the validation split for model selection). We fine-tune 5 models: VGG11 ([Simonyan & Zisserman, 2014](#)), ResNet50 ([He et al., 2016](#)), DenseNet121 ([Huang et al., 2017](#)), ViT-B/32 ([Dosovitskiy et al., 2021](#)), and ConvNeXt-T ([Liu et al., 2022](#)), where the ViT model is pre-trained on ImageNet 21K and the others are trained on ImageNet 1K ([Deng et al., 2009](#)).

Normalized performance metric. Since the performance metrics in our testbed vary substantially across tasks and models, they are challenging to compare in aggregate. To address this, we consider the following notion of *relative error difference* (RED), that provides a normalized measure of a difference between two model’s performance. In particular, given a task and a model architecture, let err_x be the error⁶ of the model when trained with optimizer x (Adam or SGD with a certain learning rate, or L-DOG) and let err_{DOG} be the error when trained with DOG. Then

$$\text{RED}(\text{err}_x, \text{err}_{\text{DOG}}) := \frac{\text{err}_{\text{DOG}} - \text{err}_x}{\text{err}_{\text{DOG}}}.$$

A positive RED value indicates that optimizer x is better than DOG, and a negative value indicates the opposite. When the absolute value of RED is beneath a few percentage points, the compared methods are nearly equivalent.

Setting r_ϵ . Our theoretical analysis suggest that the particular choice of r_ϵ does not matter as long as it is sufficiently small relative to the distance between the weight initialization x_0 to the optimum. Consequently, for vision experiments we set $r_\epsilon = \alpha \cdot (1 + \|x_0\|)$ for $\alpha = 10^{-4}$, assuming that the distance to the optimum is more than 0.01% of the initialization norm. For language experiments, this assumption turned out to be wrong (causing DOG to diverge in some cases), and we decreased α to 10^{-6} for DOG and to 10^{-8} for L-DOG, where the additive 10^{-6} term was too large in some layers. We believe that 10^{-6} and 10^{-8} should be good defaults for DOG and L-DOG, respectively, though networks with batch normalization or different initialization schemes could require a larger value; see Section 4.3 for additional discussion.

4.2. Comparison of fine-tuning performance

Figure 2 depicts the median, IQR (inter-quartile range) and mean RED of each model,⁷ when trained with SGD and

⁶We consider the error to be 1 minus the respective performance metric, as detailed in Table 2.

⁷When aggregating results over tasks, we always report the RED statistics across tasks, where for each task we average the RED values over seeds. See Appendix E.5 for details.

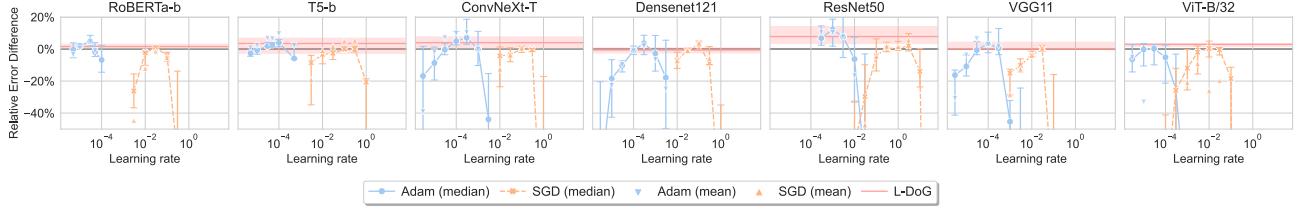


Figure 2. Relative error difference statistics (median, mean, and error bars showing IQR) across tasks for each model, as a function of peak learning rate. The red horizontal line and shaded region indicate the median and IQR RED for L-DoG, respectively.

Adam with different peak learning rates. The figure shows that, when comparing across models, there is no good default learning rate for neither SGD nor Adam. Moreover, even for a single model only very specific SGD learning rate performs well, while most are considerably inferior to using DOG. Even when tuned to the best *fixed* learning-rate value per model (which we refer to as *model tuned LR*), some tasks may still fail (compared to DOG) as indicated by the large IQR and the gap between the mean (triangles) and the median RED (circles) in models such as ViT-B/32 and Densenet121. While Adam also requires tuning, it is somewhat less sensitive than SGD to the choice of peak learning rate. For a full breakdown of performance per task, see Figure 5 and Tables 4 and 5 in Appendix F.1.

DOG performs similarly to well-tuned SGD in 79 out of the 80 model/task combinations in our testbed. The one exception is tuning T5-b on CoLA, where DOG behaves erratically while SGD succeeds only with a few learning rates. In contrast, both Adam and L-DOG achieved reasonable performance consistently. DOG’s poor performance on CoLA results in high RED measures for this case, which draw the mean RED (triangles) above the median one in Figure 2 for T5-b. We further analyze this exception in Appendix F.3 and show that choosing significantly smaller r_e for DOG alleviates the problem.

Figure 3 (top) compares DOG to SGD with model tuned LR as defined above, as well as *instance tuned LR*, where for each model/task pair we select the best learning rate, at a computational expense 6–7 times larger than running DOG. The performance of DOG remains close to that of SGD with instance-tuned LR, with the largest median RED observed for ResNet50 and ViT-B/32.

Figure 3 (bottom) compares DOG to model-tuned and instance-tuned Adam, as well as to L-DOG. In a few cases (namely ResNet50 and ConvNeXt-T) the gaps between DOG and Adam are significant, and favor Adam. We hypothesize this is due to Adam’s per-parameter step-sizes and momentum mechanisms, which DOG does not exploit. L-DOG, which has per-layer steps, has positive median RED for all models, and narrows the gap between DOG and Adam, particularly for ResNet50.

The instance-tuned baselines consume significantly more

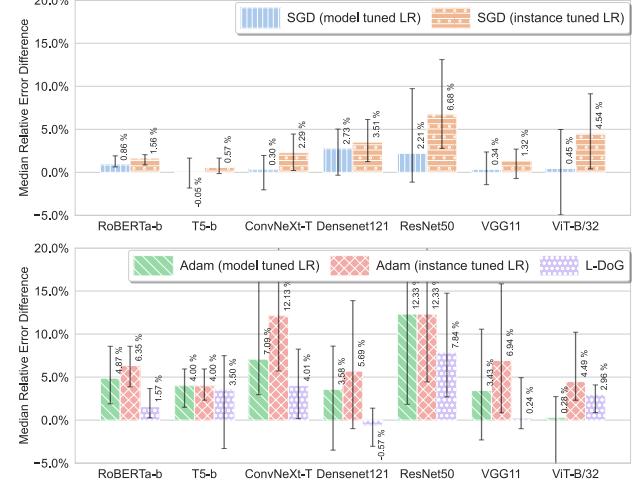


Figure 3. RED median (bar chart) and IQR (error bars) of each model on the set of applicable tasks. **Top:** Comparison with SGD when the LR is optimally tuned per model (*model tuned LR*) or per task (*instance tuned LR*). DOG is competitive with model-tuned SGD and often performs nearly as well as instance-tuned SGD. **Bottom:** Comparison of DOG with adaptive optimizers. L-DOG closes most of the gap to Adam.

compute than DOG and L-DOG. In Appendix F.2 we equalize the compute budget by reducing the number of steps for SGD and Adam. This makes DOG outperform instance-tune SGD in most cases, and brings L-DOG substantially closer to Adam.

4.3. Sensitivity of DOG’s fixed parameters

Initial movement size r_e . Our theory suggests that all sufficiently small choices of r_e should perform similarly, but choosing r_e too large (compared to the initial distance to the optimum) can hurt the performance of the algorithm. In Figure 4 (left) we plot the test performance as a function of r_e for 8 model/task combinations. For 7 out of the 8, DOG is highly robust to the value of r_e as long as it small enough, as predicted. However, ResNet50 on CIFAR-100 (bottom left) is an exception, where smaller values of r_e result in an accuracy drop. We hypothesize this is due to scale invariance introduced by batch normalization (BN), and provide supporting evidence for that in Appendix F.4 (Figure 8), where we show that DOG is insensitive to r_e

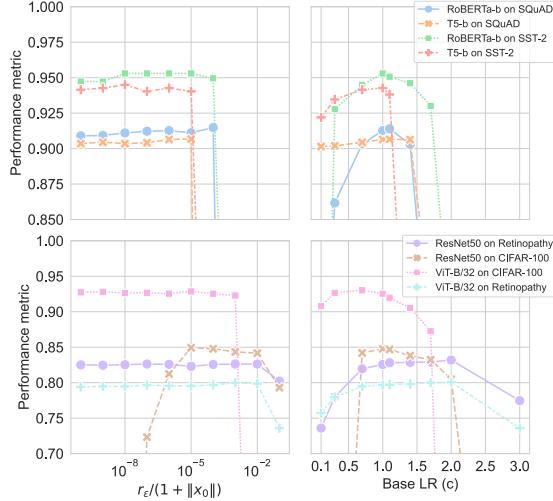


Figure 4. Performance metrics of models trained with DOG as a function of η_0 (left) or the base learning rate (right).

when we turn off BN. In the appendix we also provide a complementary diagnostic for r_ϵ sensitivity by plotting η_t vs. η_0 for different values of t (see Figure 6).

Base learning rate. For this experiment only, we consider variants of DOG with different values of base learning, i.e., step sizes of the form $\eta_t = c \cdot \frac{\max_{i \leq t} \|x_i - x_0\|}{\sqrt{\sum_{i \leq t} \|g_i\|^2}}$ with different values of c . We expect optimal performance when c is close to 1. More specifically, we expect the algorithm to be unstable when $c > 1$ and to be slower to converge (and less likely to generalize well) when $c < 1$. As can be observed in Figure 4 (right), values around $c = 1$ perform well for all models. For smaller values, there is indeed inferior performance in some models (mainly ResNet50 and RoBERTa-b)—indicating T-DOG would not work well in practice—while larger values result in divergence (in 6 out of 8 cases). Hence, the useful range for c is very narrow (about $[0.5, 1.5]$) and tuning it is not likely to produce significant improvements. This is in contrast to Adam and SGD which generally require searching over a space spanning a few orders of magnitude to properly train a model.

4.4. Convex optimization

We also evaluate DOG on convex optimization tasks, matching the assumptions of our theoretical analysis. To do so, we perform multi-class logistic regression on features obtained from the computer vision models in our testbed, i.e., linear probes. We find that model-tuned SGD performs on par or worse than DOG, while instance-tuned SGD barely gains any advantage (Figure 9), with RED values well under 1% (corresponding to the difference between accuracies 90% and 90.1%). Moreover, even in this simple setting, SGD is sensitive to the choice of learning rate, which differ significantly between models (Figure 10).

Table 1. ImageNet top-1 validation accuracies after fine-tuning a CLIP ViT-B/32 model for 25K training steps, with and without polynomial decay averaging (see Section 4.5).

Algorithm	LR	Acc. w/o avg.	Acc. w/ avg.
SGD	1e-03	60.70%	60.49%
	3e-03	73.62%	73.54%
	1e-02	76.82%	76.80%
	3e-02	77.51%	77.54%
	1e-01	75.73%	75.71%
DoG	-	74.78%	77.22%
AdamW	1e-05	78.23%	78.25%
	3e-05	79.04%	79.01%
	1e-04	75.02%	74.97%
L-DOG	-	78.20%	80.12%

4.5. Fine-tuning on ImageNet

To complement our main fine-tuning testbed, we perform a more limited experiment involving ImageNet as a downstream task, which is more expensive to tune due its larger scale. We fine-tune a ViT-B/32 CLIP model (Radford et al., 2021) and compare DOG and L-DOG to training with SGD as well as to an AdamW (Loshchilov & Hutter, 2019) training prescription similar to Wortsman et al. (2022); see Appendix E.7 for additional details. Table 1 shows the ImageNet top-1 validation accuracies of the final model checkpoints, with and without the polynomial decay averaging used throughout our experiments. DOG performs similarly to SGD, but both algorithm perform significantly worse than AdamW, perhaps due to an insufficient iteration budget. L-DOG performs well in this setting, improving on AdamW by a little over 1 point.

4.6. Training from scratch

We conduct a preliminary experiment with training a model from scratch, specifically a Wide ResNet 28-10 (Zagoruyko & Komodakis, 2016) on CIFAR-10 (Krizhevsky, 2009); see Appendix E.8 for details. Table 3 shows the test accuracy of the final checkpoint, with and without the polynomial averaging used throughout our experiments. Here DOG performs on par with the setting’s canonical training prescription of SGD with momentum 0.9 and learning rate 0.1 (Cubuk et al., 2019). In this setting Adam produces poorer results, and L-DOG is 0.5 point worse than tuned Adam with the best learning rate, perhaps due to not reaching convergence.

4.7. Comparison to other tuning-free methods

We perform preliminary comparisons between DOG and L-DOG and other methods for removing the learning rate parameter: the Stochastic Polyak Step (Loizou et al., 2021), and D-Adaptation (Defazio & Mishchenko, 2022). In both

cases, we find that DOG and L-DOG provide better performance on most tasks and on average (see Tables 6 and 7). We provide detailed results in Appendix G, where we also discuss the practical prospects of the bisection procedure of Carmon & Hinder (2022).

5. Related Work

Previous attempts to design theoretically principled and practical optimization algorithms that do not require learning rate tuning approach the problem from a variety of perspectives, resulting in a large variety of proposed algorithms. Rolinek & Martius (2018); Vaswani et al. (2019); Paquette & Scheinberg (2020) lift classical line search technique from non-stochastic optimization to the stochastic setting, while Berrada et al. (2020); Loizou et al. (2021) do the same for the classical Polyak step size (Polyak, 1987; Hazan & Kakade, 2019). Asi & Duchi (2019) develop a class of algorithms based on stochastic proximal methods and demonstrate their improved robustness both theoretically and empirically. Schaul et al. (2013) use a stochastic quadratic approximation for designing learning rates that maximize the expected one-step objective decrease. Chandra et al. (2022) nest hypergradient descent to make a method that is insensitive to initial hyper-parameter choices. However, none of these results are parameter-free in the same sense as DOG: they either do not have converges guarantees, or have suboptimality bounds that blow up polynomially when the method’s parameters do not match a problem-dependent value. In contrast, parameter-free methods have converges rates that depend at most logarithmically on algorithmic parameters.

While the parameter-free optimization literature has focused mainly on theoretical schemes, a number of works also include empirical studies (Orabona, 2014; Orabona & Tommasi, 2017; Kempka et al., 2019; Chen et al., 2022). In particular, Orabona & Tommasi (2017) build on coin-betting schemes to design an algorithm for training neural networks that has AdaGrad-style convergence guarantees for quasi-convex functions, showing promising results on neural network training problems. In recent work Chen et al. (2022) obtain improved empirical results with an algorithm that leverages coin betting and truncated linear models. However, this method lacks theoretical guarantees.

In recent independent work Defazio & Mishchenko (2022) propose a parameter-free dynamic step size schedule of dual averaging. While our work has the same motivation and shares a number of technical similarities (including the use of weighted regret bounds and an independently obtained Lemma 3.7) the proposed algorithms are quite different, and dual averaging is rarely used in training neural networks. (See additional discussion in Appendix G.3). Moreover, Defazio & Mishchenko (2022)

only prove parameter-free rates of convergence in the non-stochastic setting, while we establish high probability guarantees in the stochastic setting. Concurrently with our work, Defazio & Mishchenko (2023) heuristically extended their dual averaging scheme to SGD- and Adam-like algorithms, reporting promising experimental results.

Finally, a number neural network optimization methods—LARS (You et al., 2017a), LAMB (You et al., 2017b), Adafactor (Shazeer & Stern, 2018), and Fromage (Bernstein et al., 2020)—use the norm of neural network weights to scale the step size. DOG and L-DOG are similar in also using a norm to scale their step size, but they differ from prior work by considering the distance from initialization rather than the norm of the weights. We believe that this difference is crucial in making DOG parameter-free, while the above-mentioned method have a learning-rate parameter to tune (though Bernstein et al. (2020) report that a single default value works well for across different tasks).

6. Limitations and Outlook

Our theoretical and empirical results place DOG as a promising step toward a new generation of principled and efficient tuning-free optimization algorithms. However, much additional work is necessary for these algorithms to become ubiquitous. First, it is important to understand how to correctly combine DOG with proven technique such as momentum, per-parameter learning rates, and learning rate annealing—this is a challenge both from a theoretical and a practical perspective. Second, it is important to gain a better understanding of situations where DOG is more sensitive to the choice of r_ϵ than theory would have us expect. Our preliminary investigations suggest a connection to batch normalization, and following that lead could lead to even more robust training methods. Finally, while our experiments aim to cover a broad range of tasks and architectures, future work needs to explore DOG in additional settings, particularly those involving training from scratch.

Acknowledgments

We thank Mitchell Wortsman, Simon Kornblith, Francesco Orabona and our anonymous reviewers for their insightful comments. This work was supported by the NSF-BSF program, under NSF grant #2239527 and BSF grant #2022663. MI acknowledges support from the Israeli council of higher education. OH acknowledges support from Pitt Momentum Funds, and AFOSR grant #FA9550-23-1-0242. YC acknowledges support from the Israeli Science Foundation (ISF) grant no. 2486/21, the Alon Fellowship, the Yandex Initiative for Machine Learning, and the Len Blavatnik and the Blavatnik Family Foundation.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.

Agarwal, A., Bartlett, P. L., Ravikumar, P., and Wainwright, M. J. Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization. *IEEE Transactions on Information Theory*, 58(5):3235–3249, 2012.

Arrow, K. J. and Enthoven, A. C. Quasi-concave programming. *Econometrica: Journal of the Econometric Society*, pp. 779–800, 1961.

Asi, H. and Duchi, J. C. The importance of better models in stochastic optimization. *Proceedings of the National Academy of Sciences*, 116(46):22924–22930, 2019.

Bar Haim, R., Dagan, I., Dolan, B., Ferro, L., Giampiccolo, D., Magnini, B., and Szpektor, I. The second PASCAL recognising textual entailment challenge. In *Proceedings of the Second PASCAL Challenges Workshop on Recognising Textual Entailment*, 2006.

Beattie, C., Leibo, J. Z., Teplyashin, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., et al. Deepmind lab. *arXiv:1612.03801*, 2016.

Bentivogli, L., Dagan, I., Dang, H. T., Giampiccolo, D., and Magnini, B. The fifth PASCAL recognizing textual entailment challenge. In *Text Analysis Conference (TAC)*, 2009.

Bernstein, J. R., Vahdat, A., Yue, Y., and Liu, M.-Y. On the distance between two neural networks and the stability of learning. *arXiv:2002.03432*, 2020.

Berrada, L., Zisserman, A., and Kumar, M. P. Training neural networks for and by interpolation. In *International Conference on Machine Learning (ICML)*, 2020.

Bhaskara, A., Cutkosky, A., Kumar, R., and Purohit, M. Online learning with imperfect hints. In *International Conference on Machine Learning (ICML)*, 2020.

Carmon, Y. and Hinder, O. Making SGD parameter-free. In *Conference on Learning Theory (COLT)*, 2022.

Carmon, Y., Raghunathan, A., Schmidt, L., Duchi, J. C., and Liang, P. S. Unlabeled data improves adversarial robustness. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Cer, D. M., Diab, M. T., Agirre, E., Lopez-Gazpio, I., and Specia, L. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *International Workshop on Semantic Evaluation*, 2017.

Chandra, K., Xie, A., Ragan-Kelley, J., and Meijer, E. Gradient descent: The ultimate optimizer. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Chen, K., Langford, J., and Orabona, F. Better parameter-free stochastic optimization with ODE updates for coin-betting. In *AAAI Conference on Artificial Intelligence*, 2022.

Cheng, G., Han, J., and Lu, X. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017.

Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., and Vedaldi, A. Describing textures in the wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.

Comet.ML. Comet.ML home page, 2021. URL <https://www.comet.ml/>.

Cubuk, E. D., Zoph, B., Mane, D., Vasudevan, V., and Le, Q. V. AutoAugment: Learning augmentation strategies from data. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

Cutkosky, A. Artificial constraints and hints for unbounded online learning. In *Conference on Learning Theory (COLT)*, 2019.

Cutkosky, A. and Orabona, F. Black-box reductions for parameter-free online learning in Banach spaces. In *Conference on Learning Theory (COLT)*, 2018.

Dagan, I., Glickman, O., and Magnini, B. The PASCAL recognising textual entailment challenge. In *Machine learning challenges. Evaluating predictive uncertainty, visual object classification, and recognising textual entailment*. Springer, 2006.

Davis, D. and Drusvyatskiy, D. Stochastic model-based minimization of weakly convex functions. *SIAM Journal on Optimization*, 29(1):207–239, 2019.

Defazio, A. and Mishchenko, K. Parameter free dual averaging: Optimizing lipschitz functions in a single pass. In *OPT 2022: NeurIPS Workshop on Optimization for Machine Learning*, 2022.

Defazio, A. and Mishchenko, K. Learning-rate-free learning by D-adaptation. In *International Conference on Machine Learning (ICML)*, 2023.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.

Dolan, W. B. and Brockett, C. Automatically constructing a corpus of sentential paraphrases. In *International Workshop on Paraphrasing (IWP2005)*, 2005.

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2021.

Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7), 2011.

Faw, M., Tziotis, I., Caramanis, C., Mokhtari, A., Shakkottai, S., and Ward, R. The power of adaptivity in SGD: Self-tuning step sizes with unbounded gradients and affine variance. In *Conference on Learning Theory (COLT)*, 2022.

Fei-Fei, L., Fergus, R., and Perona, P. Learning generative visual models from few training examples: An incremental Bayesian approach tested on 101 object categories. In *CVPR Workshop*, 2004.

Giampiccolo, D., Magnini, B., Dagan, I., and Dolan, B. The third PASCAL recognizing textual entailment challenge. In *ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, 2007.

Gupta, V., Koren, T., and Singer, Y. A unified approach to adaptive regularization in online and stochastic optimization. *arXiv:1706.06569*, 2017.

Gupta, V., Koren, T., and Singer, Y. Shampoo: Preconditioned stochastic tensor optimization. In *International Conference on Machine Learning (ICML)*, 2018.

Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020.

Hazan, E. and Kakade, S. Revisiting the Polyak step size. *arXiv:1905.00313*, 2019.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

Hinder, O., Sidford, A., and Sohoni, N. Near-optimal methods for minimizing star-convex functions and beyond. In *Conference on Learning Theory (COLT)*, 2020.

Howard, S. R., Ramdas, A., McAuliffe, J., and Sekhon, J. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021.

Huang, G., Liu, Z., and Weinberger, K. Q. Densely connected convolutional networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

Iyer, S., Dandekar, N., and Csernai, K. First quora dataset release: Question pairs, 2017. URL <https://data.quora.com/First-Quora-Dataset-Release-Question-Pairs>.

Jacobsen, A. and Cutkosky, A. Parameter-free mirror descent. In *Conference on Learning Theory (COLT)*, 2022.

Johnson, J., Hariharan, B., Van Der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., and Girshick, R. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

Kaggle and EyePacs. Kaggle diabetic retinopathy detection, 2015. URL <https://www.kaggle.com/c/diabetic-retinopathy-detection/data>.

Karimi, H., Nutini, J., and Schmidt, M. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD)*, 2016.

Kempka, M., Kotlowski, W., and Warmuth, M. K. Adaptive scale-invariant online algorithms for learning linear models. In *International Conference on Machine Learning (ICML)*, 2019.

Kingma, D. P. and Ba, J. ADAM: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.

Kleinberg, B., Li, Y., and Yuan, Y. An alternative view: When does SGD escape local minima? In *International Conference on Machine Learning (ICML)*, 2018.

Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

Levesque, H. J., Davis, E., and Morgenstern, L. The Winograd schema challenge. In *International Conference on Principles of Knowledge Representation and Reasoning*, 2011.

Levy, D., Carmon, Y., Duchi, J. C., and Sidford, A. Large-scale methods for distributionally robust optimization. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Lhoest, Q., Villanova del Moral, A., Jernite, Y., Thakur, A., von Platen, P., Patil, S., Chaumond, J., Drame, M., Plu, J., Tunstall, L., Davison, J., Šaško, M., Chhablani, G., Malik, B., Brandeis, S., Le Scao, T., Sanh, V., Xu, C., Patry, N., McMillan-Major, A., Schmid, P., Gugger, S., Delangue, C., Matussière, T., Debut, L., Bekman, S., Cistac, P., Goehringer, T., Mustar, V., Lagunas, F., Rush, A., and Wolf, T. Datasets: A community library for natural language processing. In *Conference on Empirical Methods in Natural Language Processing (EMNLP): System Demonstrations*, 2021.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv:1907.11692*, 2019.

Liu, Z., Mao, H., Wu, C., Feichtenhofer, C., Darrell, T., and Xie, S. A ConvNet for the 2020s. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

Loizou, N., Vaswani, S., Laradji, I. H., and Lacoste-Julien, S. Stochastic Polyak step-size for SGD: An adaptive learning rate for fast convergence. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.

Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019.

Luo, H. and Schapire, R. E. Achieving all with no parameters: AdaNormalHedge. In *Conference on Learning Theory (COLT)*, 2015.

Mangasarian, O. L. Pseudo-convex functions. In *Stochastic optimization models in finance*. Elsevier, 1975.

Matthey, L., Higgins, I., Hassabis, D., and Lerchner, A. dSprites: Disentanglement testing Sprites dataset. <https://github.com/deepmind/dsprites-dataset/>, 2017.

Mhammedi, Z. and Koolen, W. M. Lipschitz and comparator-norm adaptivity in online learning. In *Conference on Learning Theory (COLT)*, 2020.

Nemirovski, A. On parallel complexity of nonsmooth convex optimization. *Journal of Complexity*, 10(4):451–463, 1994.

Nemirovski, A. and Yudin, D. *Problem complexity and method efficiency in optimization*. Wiley-Interscience, 1983.

Nemirovski, A., Juditsky, A., Lan, G., and Shapiro, A. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.

Nesterov, Y. and Polyak, B. T. Cubic regularization of Newton method and its global performance. *Mathematical Programming*, 108(1):177–205, 2006.

Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011.

Nilsback, M.-E. and Zisserman, A. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP)*, 2008.

Orabona, F. Simultaneous model selection and optimization through parameter-free stochastic learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.

Orabona, F. and Pál, D. Coin betting and parameter-free online learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.

Orabona, F. and Pál, D. Parameter-free stochastic optimization of variationally coherent functions. *arXiv:2102.00236*, 2021.

Orabona, F. and Tommasi, T. Training deep networks without learning rates through coin betting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

Paquette, C. and Scheinberg, K. A stochastic line search method with expected complexity analysis. *SIAM Journal on Optimization*, 30(1):349–376, 2020.

Parkhi, O. M., Vedaldi, A., Zisserman, A., and Jawahar, C. Cats and dogs. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

Polyak, B. T. *Introduction to Optimization*. Optimization Software, Inc, 1987.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, 2021.

Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67, 2020.

Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. SQuAD: 100,000+ questions for machine comprehension of text. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.

Reddi, S. J., Kale, S., and Kumar, S. On the convergence of Adam and beyond. In *International Conference on Learning Representations (ICLR)*, 2018.

Rolinek, M. and Martius, G. L4: Practical loss-based step-size adaptation for deep learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Schaul, T., Zhang, S., and LeCun, Y. No more pesky learning rates. In *International Conference on Machine Learning (ICML)*, 2013.

Shamir, O. and Zhang, T. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *International Conference on Machine Learning (ICML)*, 2013.

Shazeer, N. and Stern, M. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning (ICML)*, 2018.

Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014.

Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., and Potts, C. Recursive deep models for semantic compositionality over a sentiment treebank. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2013.

Streeter, M. and McMahan, H. B. No-regret algorithms for unconstrained online convex optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2012.

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. Going deeper with convolutions. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

Vaswani, S., Mishkin, A., Laradji, I., Schmidt, M., Gidel, G., and Lacoste-Julien, S. Painless stochastic gradient: Interpolation, line-search, and convergence rates. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, I., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

Vovk, V. On-line regression competitive with reproducing kernel hilbert spaces. In *Theory and Applications of Models of Computation (TAMC)*, 2006.

Wang, A., Pruksachatkun, Y., Nangia, N., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. SuperGLUE: A stickier benchmark for general-purpose language understanding. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019a.

Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *International Conference on Learning Representations (ICLR)*, 2019b.

Ward, R., Wu, X., and Bottou, L. AdaGrad stepsizes: Sharp convergence over nonconvex landscapes. In *International Conference on Machine Learning (ICML)*, 2019.

Warstadt, A., Singh, A., and Bowman, S. R. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641, 2019.

Wes McKinney. Data Structures for Statistical Computing in Python. In *Proceedings of the 9th Python in Science Conference*, 2010.

Wightman, R. PyTorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019.

Williams, A., Nangia, N., and Bowman, S. A broad-coverage challenge corpus for sentence understanding through inference. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2018.

Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. Transformers: State-of-the-art natural language processing. In *Conference on Empirical Methods in Natural Language Processing (EMNLP): System Demonstrations*, 2020.

Wortsman, M., Ilharco, G., Gadre, S. Y., Roelofs, R., Gontijo-Lopes, R., Morcos, A. S., Namkoong, H., Farhadi, A., Carmon, Y., Kornblith, S., et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning (ICML)*, 2022.

Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., and Torralba, A. Sun database: Large-scale scene recognition from abbey to zoo. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.

Xiao, J., Ehinger, K. A., Hays, J., Torralba, A., and Oliva, A. Sun database: Exploring a large collection of scene categories. *International Journal of Computer Vision*, 119(1):3–22, 2016.

You, Y., Gitman, I., and Ginsburg, B. Large batch training of convolutional networks. *arXiv:1708.03888*, 2017a.

You, Y., Gitman, I., and Ginsburg, B. Scaling SGD batch size to 32k for ImageNet training. *arXiv:1708.03888*, 2017b.

You, Y., Li, J., Reddi, S., Hseu, J., Kumar, S., Bhojanapalli, S., Song, X., Demmel, J., Keutzer, K., and Hsieh, C.-J. Large batch optimization for deep learning: Training BERT in 76 minutes. In *International Conference on Learning Representations (ICLR)*, 2020.

Zagoruyko, S. and Komodakis, N. Wide residual networks. In *British Machine Vision Conference (BMVC)*, 2016.

Zhai, X., Puigcerver, J., Kolesnikov, A., Ruyssen, P., Riquelme, C., Lucic, M., Djolonga, J., Pinto, A. S., Neumann, M., Dosovitskiy, A., Beyer, L., Bachem, O., Tschanne, M., Michalski, M., Bousquet, O., Gelly, S., and Houlsby, N. A large-scale study of representation learning with the visual task adaptation benchmark. *arXiv:1910.04867*, 2019.

Zhang, J. and Cutkosky, A. Parameter-free regret in high probability with heavy tails. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Zhang, Z., Cutkosky, A., and Paschalidis, I. PDE-based optimal strategy for unconstrained online learning. In *International Conference on Machine Learning (ICML)*, 2022.

Zhou, Y., Yang, J., Zhang, H., Liang, Y., and Tarokh, V. SGD converges to global minimum in deep learning via star-convex path. In *International Conference on Learning Representations (ICLR)*, 2019.

A. Relaxing the Convexity Assumption

This section describes relaxations of convexity under which our main theoretical results still hold. In particular, our results naturally extend to star-convex functions (Nesterov & Polyak, 2006) which satisfy

$$f(x) - f_* \leq \langle \nabla f(x), x - x_* \rangle \quad \text{for all } x \in \mathcal{X}.$$

Our results also extend (with changed constants) to quasarconvex functions (Hinder et al., 2020), which require that $f(x) - f_* \leq c \langle \nabla f(x), x - x_* \rangle$ holds for some $c < \infty$ and all $x \in \mathcal{X}$. A further relaxation of star convexity requires it to hold only along the optimization trajectory:

Assumption A.1 (Zhou et al. (2019, Definition 2)). There exists $x_* \in \arg \min_x f(x)$ and constant $c < \infty$ such that the iterates of SGD satisfy

$$f(x_k) - f_* \leq c \langle \nabla f(x_k), x_k - x_* \rangle \quad \text{for all } k$$

almost surely.

Zhou et al. (2019) introduce this notion of a “star-convex path” and provide some empirical evidence that it may hold when training deep neural networks with SGD (see also Kleinberg et al. (2018) for a related assumption). Zhou et al. (2019) also prove that the assumption suffices to prove that SGD converges to the global minimizer; it suffices for DOG for similar reasons.

When substituting Assumption 3.1 with Assumption A.1 our analysis goes through unchanged, except we can no longer use Jensen’s inequality to argue directly about the suboptimality of the point \bar{x}_τ . Instead, Theorem 3.10 with Assumption A.1 says that, with probability at least $1 - \delta$,

$$\sum_{k=0}^{\tau-1} \omega_k (f(x_k) - f_*) \leq O\left(c_{\delta, r_\epsilon, T} \cdot \frac{d_0 \sqrt{G_\tau + L^2}}{T}\right),$$

with $\omega_k := \frac{\bar{r}_k}{\sum_{i=0}^{\tau-1} \bar{r}_i}$ and τ and $c_{\delta, r_\epsilon, T}$ as defined in Theorem 3.10. (Note that Assumption A.1 implies $\sum_{k=0}^{\tau-1} \omega_k (f(x_k) - f_*) \leq \sum_{k=0}^{\tau-1} \omega_k \langle \nabla f(x_k), x_k - x_* \rangle$ which replaces (3)).

We can turn the above bound into a constant-probability guarantee for a specific T-DOG iterate x_K by sampling $K \sim \omega$ and using Markov’s inequality:

$$\mathbb{P}\left(f(x_K) - f_* \leq e \sum_{k=0}^{\tau-1} \omega_k (f(x_k) - f_*)\right) \leq e^{-1}.$$

To obtain a high probability guarantee, we can make $l = \lceil \log \frac{1}{\delta} \rceil$ independent draws from ω , denoted K_1, \dots, K_l and use the fact that

$$\mathbb{P}\left(\min_{i \leq l} f(x_{K_i}) - f_* \leq e \sum_{k=0}^{\tau-1} \omega_k (f(x_k) - f_*)\right) \leq \delta.$$

Finding the i that minimizes $f(x_{K_i})$ requires a logarithmic number of evaluations of the exact objective. When this is not feasible, we can instead consider a statistical learning setup where we have sample access to stochastic functions $F(x)$ such that $\mathbb{E}F(x) = f(x)$ for all x and, almost surely, F is L_* Lipschitz in a ball of radius $3d_0$ around x_0 . (The stochastic subgradient oracle $\mathcal{G}(x)$ is then implemented by sampling F and returning its subgradient at x). We can then sample T new stochastic functions F_1, \dots, F_T and select $K^* \in \arg \min_{k \in \{K_1, \dots, K_l\}} \sum_{i=1}^T F_i(x_k)$. Straightforward application of Hoeffding’s inequality shows that (when $\bar{r}_T \leq 3d_0$)

$$f(x_{K^*}) - f_* \leq \min_{i \leq l} f(x_{K_i}) - f_* + O\left(\frac{L_* d_0}{\sqrt{T}} \sqrt{\log \frac{1}{\delta}}\right)$$

with probability at least $1 - \delta$.

We remark that the literature contains a plethora of other convexity relaxations such as quasiconvexity (Arrow & Enthoven, 1961), pseudoconvexity (Mangasarian, 1975), Polyak-Łojasiewicz conditions (Karimi et al., 2016) and weak convexity (Davis & Drusvyatskiy, 2019). Exploring the convergence of DOG under these additional convexity relaxations is left to future work.

B. Relaxing the Global Stochastic Gradient Bound Assumption

Our results continue to hold essentially unchanged if we replace Assumption 3.2 (globally bounded stochastic gradient norm) with the following.

Assumption B.1 (Pointwise bounded stochastic gradients). There exists some continuous function $\ell : \mathcal{X} \rightarrow \mathbb{R}$ such that $\|\mathcal{G}(x)\| \leq \ell(x)$ almost surely.

In particular, if we set

$$L_* := \max_{x \in \mathcal{X}: \|x - x_0\| \leq 3\|x_0 - x_*\|} \ell(x)$$

then Theorem 3.10 holds under Assumption B.1 for any $L \geq L_*$. For full details and proof, refer to the arXiv version of our paper, available at <https://arxiv.org/abs/2302.12022>, where we use Assumption B.1 throughout.

Assumption B.1 meaningfully relaxes Assumption 3.2 when ℓ can grow very large in \mathcal{X} . For example, consider unconstrained least squares problems (with $\mathcal{X} = \mathbb{R}^m$) with stochastic gradient oracle $\mathcal{G}(x) = (\langle a, x \rangle - b)a$ for random $a \in \mathbb{R}^m$ and $b \in \mathbb{R}$, such that $\|a\| \leq 1$ and $|b| \leq 1$ hold with probability 1. In this case, there is no global upper bound on $\|\mathcal{G}(x)\|$, but Assumption B.1 holds with $\ell(x) = \|x\| + 1$ and $L_* = 3\|x_*\| + 1$ (assuming $x_0 = 0$).

However, to apply the (T-DOG) formula, one still requires an a-priori upper bound on L_* . In the above least-squares example, it is possible to bound L_* given an upper bound D on $d_0 = \|x_*\|$. However, if such a bound is available, then it is also possible to constrain the domain to a ball of radius D , where a global norm bound holds. Indeed, for all examples that we are aware of that have an a-priori bound on L_* , this bound on L_* arises from a bound on d_0 .⁸ In the arXiv version of our paper we solve this issue by changing the T-DOG formula to $\eta_t = \bar{r}_t / \sqrt{G'_t}$ with

$$G'_t = 8^4 \theta_{T,\delta}^2 \log_+^2 \left(\frac{t\bar{\ell}_t^2}{\bar{\ell}_0^2} \right) (G_{t-1} + 16\bar{\ell}_t^2)$$

where $\bar{\ell}_t := \max_{i \leq t} \ell(x_i)$. This variant of T-DOG attains essentially the same guarantees as the one presented here, but requires no prior knowledge of L_* .

We note that Assumption B.1 and techniques similar to the T-DOG variant described above have previously appeared in the parameter-free online learning literature (Cutkosky, 2019; Mhammedi & Koolen, 2020). However, these works do not also guarantee that the iterates remain close to the optimal solution and therefore do not obtain our dependence on the local Lipschitz constant L_* .

C. Useful Algebraic Facts

C.1. Lemma C.1

Lemma C.1. *Let a_0, \dots, a_t be a nondecreasing sequence of nonnegative numbers. Then*

$$\sum_{k=1}^t \frac{a_k - a_{k-1}}{\sqrt{a_k}} \leq 2(\sqrt{a_t} - \sqrt{a_0}).$$

Proof. We have

$$\sum_{k=1}^t \frac{a_k - a_{k-1}}{\sqrt{a_k}} = \sum_{k=1}^t \frac{(\sqrt{a_k} - \sqrt{a_{k-1}})(\sqrt{a_k} + \sqrt{a_{k-1}})}{\sqrt{a_k}} \leq 2 \sum_{k=1}^t (\sqrt{a_k} - \sqrt{a_{k-1}}) \leq 2(\sqrt{a_t} - \sqrt{a_0}).$$

□

C.2. Lemma C.2

Lemma C.2. *Let a_1, \dots, a_T and b_1, \dots, b_T be sequences in \mathbb{R} such that $\{a_i\}$ is nonnegative and nondecreasing. Then, for all $t \leq T$,*

$$\left| \sum_{i=1}^t a_i b_i \right| \leq 2a_t \max_{i \leq t} \left| \sum_{i=1}^t b_i \right|.$$

⁸We thank an anonymous reviewer for pointing this out.

Proof. Let $a'_i = a_i - a_{i-1}$ and $B_i = \sum_{j \leq i} b_j$. Then (by discrete integration by parts)

$$\sum_{i=1}^t a_i b_i = \sum_{i=1}^t a_i (B_i - B_{i-1}) = \sum_{i=1}^{t-1} (a_i - a_{i+1}) B_i + a_t B_t = a_t B_t - \sum_{i=1}^{t-1} a'_{i+1} B_i.$$

Therefore

$$\left| \sum_{i=1}^t a_i b_i \right| \stackrel{(i)}{\leq} |a_t B_t| + \left(\sum_{i=1}^{t-1} |a'_{i+1}| \right) \max_{i < t} |B_i| \leq \left(|a_t| + \sum_{i=1}^{t-1} |a_{i+1} - a_i| \right) \max_{i \leq t} |B_i| \stackrel{(ii)}{\leq} 2a_t \max_{i \leq t} |B_i|,$$

where we used (i) the triangle and Hölder's inequality, and (ii) that a_t is nonnegative and nondecreasing and therefore $\sum_{i=1}^{t-1} |a_{i+1} - a_i| = a_t - a_1 \leq a_t$. \square

C.3. Proof of Lemma 3.7

Proof. Define $K := \lceil \log(s_T/s_0) \rceil$, and $n := \lfloor \frac{T}{K} \rfloor$. Then, we have

$$\log \left(\frac{s_T}{s_0} \right) \geq \sum_{k=0}^{K-1} \log \left(\frac{s_{n(k+1)}}{s_{nk}} \right) \geq K \min_{k < K} \log \left(\frac{s_{n(k+1)}}{s_{nk}} \right).$$

Rearranging and using the definition of K gives

$$\min_{k < K} \log \left(\frac{s_{n(k+1)}}{s_{nk}} \right) \leq \frac{\log \left(\frac{s_T}{s_0} \right)}{K} \leq 1 \implies \min_{k < K} \frac{s_{n(k+1)}}{s_{nk}} \leq e.$$

where the implication follows from monotonicity of the exponential function. Therefore,

$$\max_{t \leq T} \sum_{i < t} \frac{s_i}{s_t} \geq \max_{t \in [n, T]} n \frac{s_{t-n}}{s_t} = \max_{k \leq K} n \frac{s_{n(k-1)}}{s_{nk}} \geq n e^{-1} = e^{-1} \left\lfloor \frac{T}{\lceil \log(s_T/s_0) \rceil} \right\rfloor \geq e^{-1} \frac{T}{\log(s_T/s_0) + 1} - e^{-1},$$

where the first inequality uses that s is positive nondecreasing sequence and the second inequality uses $\min_{k < K} \frac{s_{n(k+1)}}{s_{nk}} \leq e$ as shown above. \square

C.4. Lemma C.3

Recall that $\log_+(z) := 1 + \log(t)$.

Lemma C.3. *Let $a_{-1}, a_0, a_1, \dots, a_t$ be a nondecreasing sequence of nonnegative numbers, then*

$$\sum_{k=0}^t \frac{a_k - a_{k-1}}{a_k \log_+^2(a_k/a_{-1})} \leq 1.$$

Proof. We have

$$\begin{aligned} \sum_{k=0}^t \frac{a_k - a_{k-1}}{a_k \log_+^2(a_k/a_{-1})} &\leq \sum_{k=0}^t \int_{a_{k-1}/a_0}^{a_k/a_{-1}} \frac{d\alpha}{\alpha \log_+^2(\alpha)} = \int_1^{a_t/a_{-1}} \frac{d\alpha}{\alpha \log_+^2(\alpha)} \\ &\leq \int_1^\infty \frac{d\alpha}{\alpha \log_+^2(\alpha)} = \left[\frac{1}{1 + \log(\alpha)} \right]_1^\infty = 1. \end{aligned}$$

\square

D. Proofs for Section 3

D.1. Proof of Lemma 3.4

Proof. Using $x_{k+1} = \text{Proj}_{\mathcal{X}}(x_k - \eta_k g_k)$ we obtain the standard inequality $d_{k+1}^2 \leq \|x_k - \eta_k g_k - x_\star\|^2 = d_k^2 - 2\eta_k \langle g_k, x_k - x_\star \rangle + \eta_k^2 \|g_k\|^2$. Rearranging this gives:

$$\langle g_k, x_k - x_\star \rangle \leq \frac{d_k^2 - d_{k+1}^2}{2\eta_k} + \frac{\eta_k \|g_k\|^2}{2}. \quad (5)$$

Therefore, $\sum_{k=0}^{t-1} \bar{r}_k \langle g_k, x_k - x_\star \rangle$ is at most

$$\underbrace{\frac{1}{2} \sum_{k=0}^{t-1} \frac{\bar{r}_k}{\eta_k} (d_k^2 - d_{k+1}^2)}_{(A)} + \underbrace{\frac{1}{2} \sum_{k=0}^{t-1} \bar{r}_k \eta_k \|g_k\|^2}_{(B)}.$$

We will bound the terms (A) and (B) in turn, beginning with the former:

$$\begin{aligned} (A) &= \sum_{k=0}^{t-1} \sqrt{G'_k} (d_k^2 - d_{k+1}^2) = d_0^2 \sqrt{G'_0} - d_t^2 \sqrt{G'_{t-1}} + \sum_{k=1}^{t-1} d_k^2 \left(\sqrt{G'_k} - \sqrt{G'_{k-1}} \right) \\ &\stackrel{(i)}{\leq} \bar{d}_t^2 \sqrt{G'_0} - d_t^2 \sqrt{G'_{t-1}} + \bar{d}_t^2 \sum_{k=1}^{t-1} \left(\sqrt{G'_k} - \sqrt{G'_{k-1}} \right) = \sqrt{G'_{t-1}} (\bar{d}_t^2 - d_t^2) \stackrel{(ii)}{\leq} 4\bar{r}_t \bar{d}_t \sqrt{G'_{t-1}}. \end{aligned}$$

Inequality (i) uses $d_k \leq \bar{d}_t$ and that G'_k is nondecreasing as per Definition 3.3. Inequality (ii) holds since, for $s \in \arg \max_{k \leq t} d_k$, we have $\bar{d}_t^2 - d_t^2 = d_s^2 - d_t^2 = (d_s - d_t)(d_s + d_t) \leq \|x_s - x_t\|(d_s + d_t) \leq (\bar{r}_s + \bar{r}_t)(d_s + d_t) \leq 4\bar{r}_t \bar{d}_t$. Bounding the second term (B) , we have:

$$(B) = \sum_{k=0}^{t-1} \frac{\bar{r}_k^2 \|g_k\|^2}{\sqrt{G'_k}} \leq \sum_{k=0}^{t-1} \frac{\bar{r}_k^2 \|g_k\|^2}{\sqrt{G_k}} \leq \bar{r}_t^2 \sum_{k=0}^{t-1} \frac{\|g_k\|^2}{\sqrt{G_k}} \leq 2\bar{r}_t^2 \sqrt{G_{t-1}},$$

where the final inequality uses the standard Lemma C.1 with $a_k = G_k = \sum_{i \leq k} \|g_i\|^2$. \square

D.2. Proof of Lemma 3.5

We begin by citing the following corollary of a general bound due to Howard et al. (2021). (Recall that $\theta_{t,\delta} := \log \frac{60 \log(6t)}{\delta}$).

Corollary D.1 (Carmon & Hinder (2022, Corollary 1)). *Let $c > 0$ and X_t be a martingale difference sequence adapted to \mathcal{F}_t such that $|X_t| \leq c$ with probability 1 for all t . Then, for all $\delta \in (0, 1)$, and $\hat{X}_t \in \mathcal{F}_{t-1}$ such that $|\hat{X}_t| \leq c$ with probability 1,*

$$\mathbb{P} \left(\exists t \leq T : \left| \sum_{s=1}^t X_s \right| \geq 4 \sqrt{\theta_{t,\delta} \sum_{s=1}^t (X_s - \hat{X}_s)^2 + c^2 \theta_{t,\delta}^2} \right) \leq \delta.$$

Next, we connect Corollary D.1 with a handy algebraic fact (Lemma C.2) to obtain the following result, which underpins Lemma 3.5.

Lemma D.2. *Let S be the set of nonnegative and nondecreasing sequences. Let $c > 0$ and let X_t be a martingale difference sequence adapted to \mathcal{F}_t such that $|X_t| \leq c$ with probability 1 for all t . Then, for all $\delta \in (0, 1)$, and $\hat{X}_t \in \mathcal{F}_{t-1}$ such that $|\hat{X}_t| \leq c$ with probability 1,*

$$\mathbb{P} \left(\exists t \leq T, \exists \{y_i\}_{i=1}^\infty \in S : \left| \sum_{i=1}^t y_i X_i \right| \geq 8y_t \sqrt{\theta_{t,\delta} \sum_{i=1}^t (X_i - \hat{X}_i)^2 + c^2 \theta_{t,\delta}^2} \right) \leq \delta.$$

Proof. Follows from Lemma C.2 (with y_i and X_i taking the roles of a_i and b_i , respectively), and Corollary D.1 that bounds $\max_{i \leq t} \left| \sum_{i \leq t} X_i \right|$ for all $t \leq T$. \square

Proof of Lemma 3.5. For $k \in [T]$ define the random variables:

$$Y_k = \bar{r}_k \bar{d}_k, \quad X_k = \left\langle \Delta_k, \frac{x_k - x_\star}{\bar{d}_k} \right\rangle, \quad \text{and} \quad \hat{X}_k = - \left\langle \nabla f(x_k), \frac{x_k - x_\star}{\bar{d}_k} \right\rangle.$$

From these definitions we get

$$\sum_{k=0}^{t-1} Y_k X_k = \sum_{k=0}^{t-1} \bar{r}_k \langle \Delta_k, x_k - x_\star \rangle.$$

Therefore,

$$\begin{aligned} & \mathbb{P} \left(\exists t \leq T : \left| \sum_{k=0}^{t-1} \bar{r}_k \langle \Delta_k, x_k - x_\star \rangle \right| \geq 8\bar{r}_{t-1} \bar{d}_{t-1} \sqrt{\theta_{t,\delta} G_{t-1} + L^2 \theta_{t,\delta}^2} \right) \\ & \leq \mathbb{P} \left(\exists t \leq T : \left| \sum_{k=0}^{t-1} Y_k X_k \right| \geq 8Y_t \sqrt{\theta_{t,\delta} \sum_{k=0}^{t-1} (X_k - \hat{X}_k)^2 + L^2 \theta_{t,\delta}^2} \right) \leq \delta \end{aligned}$$

where the last inequality uses Lemma D.2. \square

D.3. Proof of Corollary 3.8

Proof. If $T > 2 \log_+ (\bar{r}_T / r_\epsilon)$ then the corollary follows by Proposition 3.6 and Lemma 3.7 with $s_t = \bar{r}_t$. For the corner case when $T \leq 2 \log_+ (\bar{r}_T / r_\epsilon)$ we use that $f(\bar{x}_\tau) - f_\star \leq O(L\bar{d}_\tau) \leq O(L(\bar{r}_\tau + d_0))$ where the first inequality uses (3), Cauchy-Schwarz and that $\|\nabla f(x_t)\| \leq L$; the second inequality uses the triangle inequality. \square

D.4. DOG can Diverge on a Pathological Instance

Consider the following variant of Nemirovski's function (Nemirovski & Yudin, 1983; Nemirovski, 1994) defined on \mathbb{R}^m :

$$f(x) = \max_{i \leq m} \max \left\{ [x]_i, -\frac{1}{\sqrt{m}} [x]_i \right\},$$

where $[x]_i$ denotes the i 'th coordinate of x and $[x_0]_i = 10r_\epsilon / \sqrt{m}$ for all i , so that $d_0 = 10r_\epsilon > r_\epsilon$. We show that applying DOG on this function gives $\bar{r}_T / d_0 = \sqrt{T}/10$ for all $T \leq m$, meaning that the ratio \bar{r}_T / d_0 can be made arbitrarily large by increasing T and m .

Define

$$i(x) := \min_{i \leq m} \arg \max \left\{ [x]_i, -\frac{[x]_i}{\sqrt{m}} \right\},$$

i.e., $i(x)$ is the smallest coordinate which is candidate for providing a subgradient. Using this notation, a valid subgradient for f is:

$$\nabla f(x) = \begin{cases} e_{i(x)} & x_{i(x)} > 0 \\ -\frac{1}{\sqrt{m}} e_{i(x)} & \text{otherwise} \end{cases}$$

where e_j is a vector with one in the j th entry and zero elsewhere. With this subgradient choice for $k \leq m$ the iterates become:

$$[x_k]_j = \begin{cases} 10r_\epsilon / \sqrt{m} - r_\epsilon & j < k \\ 10r_\epsilon / \sqrt{m} & j \geq k \end{cases} \quad (6)$$

and therefore $\bar{r}_k = \sqrt{k}r_\epsilon = \sqrt{k}d_0/10$ as claimed. We confirm (6) by induction. Since $[x_0]_i = 10r_\epsilon / \sqrt{m}$ for all i , the expression (6) holds for $k = 0$. If (6) holds for all $k \leq n < m$ then

$$\nabla f(x_k) = e_k$$

and therefore $G_n = n$ so that $\eta_n = \frac{r_\epsilon \sqrt{n}}{\sqrt{n}} = r_\epsilon$ and $x_{n+1} = x_n - \frac{\sqrt{n}}{\sqrt{n}} r_\epsilon e_n$, meaning that

$$[x_{n+1}]_j = \begin{cases} 10r_\epsilon/\sqrt{m} - r_\epsilon & j < n+1 \\ 10r_\epsilon/\sqrt{m} & j \geq n+1 \end{cases}$$

which completes the induction.

D.5. Proof of Proposition 3.9

To show iterate boundedness in the stochastic setting we define the stopping time

$$\tau = \min\{t : \bar{r}_t > 3d_0\},$$

so that the event $\{\bar{r}_T \leq 3d_0\}$ is the same as $\{\tau > T\}$. Let η_k denote the sequence of **T-DOG** step sizes (for given L, T and δ). To facilitate our analysis we also define the following truncated step size sequence:

$$\tilde{\eta}_k := \begin{cases} \eta_k & k < \tau \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

Truncating the step size allows us to rigorously handle the possibility that \bar{r}_T exceeds $3d_0$. In particular, the following holds for $\{\tilde{\eta}_k\}$ but not for $\{\eta_k\}$. (Recall that $\Delta_t := g_t - \nabla f(x_t)$).

Lemma D.3. *Under Assumption 3.2 the truncated **T-DOG** step sizes (7) satisfy, for all $t \leq T$,*

$$\tilde{\eta}_t \in \sigma(g_0, \dots, g_{t-1}), \quad (8)$$

$$|\tilde{\eta}_t \langle \gamma, x_t - x_\star \rangle| \leq \frac{6d_0^2}{8^2 \theta_{T,\delta}} \text{ for } \gamma \in \{g_t, \nabla f(x_t), \Delta_t\}, \quad (9)$$

$$\sum_{k=0}^t \tilde{\eta}_k^2 \|g_k\|^2 \leq \frac{9d_0^2}{8^4 \theta_{T,\delta}}, \text{ and} \quad (10)$$

$$\sum_{k=0}^t (\tilde{\eta}_k \langle g_k, x_k - x_\star \rangle)^2 \leq \frac{12^2 d_0^4}{8^4 \theta_{T,\delta}}. \quad (11)$$

Proof. The bound (8) holds directly from the definition of **T-DOG** and (7).

To see the bound (9), first note that $\|\Delta_k\| \leq \|g_k\| + \|\nabla f(x_k)\| \leq 2L$. Since $G'_t \geq 4^2 8^4 L^2 \theta_{T,\delta}^2$ for all t , the Cauchy-Schwartz inequality gives

$$|\tilde{\eta}_t \langle \Delta_t, x_t - x_\star \rangle| \leq \frac{\bar{r}_t}{\sqrt{G'_t}} \|\Delta_t\| d_t \leq \frac{1}{2 \cdot 8^2 \theta_{T,\delta}} \bar{r}_T d_t \leq \frac{6d_0^2}{8^2 \theta_{T,\delta}},$$

where the last inequality uses again $\bar{r}_t \leq 3d_0$ and $d_t \leq d_0 + \bar{r}_t$ by the triangle inequality. Bounds for $|\tilde{\eta}_t \langle \gamma, x_t - x_\star \rangle|$ for $\gamma \in \{g_t, \nabla f(x_t)\}$ follow by the same argument.

To establish (10), first note that $\sum_{k=0}^t \tilde{\eta}_k^2 \|g_k\|^2 \leq \sum_{k=0}^{\tau-1} \tilde{\eta}_k^2 \|g_k\|^2$ by the definition of $\tilde{\eta}_k$. Furthermore

$$\sum_{k=0}^{\tau-1} \tilde{\eta}_k^2 \|g_k\|^2 = \sum_{k=0}^{\tau-1} \frac{\bar{r}_k^2 \|g_k\|^2}{G'_k} \stackrel{(i)}{\leq} \frac{\bar{r}_{\tau-1}^2}{8^4 \theta_{T,\delta}} \sum_{k=0}^{\tau-1} \frac{G_k - G_{k-1}}{(G_k + L^2) \log_+^2 \frac{G_k + L^2}{L^2}} \stackrel{(ii)}{\leq} \frac{9d_0^2}{8^4 \theta_{T,\delta}},$$

where (i) uses that $\|g_k\|^2 = G_k - G_{k-1}$ (with the shorthand $G_{-1} := 0$) and

$$G'_k \geq 8^4 \theta_{T,\delta} (G_k + L^2) \log_+^2 (k+1) \geq 8^4 \theta_{T,\delta} (G_k + L^2) \log_+^2 \frac{G_k + L^2}{L^2};$$

since $\|g_k\| \leq L$ and $G_k \leq kL^2$, while (ii) uses Lemma C.3 with $a_k = G_k + L^2$ and $\bar{r}_{\tau-1} \leq 3d_0$.

The final bound (11) follows immediately from (10) by noting that

$$\sum_{k=0}^t (\tilde{\eta}_k \langle g_k, x_k - x_\star \rangle)^2 \leq \sum_{k=0}^t \tilde{\eta}_k^2 \|g_k\|^2 d_k^2 \leq (4d_0)^2 \sum_{k=0}^t \tilde{\eta}_k^2 \|g_k\|^2,$$

where the first inequality follows from Cauchy-Schwartz and the second inequality from the fact that only terms with $k < \tau$ contribute to the sum. \square

The above properties allow us to establish the following concentration bound.

Lemma D.4. *In the setting of Lemma D.3,*

$$\mathbb{P} \left(\exists t \leq T : \sum_{k=0}^{t-1} \tilde{\eta}_k \langle \Delta_k, x_\star - x_k \rangle > d_0^2 \right) \leq \delta.$$

Proof. Consider the filtration $\mathcal{F}_t = \sigma(g_0, \dots, g_t)$ and define $X_t = \tilde{\eta}_t \langle \Delta_t, x_\star - x_t \rangle$ and $\hat{X}_t = -\tilde{\eta}_t \langle \nabla f(x_t), x_\star - x_t \rangle$. Then, by (8) we have that X_t is martingale difference sequence adapted to \mathcal{F}_t and $\hat{X}_t \in \mathcal{F}_{t-1}$. Moreover, by (9) we have that $\max\{|X_t|, |\hat{X}_t|\} \leq c$ almost surely for $c = \frac{24d_0^2}{8^4 \theta_{T,\delta}}$. Substituting into Corollary D.1 (and shifting the start of the summation from 1 to 0) we have

$$\mathbb{P} \left(\exists t \leq T : \left| \sum_{k=0}^{t-1} X_k \right| \geq 4 \sqrt{\theta_{t,\delta} \sum_{k=0}^{t-1} (X_k - \hat{X}_k)^2 + c^2 \theta_{t,\delta}^2} \right) \leq \delta.$$

Noting that $X_t - \hat{X}_t = \tilde{\eta}_t \langle g_t, x_\star - x_t \rangle$ and substituting the definition of c and the bound (11) gives, for every $t < T$,

$$4 \sqrt{\theta_{t,\delta} \sum_{k=0}^{t-1} (X_k - \hat{X}_k)^2 + c^2 \theta_{t,\delta}^2} \leq 4 \sqrt{\theta_{t,\delta} \frac{12^2 d_0^4}{8^4 \theta_{T,\delta}} + \left(\frac{6\theta_{t,\delta} d_0^2}{8^2 \theta_{T,\delta}} \right)^2} \leq d_0^2,$$

concluding the proof of lemma. \square

Finally, we show that the event defined in Lemma D.4 implies the desired distance bound.

Lemma D.5. *In the setting of Proposition 3.9, if $\sum_{k=0}^{t-1} \tilde{\eta}_k \langle \Delta_k, x_\star - x_k \rangle \leq d_0^2$ for all $t \leq T$ then $\tau > T$, i.e., $\bar{r}_T \leq 3d_0$.*

Proof. To condense notation, let $B_t := \max_{t' \leq t} \sum_{k=0}^{t'-1} \tilde{\eta}_k \langle \Delta_k, x_\star - x_k \rangle$, so that the claim becomes $B_t \leq d_0^2$ implies $\tau > t$ for all $t \leq T$. We prove the claim by induction on t . The basis of the induction is that $\tau > 0$ always holds since $\bar{r}_0 = r_\epsilon \leq 3d_0$ by assumption. For the induction step, we assume that $B_{t-1} \leq d_0^2$ implies $\tau \geq t$ and show that $B_t \leq d_0^2$ implies $\tau > t$. To that end, we use $\langle \nabla f(x_t), x_t - x_\star \rangle \geq f(x_t) - f_\star \geq 0$ to rearrange (5) and obtain

$$d_{k+1}^2 - d_k^2 \leq \eta_k^2 \|g_k\|^2 + 2\eta_k \langle \Delta_k, x_\star - x_k \rangle$$

for all k . Summing this inequality from $k = 0$ to $k = t-1$, we get

$$d_t^2 - d_0^2 \leq \sum_{k=0}^{t-1} \eta_k^2 \|g_k\|^2 + 2 \sum_{k=0}^{t-1} \eta_k \langle \Delta_k, x_\star - x_k \rangle = \sum_{k=0}^{t-1} \tilde{\eta}_k^2 \|g_k\|^2 + 2 \sum_{k=0}^{t-1} \tilde{\eta}_k \langle \Delta_k, x_k - x_\star \rangle,$$

where the equality holds since $\tau > t-1$ and therefore $\eta_k = \tilde{\eta}_k$ for all $k \leq t-1$. Now, by the bound (10) we have $\sum_{k=0}^{t-1} \tilde{\eta}_k^2 \|g_k\|^2 \leq \frac{9^2}{8^4 \theta_{T,\delta}} d_0^2 \leq d_0^2$. Moreover, $\sum_{k=0}^{t-1} \tilde{\eta}_k \langle \Delta_k, x_k - x_\star \rangle \leq B_t \leq d_0^2$ by definition and assumption, from which we conclude that $d_t^2 \leq 4d_0^2$ and hence $r_t \leq d_0 + d_t \leq 3d_0$. Since $\bar{r}_t = \max\{\bar{r}_{t-1}, r_t\}$ and $\bar{r}_{t-1} \leq 3d_0$ by the induction assumption, we have that $\bar{r}_t \leq 3d_0$ as well, concluding the proof. \square

Proposition 3.9 follows immediately from Lemmas D.4 and D.5.

D.6. Illustrating DOG's guarantees for least squares problems

In order to illustrate the advantage of T-DOG's iterate boundedness guarantee, we now instantiate Theorem 3.10 and Proposition 3.9 for the following stochastic least squares problem. Let P be a distribution over pairs $(a, b) \in \mathbb{R}^m \times \mathbb{R}$, and for $(a, b) \sim P$ consider the gradient oracle

$$\mathcal{G}(x) = (\langle a, x \rangle - b)a + \lambda x,$$

corresponding to the objective function

$$f(x) = \frac{1}{2} \mathbb{E}_{(a,b) \sim P} (\langle a, x \rangle - b)^2 + \frac{1}{2} \lambda \|x\|^2.$$

For simplicity, we set $x_0 = 0$.

Using λ -strong-convexity of f we can *crudely* bound the initial distance to optimality by $D = \sqrt{2f(0)/\lambda} > d_0$, noting that $f(0) = \frac{1}{2} \mathbb{E}_b b^2$ can be easily estimated from samples. If we also assume that $\|a\| \leq A$ and $|b| \leq B$ with probability 1, then $L = O((A^2 + \lambda)D + AB)$ is an upper bound on the stochastic gradient norm in the set $\mathcal{X} = \{x \mid \|x\| \leq D\}$.

The crudeness of the upper bound $D = \sqrt{2f(0)/\lambda}$ only affects T-DOG's rate of convergence through a lower order term. Specifically, the bound $d_T \leq O(\|x_\star\|)$ guarantees that $G_T \leq O(((A^2 + \lambda)\|x_\star\| + AB)^2 T)$, and therefore Theorem 3.10 implies that (with high probability) T-DOG's optimality gap is

$$\tilde{O} \left(\frac{(A^2 + \lambda)\|x_\star\|^2 + AB\|x_\star\|}{\sqrt{T}} + \frac{(A^2 + \lambda)\|x_\star\|\sqrt{f(0)/\lambda}}{T} \right).$$

Since $f(0) \leq \frac{1}{2}B^2$, the L -dependent term is of lower order as long as $\lambda = \omega(A^2/T)$.

For previously-proposed general purpose⁹ parameter-free methods (e.g. Orabona & Pál, 2016; Cutkosky & Orabona, 2018; Carmon & Hinder, 2022) the iterates x_t may be anywhere inside $\mathcal{X} = \{x \mid \|x\| \leq D\}$.¹⁰ Therefore, the best bounds for G_T would still involve the crude bound D , and be of the form $G_T \leq O(((A^2 + \lambda)D + AB)^2 T)$, leading to a worse overall error bound when $\|x_\star\| \ll D$. Furthermore, as we explain in Appendix B, for T-DOG with L as given above, there is no need to explicitly constrain the domain to $\mathcal{X} = \{x \mid \|x\| \leq D\}$.

E. Experiment Details

E.1. Environment settings

All experiments were based on PyTorch (Paszke et al., 2019) (version 1.12.0).

Language experiments were done with the *transformers* (Wolf et al., 2020) library (version 4.21.0) and tracked using the *Comet.ML* (Comet.ML, 2021). All datasets were provided by the *Datasets* library (Lhoest et al., 2021) (version 2.4.0) and were left as is, including train-eval-test splits.

Vision experiments were based on the *pytorch-image-models* (timm, version 0.7.0dev0) repository (Wightman, 2019), with *TensorFlow datasets* (version 4.6.0) as a dataset backend (Abadi et al., 2015).

To support the training and analysis of the results, we used *numpy* (Harris et al., 2020), *scipy* (Virtanen et al., 2020), *pandas* (Wes McKinney, 2010) and *scikit-learn* (Pedregosa et al., 2011).

E.2. Implementation details

Whenever possible, we used existing scripts and recipes provided by *timm* and *transformers* to fine-tune the models. We implemented DOG, L-DOG and the polynomial model averaging as a subclass of PyTorch *Optimizer* interface. We provide implementation of both in <https://github.com/formll/dog>.

⁹There exist parameter-free methods specialized to least-squares problems that obtain better without requiring an a-priori bound on the solution norm (Vovk, 2006).

¹⁰Carmon & Hinder (2022) guarantee boundedness of the point they output, but do not have guarantees on the magnitude of intermediate query points.

Table 2. Configuration used for each dataset in our testbed (Section 4.1). For all language tasks, we used the batch size as in Liu et al. (2019), and at least 150% the number of steps used there, in order to ensure convergence. Learning rate (LR) warmup and annealing refers to tuning with SGD and Adam. In all cases, both DOG and L-DOG used neither warmup nor annealing.

Task	Batch size	Steps	Metric	LR warmup	LR annealing	Grad. clipping
VTAB datasets	128	20K	Accuracy	None	Cosine	None
SQuAD	48	5475	F_1	10%	Cosine	1
SST-2	32	31407	Accuracy	10%	Cosine	1
CoLA	32	10000	Matthews correlation	10%	Cosine	1
MRPC	32	1734	F_1	10%	Cosine	1
STS-B	32	3281	Pearson correlation	10%	Cosine	1
QNLI	32	49218	Accuracy	10%	Cosine	1
RTE	32	10000	Accuracy	10%	Cosine	1
QQP	32	160625	F_1	10%	Cosine	1
MNLI	32	184218	Accuracy	10%	Cosine	1

E.3. Datasets

The datasets used in the language experiments are: **CoLA** (Warstadt et al., 2019), **SST-2** (Socher et al., 2013), **MRPC** (Dolan & Brockett, 2005), **QQP** (Iyer et al., 2017), **STS-B** (Cer et al., 2017), **MNLI** (Williams et al., 2018), **QNLI** (Rajpurkar et al., 2016), and **RTE** (Dagan et al., 2006; Bar Haim et al., 2006; Giampiccolo et al., 2007; Bentivogli et al., 2009). Following Liu et al. (2019), we discard WNLI (Levesque et al., 2011) as it was found to be ill-defined and was reformulated differently in SuperGLUE (Wang et al., 2019a).

The datasets used in the vision experiments are: The tasks are **Caltech101** (Fei-Fei et al., 2004), **CIFAR-100** (Krizhevsky, 2009), **CLEVR-Dist** (Johnson et al., 2017), **DMLab** (Beattie et al., 2016), **dSprites-Ori** (Matthey et al., 2017), **DTD** (Cimpoi et al., 2014), **Flowers102** (Nilsback & Zisserman, 2008), **Pets** (Parkhi et al., 2012), **Resisc45** (Cheng et al., 2017), **Retinopathy** (Kaggle & EyePacs, 2015), **Sun397** (Xiao et al., 2010; 2016), and **SVHN** (Netzer et al., 2011).

E.4. Models

When fine-tuning RoBERTa (from the ‘roberta-base’ checkpoint) on classification tasks, we follow the common technique of prepending an *CLS* token to the input, and feeding its final representation to a one hidden-layer, randomly initialized MLP that is used as a classification head. For SQuAD, the classification head is tasked with multi-label classification, predicting the probability that each word (token) in the input is the beginning/end of the answer span, and we then used the span that has the maximum likelihood as the model’s output. When fine-tuning T5 (from the ‘t5-base’ checkpoint), we treated all tasks as sequence-to-sequence tasks, translating classification labels to appropriate words (e.g. 0/1 to positive/negative) and then evaluated accuracy with exact match. The computer vision pre-trained models were accessed via `timm`, and had randomly initialized classification heads. The strings used to load the models were: ‘convnext_tiny’, ‘resnet50’, ‘densenet121’, ‘vit_base_patch32_224_in21k’ and ‘vgg11’.

E.5. Hyper-parameters

We trained each model/tasks combination a fixed number of steps (see Table 2), performing evaluation every 500 update steps (except for the smaller datasets Caltech101, DTD, Flowers102 and Pets where we evaluated every 200) with both the latest checkpoint, and the polynomial averaged one (see below). We did not use any weight decay. For language models, we left dropout at its default value in the transformers library. We used batch sizes as is common practice for each task, as detailed in Table 2.

Data augmentation in vision experiments. The VTAB suite (Zhai et al., 2019) divides its datasets into three categories: natural, specialized and structured, and we use a suitable data augmentation strategy for each of the categories. In particular, for structured datasets we simply resizes the images to a (224, 224) resolution, while for the natural and specialized datasets we use the standard ‘inception crop’ (Szegedy et al., 2015) at training time and a 0.875 center crop at test time. For natural datasets we additionally applied a color jitter operation with parameter 0.4 (as implemented in `timm`). Finally, we applied a random horizontal flip for all datasets except SVHN and dSprites-Ori, where such augmentation clearly interferes with the task.

Model selection in vision experiments. For computer vision experiments, we used the VTAB evaluation splits to select the best checkpoint, and then reported performance on the training split. Unlike the experiments accompanying the VTAB suite (Zhai et al., 2019), we did not retrain selected models on the combination of training and validation data.

Repeated runs. To account for randomness, we repeated our fine-tuning experiments using multiple seeds. In most cases (with exceptions listed below) we repeated each DOG and L-DOG training 5 times. For SGD and Adam repeating the learning with all learning rates was computationally prohibitive, so instead for each task / model pair we repeated 5 times only the best-performing LR (i.e., instance-tuned LR) and the best-performing LR across all tasks for that model (i.e., model-tuned LR) according the validation split. A few experiments were too computationally expensive to repeat: for QQP and MNLI (which require a large step budget) we have only 1–3 repetitions per training configuration, and for ConvNeXt-T (which takes a long time per step) we did repeat the training runs.

Each relative error difference (RED) score combines the error of two optimization methods (one being DOG) on a particular model task combination. Given multiple seeds for each optimization method, we computed the RED scores for each possible seed combination. In Figures 2, 3, 9 and 10 (which aggregate multiple tasks) we average over those multiple RED values and compute the statistics of the average RED. In per-task breakdowns such as Figure 5 and tables 4 and 5 we report the statistics over the multiple RED values.

Baseline optimizers. For both SGD and Adam, we used cosine learning rate decay, and searched over appropriate values for peak learning rate. The base learning rate search space used when performing fine-tuning for each model/task combination can be found in Tables 4 and 5. We did not use momentum for SGD. For Adam we used $\beta_1 = 0.9$ for all experiments, and $\beta_2 = 0.999$ for language experiments and $\beta_2 = 0.99$ for vision experiments. For language models only, we used warmup of 10% of the maximum steps count, and gradient clipping at global norm 1. We did not perform learning warmup or gradient clipping for the vision experiment since we did not encounter any training stability issues there.

Setting r_ϵ . As explained in Section 4.1, setting $r_\epsilon := \alpha(1 + \|x_0\|)$ generally works well for $\alpha = 10^{-4}$. However, in some cases such as with T5, $\|x_0\|$ can be very large, causing destructively large first updates, with η_t increasing exponentially and the model diverging. This is easily detectable early during training, as usually η_t exceeds 1000 within the first 100 steps. Since the theory requires r_ϵ to be small, we simply decreased α by a factor of 100. While preliminary experiments with RoBERTa indicated that DOG also performed well with $\alpha = 10^{-4}$, for the sake of consistency we use the same values in all models of the same domain. Thus, models fine-tuned on vision tasks used $\alpha = 10^{-4}$, while language models used 10^{-6} for DOG and 10^{-8} for L-DOG.

Model averaging. As mentioned in Section 4.1, we used the polynomial decay averaging as proposed by Shamir & Zhang (2013). Namely, we kept an additional copy of the model weights, and in every update step we updated our running average of the model parameters as follows:

$$\bar{x}_t^\gamma = \left(1 - \frac{1 + \gamma}{t + \gamma}\right) \bar{x}_{t-1}^\gamma + \frac{1 + \gamma}{t + \gamma} x_t \quad (12)$$

The vector \bar{x}_t^γ roughly corresponds to an average of the last t/γ iterates preceding iteration t . For all models, we set $\gamma = 8$. We did not perform any tuning of the parameter γ ; we chose the value 8 because $1/8$ seemed like a good fraction of iterates to average, and because it worked well in the experiments of (Levy et al., 2020).

To ensure that iterate averaging is never harmful, for each optimization method we selected the best-performing checkpoint across both x_t and \bar{x}_t^γ (i.e., with or without averaging).

E.6. Figure 1 details

We generated Figure 1 as part of our fine-tuning testbed. In particular, SGD used a cosine learning rate annealing (without warmup), both algorithms use polynomial decay averaging, and we report test performance on the best checkpoint selected on a validation set.

E.7. Fine-tuning ImageNet

Our training setup mostly followed the default configuration in Wortsman et al. (2022). In particular, we used batch size 512 and the default `timm` augmentation (as in our main computer vision experiments) which Wortsman et al. (2022) refer to as ‘medium aug.’ We trained for 25K steps, corresponding to roughly 10 passes over the data. However (keeping with

Table 3. CIFAR-10 test accuracies after training a Wide ResNet 28-10 model from scratch for 200 epochs, with and without polynomial decay averaging (see Section 4.6). \dagger denotes the standard training configuration (cf. [Cubuk et al., 2019](#), Table 2).

Algorithm	LR	Acc. w/o avg.	Acc. w/ avg.
SGD	0.1	94.9%	94.9%
	0.3	95.8%	95.6%
	1	96.4%	84.4%
	3	95.9%	21.7%
	10	10.0%	10.0%
SGD w/ mom. 0.9	0.01	95.0%	95.1%
	0.03	95.8%	95.7%
	0.1 \dagger	<u>96.3%</u>	88.5%
	0.3	95.8%	27.5%
	1	42.0%	63.4%
DoG	-	85.2%	96.4%
Adam	3e-05	91.1%	91.1%
	1e-04	94.0%	94.0%
	3e-04	93.5%	93.8%
	1e-03	91.4%	91.6%
L-DoG	-	83.2%	93.5%

our computer vision testbed setting) we did not perform learning rate warmup or gradient clipping, and we initialized the classification head to be random.

For AdamW ([Loshchilov & Hutter, 2019](#)) we used weight decay 0.1 and cosine learning rate annealing as in [Wortsman et al. \(2022\)](#). We obtained accuracies within 0.5% of the numbers reported in Appendix L of [Wortsman et al. \(2022\)](#).

DOG and L-DOG we used weight decay 0 since the value 0.1 is meant for decoupled weight decay and we did not wish to re-tune a weight decay parameter. We set r_ϵ to be $10^{-6} \cdot (1 + \|x_0\|)$ without trying different values of this parameter.

For SGD we used cosine learning rate annealing and set weight decay to 0 for a more direct comparison to DOG.

E.8. Training from scratch

Our training setup followed the basic training configuration of [Cubuk et al. \(2019\)](#), which is typical for training ResNets on CIFAR-10: data augmentations comprising a random crop after 4 pixel padding and random horizontal flip, batch size of 128, weight decay of 0.0005 and 200 epochs of training. SGD used cosine learning weight annealing and (when applicable) Nesterov momentum. We did not use dropout or any other additional form of regularization. For DOG and L-DOG, we set $r_\epsilon = 10^{-4} \cdot (1 + \|x_0\|)$ without trying different values of this parameter. The accuracies we obtained using SGD and DOG are consistent (and slightly better) than the baseline number reported in Table 2 of [Cubuk et al. \(2019\)](#) and within 0.1% of the one reported in Table 1 of [Carmon et al. \(2019\)](#).

F. Additional experiment results

F.1. Full breakdown of main experiment results

Figure 5 as well as Tables 4 and 5 provide the full breakdown of our main fine-tuning results, comparing DOG and L-DOG to SGD and Adam with different learning rates for each model/task combination.

F.2. Comparison with equalized compute budget

Throughout the paper, our experiments focus on comparing different methods by the final test performance they are able to reach given sufficient compute budget to essentially fully converge. As a consequence, SGD and Adam—which require learning rate tuning—use up significantly more computation than DOG and L-DOG. More specifically, for each model/task

Table 4. Average (std) performance of RoBERTa-b and T5-b on language tasks, when fine-tuned with different optimization algorithms and their respective base learning rate (when applicable). DOG uses $r_\epsilon = 10^{-6}(1 + \|x_0\|)$ and L-DOG uses $r_\epsilon = 10^{-8}(1 + \|x_0\|)$. Scores are reported as mean across seeds, measured in the corresponding performance metric as detailed in Table 2.

Model	Optimizer	LR	CoLA	MNLI	MRPC	QNLI	QQP	RTE	SQuAD	SST-2	STS-B	Avg.
RoBERTa-b	Adam	5e-06	60.8	87.8	89.8	93.0	88.7	77.6	90.3	95.1 (0.34)	90.4	85.46
		1e-05	63.4	87.9 (0.05)	90.5	93.1 (0.22)	89.0 (0.11)	77.6	91.5	95.1	90.8	86.22
		3e-05	63.5 (1.33)	87.3	92.3 (0.44)	92.9 (0.15)	88.8	80.6 (1.19)	92.3 (0.05)	94.8 (0.16)	91.1 (0.19)	86.94
		5e-05	61.8	86.8	92.0	92.3	88.0	78.7	92.4 (0.07)	94.3	90.9	85.93
		0.0001	57.5	86.2	91.8	91.3	0.0	79.4	91.9	94.7	89.8	75.33
	SGD	0.003	56.3	86.4	81.9	91.5	85.1	75.5	79.4	93.6	87.0	81.44
		0.01	59.1	87.4	89.7	92.5	87.0	79.0 (0.69)	86.2	94.8	90.3 (0.25)	84.81
		0.03	62.3 (1.38)	87.8 (0.04)	91.8 (0.21)	92.7 (0.18)	88.3	78.9 (0.79)	89.5 (0.12)	95.0 (0.15)	90.7 (0.12)	86.30
		0.1	58.7	87.4	91.0	92.2	88.7 (0.06)	78.3	91.0	94.0	90.4	85.52
		0.3	0.0	86.0	81.2	85.3	87.7	64.6	91.3 (0.11)	92.8	27.0	68.14
		1.0	0.0	81.5	81.2	83.8	79.4	52.7	82.8	89.7	13.0	62.22
T5-b	DOG	-	62.8 (1.17)	87.7 (0.12)	91.6 (0.29)	92.6 (0.15)	88.2 (0.02)	78.5 (2.91)	91.3 (0.17)	94.9 (0.26)	90.5 (0.33)	86.46
	L-DOG	-	63.3 (0.32)	87.7 (0.12)	91.5 (0.19)	92.8 (0.28)	88.7 (0.14)	80.1 (1.00)	91.8 (0.18)	94.8 (0.54)	90.6 (0.34)	86.81
	Adam	5e-06	53.4 (0.93)	86.8	91.4	93.4	88.0	79.8	90.3	93.9	90.4	84.82
		1e-05	56.0 (0.63)	86.9	91.2	93.5	88.2	80.5	90.4	94.2	90.6	85.33
		3e-05	58.9 (1.10)	87.1	91.6	93.4 (0.16)	88.8	82.3	90.8	94.8 (0.24)	90.7	86.12
		5e-05	58.9 (0.80)	87.3	91.8	93.3	89.0	80.9	90.7	94.8	90.8 (0.10)	85.97
		0.0001	58.3 (0.80)	86.9	92.9 (0.35)	93.5 (0.10)	89.2	82.5 (0.48)	90.9 (0.15)	94.9 (0.26)	90.8 (0.18)	86.53
	SGD	0.0005	55.4 (0.45)	86.1	92.3	92.7	88.8	81.2 (1.48)	90.0	94.6	89.7	85.29
		0.003	22.9 (1.64)	85.8	90.1	92.7	87.5	66.8	90.3	92.1	90.3	79.43
		0.01	49.4 (0.27)	86.4	92.2	93.1	87.4	80.9	90.3	93.0	90.5	84.49
		0.03	56.4 (0.52)	86.5	92.0	93.2 (0.03)	88.1	80.9	90.5	93.6	90.6	85.40
		0.1	58.9 (0.82)	86.8	91.8	93.1	88.7	84.1	90.7 (0.04)	93.7	90.6 (0.09)	86.13
		0.3	56.7 (0.83)	86.1	92.8 (0.47)	93.0 (0.13)	88.7	82.8 (1.24)	90.7 (0.05)	93.9 (0.15)	90.7 (0.21)	86.07
	DOG	-	7.3 (6.78)	86.9 (0.21)	92.8 (0.35)	93.1 (0.09)	88.5	81.7 (3.06)	90.6 (0.05)	94.1 (0.19)	90.7 (0.09)	80.58
	L-DOG	-	59.9 (1.43)	87.3 (0.10)	91.9 (0.32)	93.6 (0.02)	87.8	83.1 (0.78)	90.3 (0.02)	95.0 (0.19)	90.5 (0.05)	86.51

combination we tune the SGD and Adam learning rates over a grid of at least 5 values (and often 6 or more), resulting in computational cost increased by the same factor.

In this subsection only, we compare different optimizers using roughly the same computational budget, by measuring the performance of Adam and SGD after roughly 20% of their overall step budget.¹¹ Figure 7 shows the result of this comparison, contrasting it to our main experiment. The figure shows that DOG often exceeds the performance of instance-tuned SGD with equalized step budget.

We note a number caveats regarding the equalized compute budget comparison:

1. Since our experiments are focused on getting the best possible generalization, we substantially over-provisioned the iteration budget, and hence the performance of instance-tuned SGD and Adam declines only mildly when we cut the budget by roughly 5. Our tightest budget was for RoBERTa (150% the iterations in Liu et al. (2019)), and there we can see that performance degraded more substantially. If we were to instead take the number of iterations DOG actually needs to reach its peak performance, its advantage over equalized-compute SGD would likely be far larger.
2. Since the results reported here are obtained by re-analysis of our original experiments, the cosine learning rate schedule for SGD and Adam is not properly set for using 20% of the iteration budget; in particular, the learning rate does not decay to zero at the end of the training. Running these methods with a properly set cosine schedule would likely improve their performance. However, we noted that the addition of iterate averaging appears to partially compensate for the lack of sharp learning rate decay.
3. Given sufficient resources, it is possible to run all the different learning rates of SGD and Adam in parallel. Therefore, the comparison equalizes overall computational cost, but not necessarily wall time.
4. The comparison also does not take into account more sophisticated learning rate tuning schemes that early-stop unpromising learning rate candidates. However, that such schemes run risk choosing a suboptimal learning rate.

¹¹Since these results are just a re-analysis of our original experiments, for language experiments we take all the warmup iterates plus the first 20% of the remaining iterates, overall using 28% of the budget.

F.3. Fine-tuning CoLA

As discussed in Section 4.2, DOG with $r_\epsilon = 10^{-6}(1 + \|x_0\|)$ failed in fine-tuning T5-b on CoLA. To investigate this issue, we ran DOG and L-DOG with different choices of r_ϵ . Figure 11 depicts the results of this test as well as the performance of SGD and Adam with different learning rates. The figure shows that using lower values of r_ϵ allows DOG to reach reasonable results, but with some seeds still failing. In contrast, L-DOG shows consistent and superior performance across a large range of r_ϵ values. We leave further investigations on the cause of failure in CoLA to future work.

F.4. Sensitivity of DOG to r_ϵ and the effect of batch normalization

In Section 4.3, we discuss DOG’s insensitivity to the choice of r_ϵ as long as it is small enough. Here, we expand on this analysis by testing how the DOG step size at iteration t , denoted η_t , depends on its initial step size $\eta_0 = r_\epsilon/\|g_0\|$. For each task and in our testbed and 4 models, we perform short training runs with a large numbers of η_0 values. In Figure 6 we plot η_t vs. η_0 for $t \in \{2, 10, 100, 1000\}$. We also show a horizontal line for the best peak step size of SGD, and the $y = x$ diagonal. The figure shows that for most model/task combinations, η_t converges quickly (within the first 1000 steps) to a value near the optimal one for SGD, and mostly independent of η_0 as long as it is small enough.

However, we also observe some failure cases where η_t strongly depends on η_0 , such as fine-tuning ResNet50 on CIFAR-100. This provides a complementary perspective on the fact DOG is sensitive to r_ϵ in this setting, as already shown in Figure 3: when η_0 is too low, DOG fails to reach a suitable value of η_t in a reasonable time. We hypothesize that this is due to the batch normalization (BN) layers in the model causing many different step size to “look” like solutions to the implicit equation motivating DOG. To test this hypothesis, we repeat the CIFAR-100 training experiment but without BN (we disable BN by fine-tuning the model in evaluation mode). Figure 8(a) shows that removing BN allows DOG to recover its stabilizing behavior. Moreover, Figure 8(b) further shows that without batch normalization, the performance of DOG again becomes insensitive to the choice of r_ϵ provided it is sufficiently small. Unsurprisingly, we also observe that removing BN slightly hurts generalization performance in this task. As mentioned in Section 6, improving DOG to be more robust in the presence of normalization layers in general and batch normalization in particular is an important direction for future research.

F.5. Additional convex optimization results

Figures 9 and 10 show results for learning lines probes on our computer vision fine-tuning testbed (see Section 4.4). The figures show that DOG attains results on par with instance-tuned SGD and Adam.

F.6. The growth rate of \bar{r}_t

Figure 12 plots \bar{r}_t for DOG as a function of the iteration index t . As the figure shows, \bar{r}_t grows very rapidly and then approximately plateaus. Therefore, the quantity $\sum_{i \leq t} \frac{\bar{r}_i}{\bar{r}_t}$ grows roughly linearly in t , implying a near-optimal rate of convergence for DOG, as discussed in Section 3.2.

G. Comparison to Other Tuning-Free Methods

Section 4.7 discusses a comparison between DOG and other parameter-free optimizers. In this section, we provide further details on the experiments.

G.1. Parameter-free SGD

Carmon & Hinder (2022) propose a bisection procedure for tuning the SGD learning rate. A direct implementation of this method would need to perform at least 4 or 5 bisection steps and therefore, *in the best case*, perform similarly to our instance-tuned SGD baseline. Since our learning rate tuning employs a tight grid of values selected using some prior knowledge of the problem, and since we select learning rates based on validation set performance and not a step size certificate, instance-tuned SGD is likely a conservative upper bound on the performance of bisection approach.

Similar to instance tuned SGD, the bisection procedure has increased computational cost relative to DOG that is due to the need for multiple SGD runs. That is, performing 5 steps of bisection where each SGD call has the same step budget as DOG consumes 5 times more compute than DOG. We may also consider a situation where each bisection step uses only 20% of the DOG compute budget, leading to equal overall cost. In this setting, the “equalized compute budget” comparison

we perform in Appendix F.2 and Figure 7 provides a conservative upper bound on the bisection performance, indicating it is likely to under-perform DOG.

G.2. Stochastic Polyak step-size

We apply the Stochastic Polyak Step (SPS) proposed by Loizou et al. (2021) using their open-source implementation¹² to a subset of our fine-tuning testbed, and present the results in Tables 6 and 7. For the vision experiments, the SPS with the hyper-parameters proposed in the paper ($c = 0.2$, $\tau = 2$) and initial step size of 1.0 (the default in the code) worked reasonably well, but not as well as DoG. For the language experiments the same algorithm diverges; we find initial learning rate of 0.01 worked reasonably well, but again not as well as DoG (we also attempted an initial learning rate of 0.0001, which produced worse results). For vision tasks, similarly tuning the initial step size did not significantly improve performance. We run 5 random seeds per experiment, and average the results across seeds. DOG outperforms SPS in 22 out of 34 task/model combinations, and by 2.3 percentage points on average. L-DOG further increases this gap by outperforming SPS in 24 out of 34 pairs, with an average of 5.3 percentage points.

G.3. D-adaptation

Empirical comparison to D-adapt SGD and Adam. We perform a preliminary empirical evaluation of the practical algorithms proposed in Defazio & Mishchenko (2023) using the code they release¹³ and a subset of our fine-tuning testbed. As Tables 6 and 7 show, D-adapt SGD and D-adapt Adam perform reasonably well but slightly worse than DOG, and noticeably worse than L-DOG and Adam. DOG outperforms D-adapt SGD in 24 out of 34 task/model combinations, and by 1.9 percentage points on average. L-DOG further increases this gap by outperforming D-adapt SGD in 30 out of 34 pairs, with an average of 4.9 percentage points. D-adapt Adam is less stable on many of the tasks in our testbed, being outperformed by DOG in 26 out of 34 task/model combinations, and by L-DOG in 27, with an average of 10.8 and 13.8 percentage points respectively.

Theoretical comparison to Algorithm 1 of Defazio & Mishchenko (2023). Defazio & Mishchenko (2023) carry out their main theoretical analysis on a “Parameter Free Dual Averaging” (PFDA) method. We now provide some additional remarks comparing PFDA and DOG. The iterate x_t in PFDA is

$$x_t = x_0 - \frac{1}{\sqrt{G_t}} \sum_{i \leq t} q_i g_i$$

where $G_t = \sum_{i \leq t} \|g_i\|^2$ and q_i is a lower bound on the distance to optimality (denoted d_i in (Defazio & Mishchenko, 2023)). In contrast, the DOG iterates are

$$x_t = x_0 - \sum_{i \leq t} \frac{\bar{r}_i}{\sqrt{G_i}} g_i$$

where $\bar{r}_t = \max_{i \leq t} \|x_i - x_0\|$. While both d_t in PFDA and \bar{r}_t are lower bounds for (a constant factor times) the distance to optimality, only DoG aims to approximate $\eta_\star = \frac{\|x_0 - x_\star\|}{\sqrt{G_t}}$; PFDA instead approximates the optimal step size for dual averaging.

The dual averaging prescription of putting the factor $1/\sqrt{G_t}$ outside the summation defining x_t likely hurts performance in practice. The practical D-adapt SGD and D-adapt Adam methods that Defazio & Mishchenko (2023) propose do not follow this prescription. Consequently, these algorithms are very different from PFDA and have no theoretical guarantees.

¹²<https://github.com/IssamLaradji/sps>

¹³<https://github.com/facebookresearch/dadaptation>

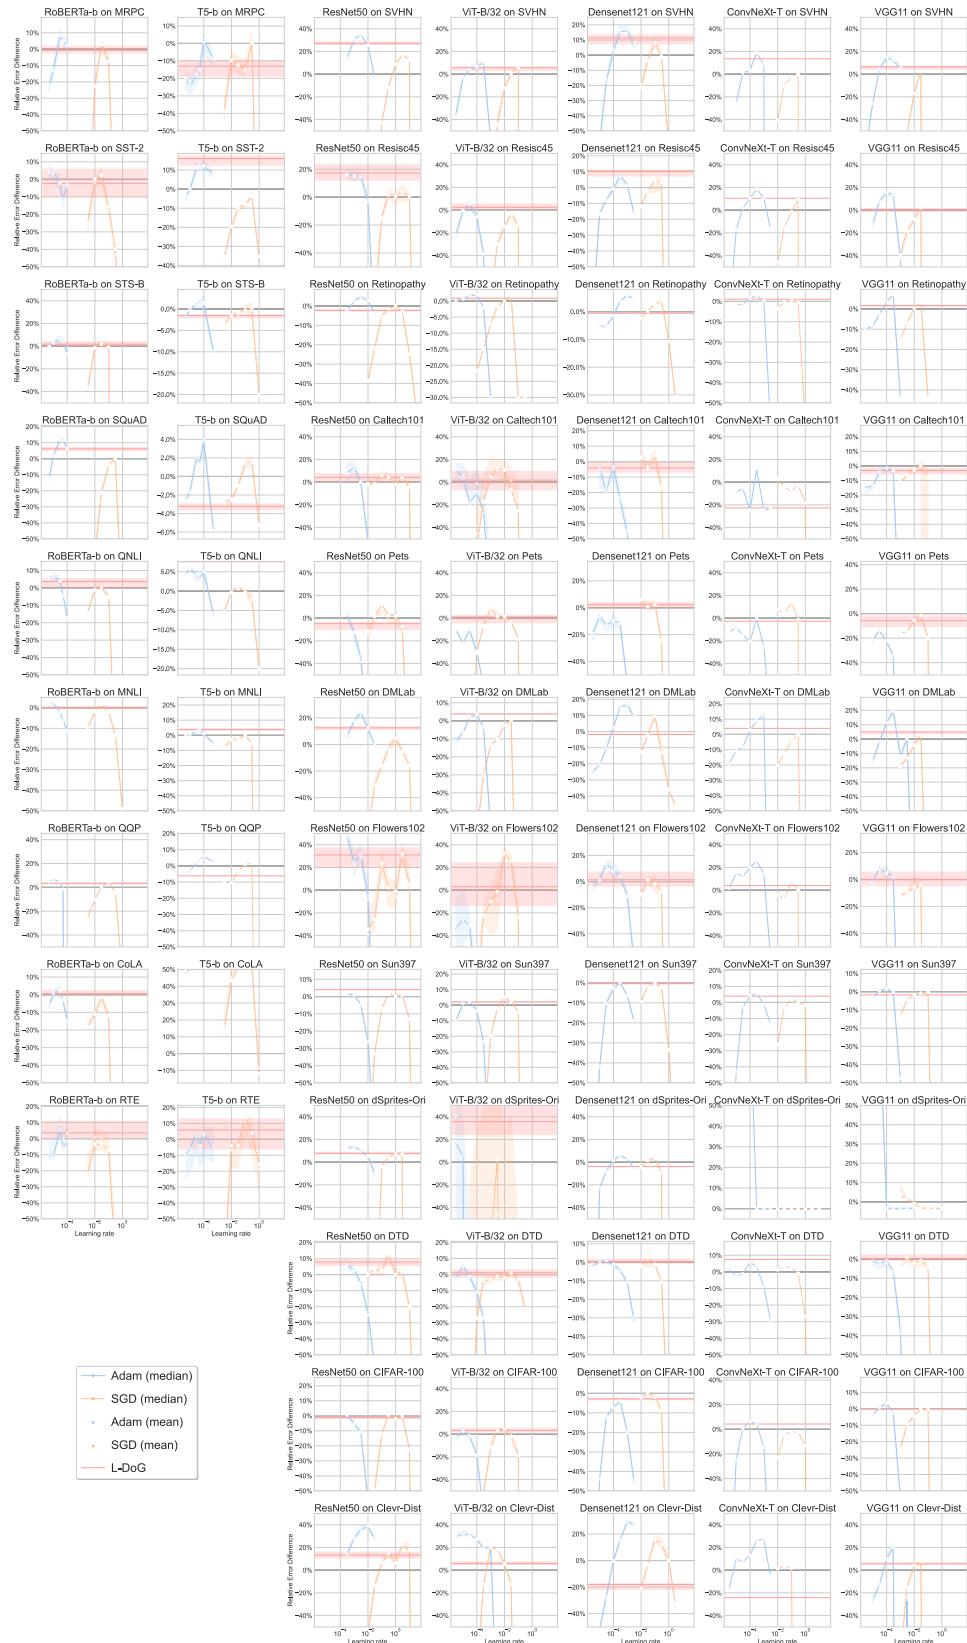


Figure 5. Relative error difference (RED) statistics across seeds (median, mean and IQR shown as shaded region) for all model/task combinations. The red horizontal line shows the median RED of L-DOG.

Table 5. Average (std) test accuracy across seeds for vision tasks, when fine-tuned with different optimization algorithms and their respective base learning rate when applicable. DOG and L-DOG use $r_\epsilon = 10^{-4}(1 + \|x_0\|)$.

Model	Optimizer	LR	Caltech101	CIFAR-100	Clevr-Dist	DMLab	dSprites-Ori	DTD	Flowers102	Pets	Resisc45	Retinopathy	Sun397	SVHN	Avg.
Adam	3e-06	-	70.0	89.2	70.5	86.5	72.1	92.5	93.1	91.2	-	36.7	-	-	-
	1e-05	89.0	84.7	91.7	72.7	95.7	70.9	93.5	93.3	94.7	83.2	64.2	96.6	85.33	
	3e-05	89.4	87.8	91.4	73.2	96.3	71.3	93.3	92.9	95.6	83.2	74.3	97.3	86.75	
	0.0001	87.5	88.5	91.9	76.0	96.4	73.3	93.9	92.6	95.9	83.7	76.0	97.4	87.25	
	0.0003	91.1	88.2	93.2	77.5	7.6	72.5	94.3	93.8	96.3	83.9	76.3	97.8	80.58	
	0.001	87.5	85.9	93.2	78.5	7.6	69.1	93.4	92.3	95.9	83.5	74.7	97.5	79.42	
ConvNeXt-T	0.003	87.6	73.7	90.2	22.2	7.6	63.4	87.9	88.3	94.7	73.6	71.7	19.6	64.50	
	0.01	90.2	85.0	90.8	70.3	7.6	72.0	92.0	94.3	93.4	82.9	68.4	96.2	78.25	
	0.03	89.5	87.2	91.1	72.2	7.6	72.6	91.8	94.1	94.8	83.4	74.2	97.0	79.25	
	0.1	89.1	87.5	90.9	74.3	7.6	72.3	92.9	94.7	95.4	83.6	75.4	97.2	79.58	
	SGD	0.3	89.7	87.4	24.5	75.2	7.6	71.4	92.4	93.8	95.9	83.4	75.1	97.3	74.00
	1.0	88.1	86.1	24.5	22.2	7.6	64.0	0.4	13.4	2.1	73.6	74.5	19.6	39.33	
DOG	3.0	0.4	1.0	20.0	22.2	7.4	2.1	0.3	2.7	2.2	73.6	0.5	6.7	11.25	
	10.0	-	1.0	20.0	22.2	7.4	2.1	0.3	2.7	2.2	-	0.5	-	-	
	-	89.9	87.7	90.7	75.3	7.6	71.6	92.4	93.8	95.5	83.5	75.0	97.3	79.50	
	L-DOG	-	87.7	88.2	88.5	76.3	96.4	73.8	92.7	93.7	95.9	83.7	75.9	97.7	86.92
	3e-06	-	62.3	84.3	65.0	84.8	65.9	88.1	88.4	90.6	-	37.9	-	-	
	1e-05	86.5 (1.24)	76.9	86.9	66.3	95.4	66.4	88.6	89.7	93.9	81.1	59.6	95.6	81.71	
SGD	3e-05	85.2	82.0	89.0	69.0	95.9	66.4 (0.76)	90.0	89.0	94.4	81.0	68.7	96.8	83.70	
	0.0001	86.7 (1.40)	82.8 (0.26)	91.4 (0.16)	72.7 (0.69)	96.3 (0.07)	65.8 (0.39)	89.4 (1.08)	89.4 (0.57)	94.8 (0.16)	81.7 (0.10)	70.8 (0.45)	97.3 (0.05)	84.92	
	0.0003	84.5	83.3	92.7	76.3	96.4 (0.03)	65.1	88.9	89.2	95.1 (0.20)	82.7	71.6	97.7 (0.08)	84.93	
	0.001	81.8	80.8	93.9	76.6 (0.39)	96.4	62.8	87.2	84.7	94.9	83.0 (0.08)	69.7	97.7	83.55	
	0.003	76.0	76.7	93.7 (0.17)	74.6	96.0	56.1	81.5	82.7	93.9	82.8	66.2	97.4	81.06	
	0.01	87.7 (0.74)	83.6	89.5	68.5	96.1	65.7	87.4	90.9	94.2	81.6	68.9	96.7	83.73	
DenseNet121	0.03	87.0	84.0 (0.08)	91.1	71.4	96.3 (0.07)	66.6 (1.59)	88.6	90.5 (0.47)	94.7	82.0	71.1	97.1	84.87	
	SGD	0.1	87.9 (0.57)	83.7 (0.31)	92.8 (0.26)	74.5 (0.53)	96.4 (0.05)	65.9 (0.49)	88.3 (0.82)	90.6 (0.08)	94.9 (0.36)	82.5 (0.08)	71.5 (0.21)	97.4 (0.09)	85.53
	0.3	85.5	82.5	92.4 (0.56)	69.1 (3.67)	95.9	62.9	87.6	87.9	95.0 (0.25)	82.6 (0.28)	70.9	97.2	83.67	
	1.0	61.6	65.1	91.4	61.9	7.6	35.4	55.4	64.6	82.6	80.0	62.0	95.6	63.17	
	3.0	50.1	61.6	88.5	59.2	7.6	23.4	44.1	39.8	85.8	76.5	45.6	94.7	55.92	
	DOG	-	87.4 (0.65)	84.0 (0.18)	91.4 (0.19)	71.9 (0.43)	96.2 (0.04)	66.1 (0.90)	88.5 (0.92)	90.3 (0.04)	94.8 (0.12)	82.0 (0.08)	71.6 (0.22)	97.2 (0.13)	85.12
L-DOG	-	86.9 (0.27)	83.5 (0.18)	89.8 (0.31)	71.5 (0.24)	96.1 (0.03)	66.4 (0.61)	88.7 (0.70)	90.6 (0.13)	95.3 (0.17)	81.9 (0.10)	71.4 (0.25)	97.5 (0.09)	84.97	
	0.0003	87.8 (1.52)	84.8	91.0	73.3	96.2	68.5 (0.97)	92.7	93.1 (0.27)	95.5	82.2	73.9	97.0	86.03	
	0.001	88.3 (1.18)	83.9 (0.31)	92.4 (0.21)	76.5 (0.53)	96.3 (0.04)	67.2 (0.94)	89.2 (1.83)	92.0 (0.37)	95.5 (0.25)	83.0 (0.18)	73.7 (0.33)	97.6 (0.07)	86.30	
	Adam	0.003	86.9	83.1	93.3 (0.37)	78.3 (0.30)	96.1	64.4	90.1	90.2	95.4	83.4 (0.15)	72.2	97.7	85.67
	0.01	79.9	75.9	93.3	75.0	95.9	58.0	85.1	85.0	94.4	83.0	66.5	97.4	82.08	
	0.03	62.8	64.9	92.4	71.3	95.3	46.9	63.1	67.0	89.3	82.0	50.8	96.4	73.08	
ResNet50	0.01	86.7	62.9	83.7	52.5	68.3	66.2	80.9	91.9	84.0	75.9	48.0	80.2	72.92	
	0.03	86.7	77.0	88.0	63.0	88.5	67.2	82.1	92.8	90.5	78.8	64.7	91.4	80.50	
	SGD	0.1	87.6 (0.66)	82.8	90.2	67.0	95.5	67.5 (1.71)	89.2	93.5	93.8	81.7	71.6	95.0	84.26
	0.3	87.4	84.8	91.1	70.7	95.9	70.5	85.8 (2.54)	92.9	94.7	82.3	73.9	96.1	84.98	
	1.0	86.6 (0.50)	84.5 (0.46)	90.1 (0.51)	72.6 (1.96)	96.0 (0.08)	66.4 (1.16)	85.5 (2.39)	93.0 (0.30)	94.6 (0.28)	82.6 (0.09)	73.6 (0.44)	96.8 (0.06)	85.19	
	3.0	87.4	84.7	91.8	70.1	96.0	66.9	90.3	92.1	94.8 (0.49)	82.1	73.7	97.0 (0.10)	85.23	
VGG11	10.0	86.2	81.2	91.4 (0.71)	67.2	7.6	59.4	86.6	82.2	94.6	78.2	70.0	96.9	74.78	
	DOG	-	86.8 (0.62)	84.8 (0.37)	89.3 (0.58)	71.4 (0.77)	95.7 (0.09)	66.4 (1.48)	85.8 (2.64)	92.9 (0.38)	94.6 (0.33)	82.6 (0.16)	73.5 (0.41)	96.5 (0.10)	85.02
	L-DOG	-	87.6 (1.15)	84.6 (0.38)	90.8 (0.38)	75.0 (0.44)	96.0 (0.05)	69.1 (1.30)	90.2 (2.02)	92.4 (0.36)	95.6 (0.46)	82.2 (0.15)	74.6 (0.32)	97.4 (0.08)	86.29
	3e-06	80.2 (0.71)	-	-	-	-	-	-	-	-	79.8 (0.07)	-	93.8 (0.15)	-	
	1e-05	80.5	73.7	89.3	63.5	94.3	61.5	82.1	87.3	91.4	80.0	65.4	95.3	80.00	
	3e-05	82.0 (0.48)	74.9	90.5	67.7	96.1 (0.06)	61.4 (0.84)	84.0	88.1 (0.17)	92.1 (0.08)	91.1	66.6	96.4	81.48	
Adam	0.0001	82.6	75.6 (0.23)	92.4	71.9	7.6	61.6	83.5 (1.06)	87.2	93.5 (0.23)	82.3	66.9 (0.12)	96.8	74.79	
	0.0003	82.3	74.1	93.2 (0.25)	74.4 (0.47)	7.5	59.6	83.0	86.0	93.4	82.9 (0.08)	66.3	96.8 (0.14)	74.77	
	0.001	61.7	61.1	24.5	64.6	7.5	48.0	53.7	65.0	89.1	73.6	50.5	96.5	57.58	
	0.003	-	1.0	89.4	68.4	7.6	2.1	0.5	2.7	76.2	-	30.5	-	-	
	0.01	-	1.0	24.5	22.2	7.6	2.1	1.2	2.7	2.2	-	2.0	-	-	
	0.03	-	81.0	68.5	85.0	62.3	14.8 (3.59)	61.3	80.3	88.9	78.9	61.9	92.0	71.65	
SGD	0.003	81.8	72.4	90.3	64.2	12.3	62.2	80.9	88.0	90.9	80.3	64.9	94.4	73.08	
	0.01	82.4	73.5	91.9 (0.19)	67.0	9.8	61.2	82.0 (0.99)	89.0 (0.37)	91.7	81.7	65.8	95.7	73.91	
	0.03	83.0 (0.50)	74.7 (0.21)	92.1 (0.19)	69.1 (0.36)	7.6 (0.04)	61.9 (0.41)	82.3 (1.09)	89.4 (0.32)	92.5 (0.28)	82.2 (0.06)	66.1 (0.10)	96.4 (0.08)	74.77	
	0.1	49.6 (4.49)	74.6 (0.08)	20.0	22.2	7.6	60.5	0.3	87.5	2.2	73.6	53.2 (29.49)	6.7	37.87	
	0.3	-	1.0	20.0	22.2	7.4	2.1	0.3	2.7	2.2	-	0.5	-	-	
	1.0	-	1.0	20.0	22.2	7.4	2.1	0.3	2.7	2.2	-	0.5	-	-	
VIT-B/32	DOG	-	82.9 (0.45)	74.7 (0.22)	91.5 (0.17)	68.4 (0.53)	10.4 (0.02)	62.5 (1.19)	82.8 (0.95)	89.5 (0.23)	92.4 (0.22)	81.6 (0.07)	66.3 (0.33)	96.3 (0.09)	74.94
	L-DOG	-	82.4 (0.60)	74.8 (0.16)	92.0 (0.11)	69.9 (0.44)	92.1 (8.59)	62.6 (0.95)	82.7 (1.15)	88.9 (0.62)	92.4 (0.15)	81.9 (0.18)	65.8 (0.09)	96.5 (0.12)	81.83
	3e-06	90.7 (0.76)	92.3 (0.22)	89.5 (0.35)	66.0 (0.69)	94.0 (0.24)	74.9 (0.24)	98.5 (0.53)	91.6 (0.11)	95.5 (0.07)	79.7 (0.08)	75.5 (0.23)	96.8 (0.06)	87.08	
	1e-05	92.8 (0.21)	89.9 (0.52)	67.5	51.5 (48.10)	76.9	98.7 (0.29)	90.9	96.3	79.8	78.0 (0.12)	97.6	83.84		
	3e-05	88.6	92.5	89.7	70.1	7.6	75.3 (0.24)	98.7	91.6 (0.21)	96.5 (0.13)	80.1 (0.09)	78.3	97.7	80.21	
	Adam	0.0001	89.5	91.2	89.1	70.8 (0.32)	7.6	72.9	98.1	90.0	96.1	80.1	77.0	97.8 (0.07)	79.80
SGD	0.0003	87.9	87.3	87.9	68.5	7.6	69.0	94.9	87.9	94.9	79.3	72.5	97.9	77.33	
	0.001	80.3	62.1	88.1	51.2	7.6	50.4	71.0	58.3	88.6	73.6	53.6	96.3	64.75	
	0.003	13.5	64.6	29.3	7.6	14.1	26.7	12.1	71.9	-	19.5	-	-	-	
	3e-05	-	6.1	52.7	40.8	32.7	45.9	61.4	80.3	59.7	-	5.2	-	-	
	0.0001	85.0	73.7	71.3	50.4	57.0	69.0	98.1	90						

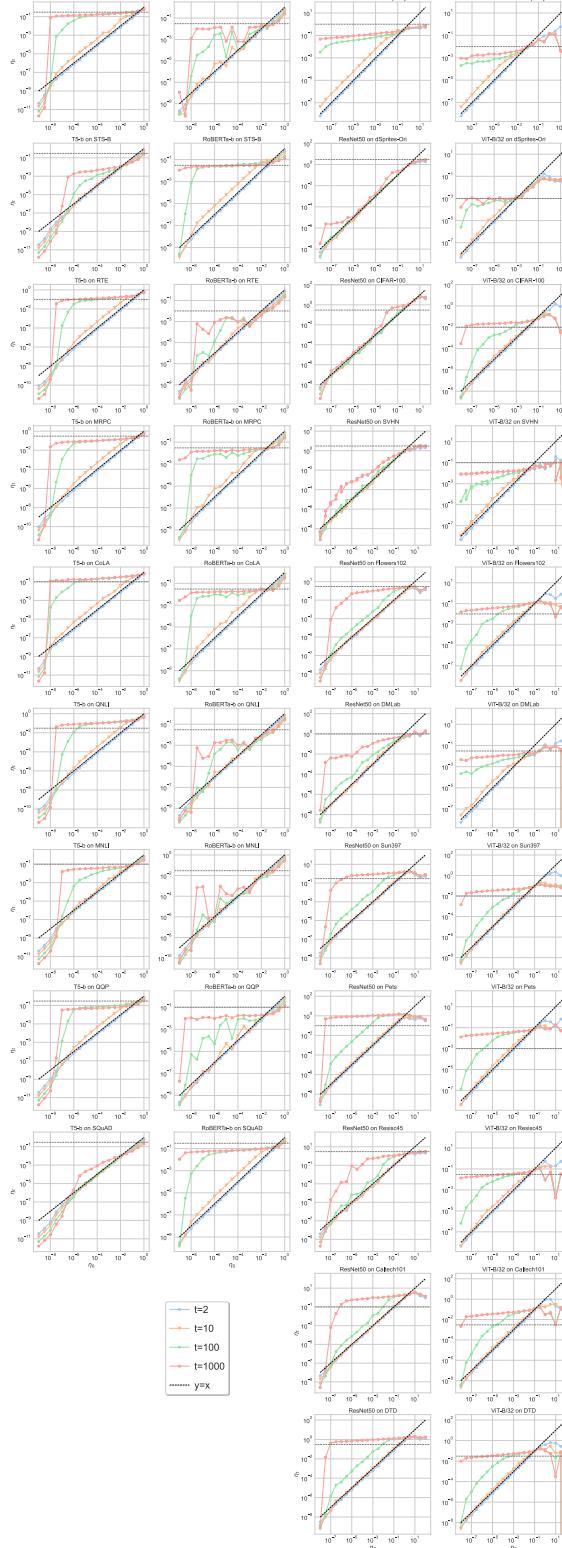


Figure 6. Stabilizing behavior of DOG on η_t as a function of η_0 (x-axis) and t (color). In most cases η_t quickly stabilizes around a value close to the optimal SGD base learning rate (dashed horizontal line) for all sufficiently small $\eta_0 = r_\epsilon / \|g_0\|$. The main exceptions (where η_t depends strongly on η_0) are dSprites-Ori, CIFAR-100 and SVHN when trained with ResNet50; see F.4 for further discussion.

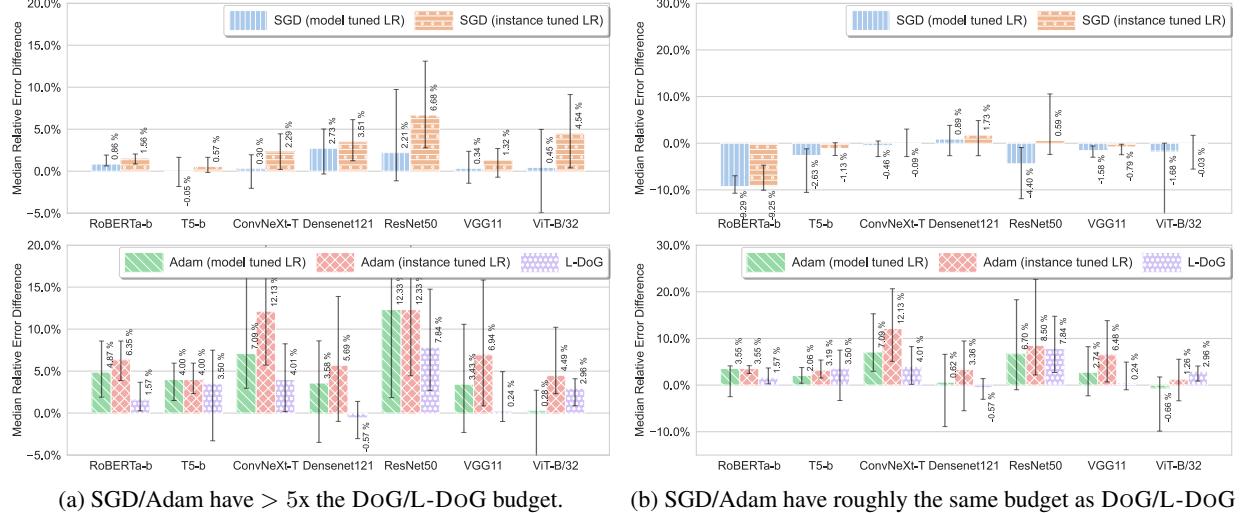


Figure 7. RED median (bar chart) and IQR (error bars) of each model on the set of applicable tasks, where we either (a) give the same iteration budget for each optimizer run, resulting in SGD and Adam using more than 5x total compute than DOG and L-DoG due to learning rate tuning (this is a reproduction of Figure 3), or (b) give each algorithm *roughly equal compute budget* by running SGD and Adam (with 5 or more learning rates) for roughly 20% of the steps that DOG and L-DoG use. With equalized compute budget, DOG outperforms model-tuned SGD almost always and often outperforms the instance-tuned SGD as well, while L-DoG closes most of the gap to Adam.

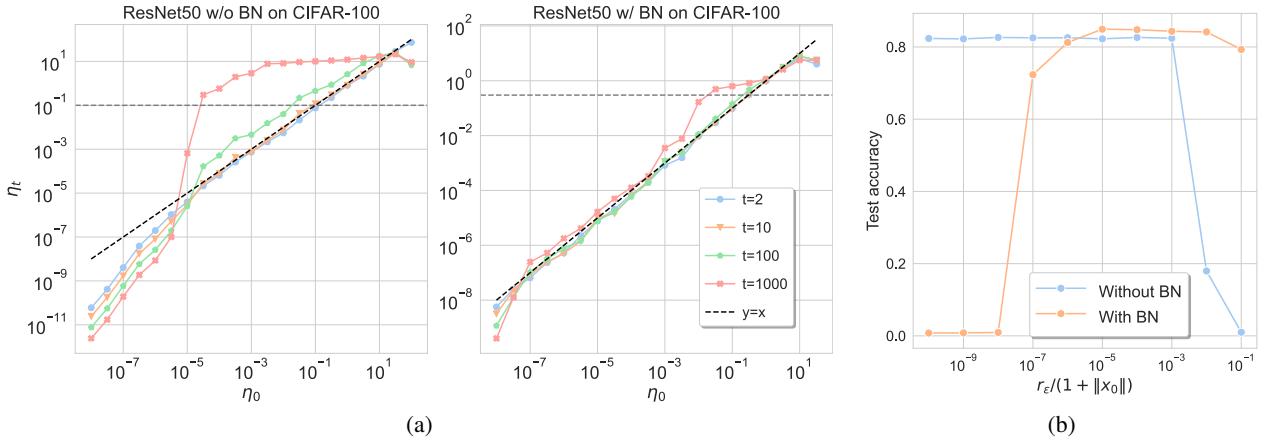


Figure 8. ResNet50 fine-tuned on CIFAR-100 with and without batch normalization. The dashed horizontal line indicates the best SGD learning rate. (a) Stabilizing behavior of DOG on η_t as a function of η_0 (x-axis) and t (color). Turning off batch normalization (left) mitigates the sensitivity of η_t to η_0 observed in batch normalized model (right). (b) Accuracies of models trained with DOG (for 20K steps) as a function of r_ϵ . Without batch normalization, DOG is robust to smaller values of r_ϵ .

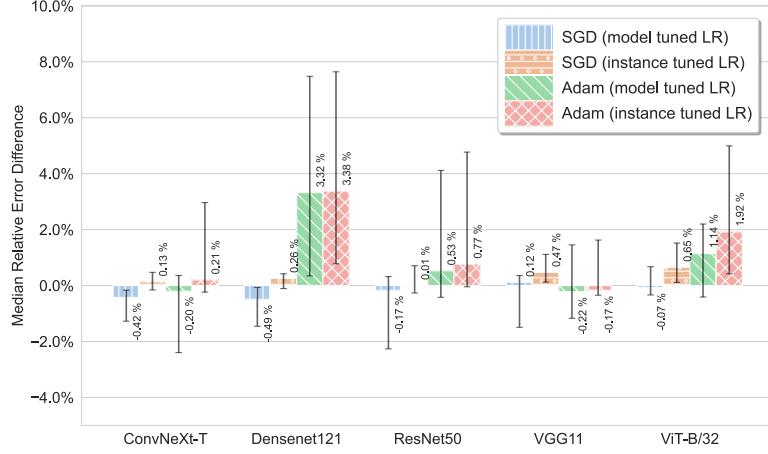


Figure 9. RED median and IQR (as in Figure 3) in the *convex optimization* setting (Section 4.4).

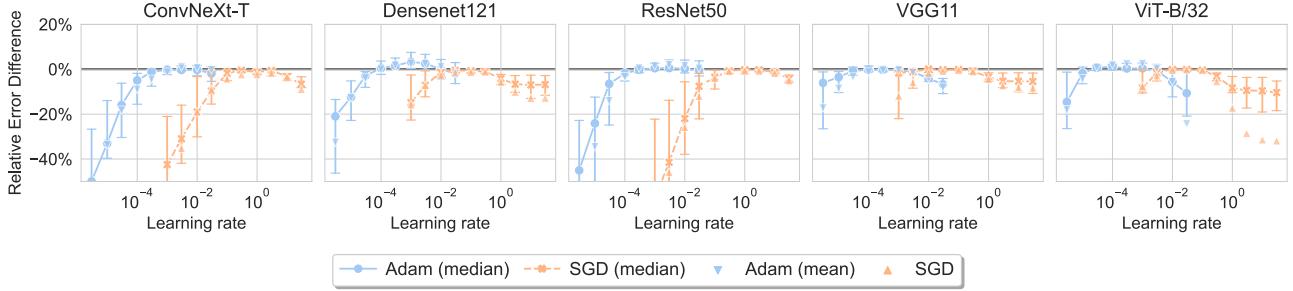


Figure 10. Per-learning rate RED statistics (as in Figure 2) in the *convex optimization* setting (Section 4.4).

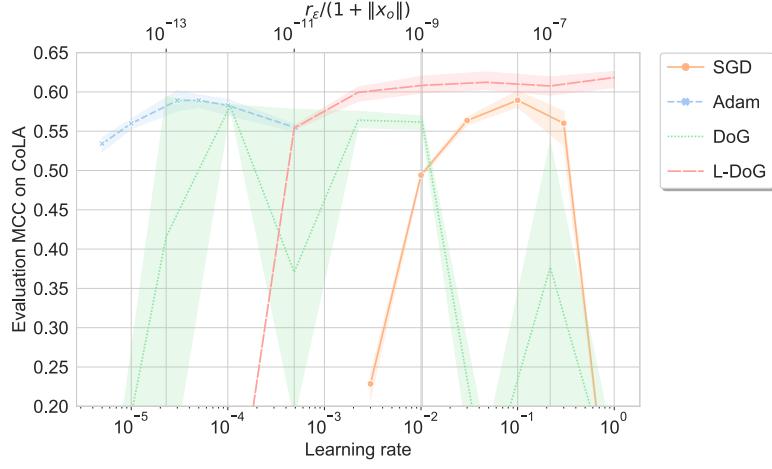


Figure 11. Matthews correlation of T5-base fine-tuned on CoLA with SGD and Adam with different base learning rates (bottom axis), as well as with DoG and L-DOG with different r_ϵ (top axis). Only L-DOG and Adam perform consistently well across different parameters. The lines and shaded regions show the average Matthews correlation and the min-max range, respectively, computed over 3 seeds.

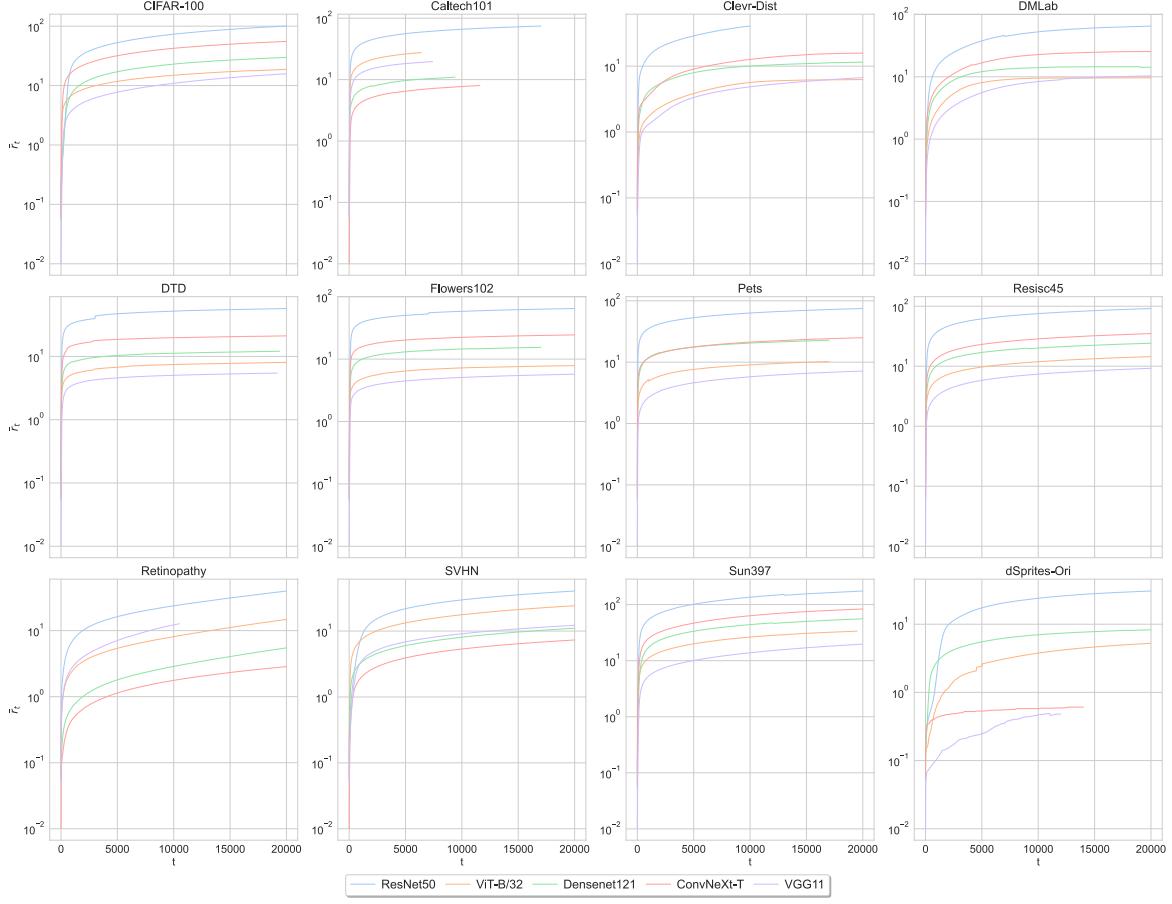


Figure 12. The quantity $\bar{r}_t = \max_{i \leq t} \|x_i - x_0\|$ as a function of the number of steps t in our computer vision testbed. The value of \bar{r}_t grows rapidly at first and then almost plateaus.

Table 6. Average (std) performance of RoBERTa-b and T5-b on language tasks, when fine-tuned with different optimization algorithms. DOG uses $r_\epsilon = 10^{-6}(1 + \|x_0\|)$ and L-DOG uses $r_\epsilon = 10^{-8}(1 + \|x_0\|)$. SPS uses $c = 0.2, \tau = 2$ as recommended by Loizou et al. (2021), but initial step size of 0.01 as the recommended value 1.0 diverged for some tasks. Still, When tuning RoBERTa-b on RTE, 3 out of 5 diverged; for this case we report the mean of the two successful and omit the standard deviation. We measure performance as detailed in Table 2.

Model	Optimizer	CoLA	MRPC	QNLI	RTE	SQuAD	SST-2	Avg.
RoBERTa-b	D-Adapt (Adam)	0.0 (0.00)	83.4 (4.93)	66.7 (22.13)	68.4 (14.40)	7.0 (0.16)	58.2 (17.88)	47.28
	D-Adapt (SGD)	49.4 (27.59)	91.6 (0.43)	91.9 (0.67)	81.4 (1.04)	91.5 (0.13)	94.1 (0.38)	83.32
	SPS	49.8 (6.79)	88.4 (3.21)	91.8 (0.21)	52.7	90.1 (0.10)	94.1 (0.28)	77.70
	DoG	62.8 (1.17)	91.6 (0.29)	92.6 (0.15)	78.5 (2.91)	91.3 (0.17)	94.9 (0.26)	85.28
	L-DOG	63.3 (0.32)	91.5 (0.19)	92.8 (0.28)	80.1 (1.00)	91.8 (0.18)	94.8 (0.54)	85.72
T5-b	D-Adapt (Adam)	11.1 (2.51)	81.8 (0.65)	85.0 (1.91)	58.1 (1.35)	90.4 (0.06)	87.1 (3.00)	68.92
	D-Adapt (SGD)	0.0 (0.00)	92.5 (0.59)	92.9 (0.04)	81.2 (0.78)	90.4 (0.06)	80.3 (2.67)	72.88
	SPS	39.1 (2.35)	92.9 (1.21)	93.2 (0.10)	80.9 (1.82)	90.5 (0.04)	94.3 (0.31)	81.82
	DoG	7.3 (6.78)	92.8 (0.35)	93.1 (0.09)	81.7 (3.06)	90.6 (0.05)	94.1 (0.19)	76.60
	L-DOG	59.9 (1.43)	91.9 (0.32)	93.6 (0.02)	83.1 (0.78)	90.3 (0.02)	95.0 (0.19)	85.63

Table 7. Average (std) test accuracy across seeds for vision tasks, when fine-tuned with different optimization algorithms. DOG and L-DOG use $r_\epsilon = 10^{-4}(1 + \|x_0\|)$. SPS uses $c = 0.2, \tau = 2$ and initial step size of 1 as recommended by Loizou et al. (2021).

Model	Optimizer	Caltech101	CIFAR-100	Clevr-Dist	DMLab	dSprites-Ori	DTD	Flowers102	Pets	Resisc45	Retinopathy	Sun397	SVHN	Avg.
ResNet50	D-Adapt (Adam)	79.0 (1.60)	82.9 (0.23)	92.3 (0.50)	75.2 (0.97)	96.3 (0.01)	60.1 (0.51)	84.6 (1.73)	85.1 (0.95)	94.1 (0.28)	83.4 (0.07)	72.9 (0.11)	97.0 (0.21)	83.58
	D-Adapt (SGD)	86.2 (0.93)	84.3 (0.15)	89.6 (0.60)	74.8 (0.73)	95.9 (0.04)	64.4 (1.31)	84.9 (1.12)	92.0 (0.25)	94.7 (0.17)	82.7 (0.12)	73.0 (0.09)	97.1 (0.06)	84.97
	SPS	86.4 (1.14)	80.2 (1.20)	92.9 (0.25)	76.2 (0.60)	96.0 (0.05)	63.7 (0.83)	84.3 (1.45)	92.4 (0.19)	94.9 (0.22)	83.3 (0.17)	70.9 (0.82)	97.5 (0.03)	84.89
	DOG	86.8 (0.62)	84.8 (0.37)	89.3 (0.58)	71.4 (0.77)	95.7 (0.09)	66.4 (1.48)	85.8 (2.64)	92.9 (0.38)	94.6 (0.33)	82.6 (0.16)	73.5 (0.41)	96.5 (0.10)	85.02
	L-DOG	87.6 (1.15)	84.6 (0.38)	90.8 (0.38)	75.0 (0.44)	96.0 (0.05)	69.1 (1.30)	90.2 (2.02)	92.4 (0.36)	95.6 (0.46)	82.2 (0.15)	74.6 (0.32)	97.4 (0.08)	86.29
ViT-B/32	D-Adapt (Adam)	87.6 (0.48)	91.8 (0.31)	89.0 (0.42)	70.7 (0.48)	7.6 (0.00)	69.4 (1.55)	94.7 (0.92)	86.6 (1.42)	95.5 (0.32)	80.1 (0.09)	76.2 (0.38)	97.9 (0.05)	78.92
	D-Adapt (SGD)	88.7 (0.58)	91.8 (0.20)	84.2 (2.39)	70.5 (0.17)	44.0 (31.87)	74.9 (0.38)	95.2 (0.43)	91.5 (0.72)	96.3 (0.06)	79.8 (0.05)	76.3 (0.13)	97.7 (0.03)	82.82
	SPS	89.7 (0.51)	91.4 (0.18)	84.6 (1.18)	71.3 (0.27)	7.6 (0.00)	74.4 (0.96)	98.4 (0.33)	91.1 (0.49)	96.2 (0.17)	79.9 (0.09)	77.5 (0.46)	98.0 (0.04)	80.01
	DOG	89.5 (1.26)	92.5 (0.22)	85.0 (0.27)	69.5 (0.11)	67.7 (36.66)	75.5 (0.71)	98.9 (0.25)	92.4 (0.16)	96.4 (0.10)	79.7 (0.01)	77.8 (0.13)	97.7 (0.08)	85.22
	L-DOG	89.6 (0.81)	92.8 (0.15)	86.0 (0.40)	70.7 (0.29)	95.3 (0.08)	75.8 (0.71)	99.0 (0.26)	92.3 (0.45)	96.5 (0.17)	79.8 (0.07)	78.3 (0.26)	97.8 (0.04)	87.82